

Distributed Maximal Matching and Maximal Independent Set on Hypergraphs

Alkida Balliu* Sebastian Brandt† Fabian Kuhn‡ Dennis Olivetti§

Abstract

We investigate the distributed complexity of maximal matching and maximal independent set (MIS) in hypergraphs in the LOCAL model. A maximal matching of a hypergraph $H = (V_H, E_H)$ is a maximal disjoint set $M \subseteq E_H$ of hyperedges and an MIS $S \subseteq V_H$ is a maximal set of nodes such that no hyperedge is fully contained in S . Both problems can be solved by a simple sequential greedy algorithm, which can be implemented naïvely in $O(\Delta r + \log^* n)$ rounds, where Δ is the maximum degree, r is the rank, and n is the number of nodes of the hypergraph.

We show that for maximal matching, this naïve algorithm is optimal in the following sense. Any deterministic algorithm for solving the problem requires $\Omega(\min\{\Delta r, \log_{\Delta r} n\})$ rounds, and any randomized one requires $\Omega(\min\{\Delta r, \log_{\Delta r} n\})$ rounds. Hence, for any algorithm with a complexity of the form $O(f(\Delta, r) + g(n))$, we have $f(\Delta, r) \in \Omega(\Delta r)$ if $g(n)$ is not too large, and in particular if $g(n) = \log^* n$ (which is the optimal asymptotic dependency on n due to Linial's lower bound [FOCS'87]). Our lower bound proof is based on the round elimination framework, and its structure is inspired by a new round elimination fixed point that we give for the Δ -vertex coloring problem in hypergraphs, where nodes need to be colored such that there are no monochromatic hyperedges.

For the MIS problem on hypergraphs, we show that for $\Delta \ll r$, there are significant improvements over the naïve $O(\Delta r + \log^* n)$ -round algorithm. We give two deterministic algorithms for the problem. We show that a hypergraph MIS can be computed in $O(\Delta^2 \cdot \log r + \Delta \cdot \log r \cdot \log^* r + \log^* n)$ rounds. We further show that at the cost of a much worse dependency on Δ , the dependency on r can be removed almost entirely, by giving an algorithm with round complexity $\Delta^{O(\Delta)} \cdot \log^* r + O(\log^* n)$.

1 Introduction and Related Work

In the area of distributed graph algorithms, we have a network that is represented by a graph $G = (V, E)$, where nodes represent machines and edges represent communication links, and the goal is to solve some graph problem on the graph representing the network. The nodes of G can communicate if they are neighbors, and in the LOCAL model of distributed computing the computation proceeds in synchronous rounds and messages are allowed to be arbitrarily large. The complexity of a problem is the minimum number of rounds required, in the worst case, to solve it.

Given a graph $G = (V, E)$, a *maximal independent set* (MIS) of G is an inclusion-wise maximal set $S \subseteq V$ of nodes such that no two neighbors are in S and a *maximal matching* of G is an inclusion-wise maximal set $M \subseteq E$ of edges such that no two adjacent edges are in M . Maximal matching and MIS are two of the classic local symmetry breaking problems considered in the area of distributed graph algorithms.

The distributed complexity of the two problems has been studied in a long line of research, see, e.g., [2, 57, 43, 54, 3, 63, 33, 34, 62, 48, 46, 52, 12, 11, 28, 49, 25, 5, 64, 18, 8, 29, 6, 7]. Researchers have investigated the computational time complexity of these two problems in two different ways: (i) as a function of the maximum degree Δ of the graph, studying complexities of the form $f(\Delta) + g(n)$, where g is a much more slowly growing function than f ; (ii) as a function of the total number of nodes n in the graph, studying complexities of the form $f(n)$. The best known randomized distributed algorithms for MIS and maximal matching run in $O(\log \Delta + \log^5 \log n)$ rounds [28, 29] and $O(\log \Delta + \log^3 \log n)$ rounds [12, 25], respectively. As a function of n , the fastest known deterministic algorithms for MIS and maximal matching have time complexities $O(\log^5 n)$ [29] and $O(\log^2 \Delta \cdot \log n) = O(\log^3 n)$ [25], respectively. For both problems, there is a deterministic algorithm with a time complexity of $O(\Delta + \log^* n)$ [62, 11]. This algorithm is optimal as a function of n , as we know that it is not possible

*Gran Sasso Science Institute.

†CISPA Helmholtz Center for Information Security

‡University of Freiburg.

§Gran Sasso Science Institute.

to solve any of these problems in $f(\Delta) + o(\log^* n)$ rounds for any function f [55, 60]. Further, for both problems, there are randomized time lower bounds of $\Omega\left(\min\left\{\frac{\log \Delta}{\log \log \Delta}, \sqrt{\frac{\log n}{\log \log n}}\right\}\right)$ [48, 9] and $\Omega\left(\min\left\{\Delta, \frac{\log \log n}{\log \log \log n}\right\}\right)$ [5, 7], which hold even on tree topologies. Hence, although there are certainly some interesting remaining questions, the distributed complexity of MIS and maximal matching in graphs is relatively well understood.

Generalization to Hypergraphs. In this paper, we consider the distributed complexity of the natural generalization of MIS and maximal matchings to hypergraphs. A hypergraph $H = (V_H, E_H)$ consists of a set of nodes V_H and a set of hyperedges $E_H \subseteq 2^{V_H}$. We say that a hyperedge $e \in E_H$ is incident to a node $v \in V_H$ if $v \in e$. Similarly, a node v is incident to a hyperedge e if $v \in e$. We will omit the subscript “ H ” from the notation if H is clear from the context. The degree $\deg_H(v)$ of a node v in H is its number of incident hyperedges, while the rank $\text{rank}_H(e)$ of an edge $e \in E_H$ is the number of its incident nodes (i.e., $\text{rank}_H(e) = |e|$). The rank r_H of H is the maximum hyperedge cardinality, i.e., $r_H = \max_{e \in E_H} |e|$. We denote with Δ_H the maximum degree in H . An independent set $S \subseteq V_H$ of a hypergraph is a set of nodes such that no hyperedge is fully contained in S , i.e., $\forall e \in E_H : e \cap S \neq e$. A matching $M \subseteq E_H$ is a set of edges such that every node $v \in V_H$ is contained in at most one of the sets in M . Maximal independent sets and maximal matchings are then again inclusion-wise maximal sets with the respective property.

Generally, while graphs model pairwise dependencies or interactions, hypergraphs model dependencies or interactions between more than 2 nodes or entities. Graphs are therefore natural to model the pairwise communication in traditional wired networks. However, in general, interactions between the participants of a distributed system can certainly be more complex. Hypergraphs have a richer structure and are generally less understood than graphs (not only in the distributed context). Various problems that can naturally be formulated as problems on hypergraphs have been studied and found applications in the context of distributed graph algorithms (e.g., [24, 51, 56, 31, 26, 17]). In this context, one can also mention the beautiful line of work that uses methods from algebraic topology to prove distributed impossibility results and also develop distributed algorithms, e.g., [40, 65, 20, 39, 19]. There, simplicial complexes (which from a combinatorial point of view are hypergraphs with the property that the set of hyperedges forms a downward-closed family of subsets of nodes) are used to express the development of the state of a distributed system throughout the execution of a distributed algorithm. We hope that a better understanding of the fundamental limitations and possibilities of distributed algorithms in hypergraphs will lead to insights that prove useful in the design of future distributed systems. In the following, we discuss some concrete reasons why hypergraph MIS and maximal matching specifically are interesting and worthwhile objects to being studied from a distributed algorithms point of view.

Distributed MIS in Hypergraphs. Assume that we have a shared resource so that in each local neighborhood in the network, only a limited number of nodes can concurrently access the resource. The shared resource could for example be the common communication channel in a wireless setting that allows some network-level coding. Maybe a node can still decode a received linear combination of messages, as long as the signal consists at most κ messages for some $\kappa > 1$. The possible sets of nodes that can access the channel concurrently can then be expressed by the set of independent sets of some hypergraph. A further direct application of the MIS problem in hypergraphs is the problem of computing a minimal set cover and the problems of computing minimal dominating sets with certain properties [38]. Given a set cover instance, we can define a hypergraph H in the natural way by creating a node for each set and hyperedge for each element consisting of the sets that contain this element. The vertex covers of H correspond to the set covers of our set cover instance and the set of minimal vertex covers (and thus the set of minimal set covers) are exactly the complements of the set of maximal independent sets of H . To see this, note that clearly the component of a vertex cover (at least one node per hyperedge) is an independent set (not all nodes per hyperedge) and vice versa. Clearly minimality of a vertex cover implies maximality of the independent set and vice versa. In [38, 51], it is further shown that the minimal set cover problem can be used as a subroutine to compute sparse minimal dominating sets¹ and minimal connected dominating sets.

From a theoretical point of view, the distributed MIS problem in hypergraphs is also interesting because many of the techniques that work efficiently in the graph setting seem to fail or become much less efficient. It is for example not even clear how to obtain a randomized hypergraph MIS algorithm that is similarly efficient and simple as Luby’s algorithm [2, 57] for graphs. To the best of our knowledge, the first paper to explicitly study the

¹A minimal dominating set is called sparse if the average degree of the nodes in the dominating set is close to the average degree of the graph.

problem of computing an MIS in a hypergraph is by Kutten, Nanongkai, Pandurangan, and Robinson [51]. In the LOCAL model (i.e., with arbitrarily large messages), it is relatively straightforward to compute an MIS by using a network decomposition. By using the network decomposition algorithm of [53], one in this way obtains a randomized $O(\log^2 n)$ -round algorithm and by using the recent deterministic network decomposition of [64, 29], one obtains a deterministic $O(\log^5 n)$ -round algorithm. The focus of [51] was therefore to obtain efficient algorithms in the CONGEST model, i.e., algorithms that only use small messages.² In [51], it is shown that an MIS in a hypergraph of maximum degree Δ and rank r can be computed in the CONGEST model in time $O(\log^{(r+4)!+4} n)$ if r is small and in time $\Delta^\varepsilon \cdot \text{poly} \log n$ for any constant $\varepsilon > 0$. This upper bound has been later improved to $(\log n)^{2^{r+3}+O(1)}$ in [36]. Distributed CONGEST algorithm for a closely related problem have also been studied in [50]. All the mentioned bounds are obtained by randomized algorithms. In the present paper, we focus on deterministic algorithms and we focus on the dependency on Δ and r , while keeping the dependency on n to $O(\log^* n)$ and thus as small as possible (due to Linial's $\Omega(\log^* n)$ lower bound [55]). While there is not much work on distributed algorithms for computing an MIS in a hypergraph, the problem has been studied from a parallel algorithms point of view, see, e.g. [44, 13, 45, 58, 14, 36].

Distributed Maximal Matching in Hypergraphs. First note that the set of matchings on a hypergraph $H = (V_H, E_H)$ is equal to the set of independent sets on the line graph of H , that is, on the graph containing a node for every hyperedge in E_H and an edge for every intersecting pair of hyperedges in E_H . The maximal matching problem on hypergraphs of maximum degree Δ and rank r is therefore a special case of the MIS problem on graphs of maximum degree at most $r(\Delta - 1)$. Maximal matchings in hypergraphs have several applications as subroutines for solving standard problems in graphs, as we discuss next.

The maximum matching problem in graphs can be approximated arbitrarily well by using a classic approach of Hopcroft and Karp [41]. When starting with some matching and augmenting along a maximal set of disjoint shortest augmenting paths, one obtains a matching for which the shortest augmenting path length is strictly larger. As soon as the shortest augmenting path length is at least $2/\varepsilon$, the matching is guaranteed to be within a factor $(1 - \varepsilon)$ of an optimal matching. The step of finding a maximal set of disjoint shortest augmenting paths can be directly interpreted as a maximal matching problem in the hypergraph defined by the set of shortest augmenting paths. In the distributed context, this idea has been used, e.g., in [23, 56, 26, 30, 37]. In [42, 32], it is further shown how maximal hypergraph matching can also be used in a similar way to obtain fast distributed approximations for the weighted maximum matching problem. In [32], this was then used to deterministically compute an edge coloring with only $(1 + \varepsilon)\Delta$ colors. In [26], it was further shown that the problem of computing a $(2\Delta - 1)$ -edge coloring in graphs can be directly reduced to the problem of computing a maximal matching in a hypergraph of rank 3. Additionally, [26] also shows how to compute an edge orientation of out-degree $(1 + \varepsilon)$ times the arboricity of a given graph by reducing to maximal matching in low-rank hypergraphs. Finally, when viewing a hypergraph as a bipartite graph between nodes and hyperedges, the maximal hypergraph matching problem is related to the MIS problem on the square of this bipartite graph and understanding the complexity of maximal matching in hypergraphs can be a step towards understanding the complexity of MIS in G^2 .

Since the maximal hypergraph matching problem is a special case of the MIS problem in graphs, there are quite efficient randomized algorithms for the problem and the focus in the literature therefore so far has been on developing deterministic distributed algorithms for maximal matchings in hypergraphs [26, 30, 37]. Prior to the new efficient deterministic network decomposition algorithm of Rozhoň and Ghaffari [64], those papers lead to the first deterministic distributed polylogarithmic-time $(2\Delta - 1)$ -edge coloring algorithms. The best known deterministic distributed algorithm for computing a maximal matching in a hypergraph of maximum degree Δ and rank r is due to Harris [37] and it has a time complexity of $\tilde{O}((r^2 \log \Delta + r \log^2 \Delta) \cdot \log n)$, where $\tilde{O}(x)$ hides polylogarithmic factors in x .

The Trivial Algorithm. Both an MIS and a maximal matching of an n -node hypergraph $H = (V_H, E_H)$ with maximum degree Δ and rank r can be computed in time $O(\Delta r + \log^* n)$ in a straightforward way. The trivial algorithm for hypergraph MIS proceeds as follows. We first color the nodes in V_H with $\Delta(r - 1) + 1$ colors such that no two nodes that share a hyperedge are colored with the same color. Such a coloring can be computed in $O(\Delta r + \log^* n)$ rounds, e.g., by using the algorithm³ of [11]. Then, we iterate through the $O(\Delta r)$ color classes and

²Note that there is more than one way in which one can define the CONGEST model for hypergraphs. As we do not focus on CONGEST algorithms in this paper, we refer to [51] for a discussion of this issue.

³We could achieve a smaller runtime of $O(\sqrt{\Delta r \log(\Delta r)} + \log^* n)$ rounds for computing the coloring by using a state-of-the-art distributed (*degree* + 1)-coloring algorithm [27, 10, 59]; however, as, subsequently, we iterate through the obtained color classes, this

greedily add nodes to the solution if the addition does not violate the maximality condition. The trivial algorithm for hypergraph maximal matching proceeds analogously, where we color hyperedges instead of nodes, and there is a different maximality condition.

Some Notation. Before stating our contributions, we briefly discuss some terminology regarding hypergraphs. A hypergraph is called r -uniform if all its hyperedges are of cardinality exactly r and it is called Δ -regular if all nodes have degree Δ . A hypergraph is called *linear* if any two hyperedges intersect in at most one node. Further, a *hypertree* is a connected hypergraph $H = (V_H, E_H)$ such that there exists an underlying tree T on the nodes V_H for which every hyperedge consists of the nodes of a connected subtree of T . Finally, there is a natural representation of a hypergraph $H = (V_H, E_H)$ as a *bipartite graph* consisting of the nodes $V_H \cup E_H$ and an edge between $v \in V_H$ and $e \in E_H$ if and only if $v \in e$. We refer to this bipartite graph as the *bipartite representation of H* . Note that H is a linear hypertree if and only if its bipartite representation is a tree. We say that a linear hypertree is Δ -regular if every node is either of degree Δ or of degree 1.

1.1 Our Contributions While there is a clear relation between *graph MIS* and *hypergraph maximal matching*, or between *graph MIS* and *graph maximal matching* (it is possible to use an algorithm for the first problem to solve the second ones, as discussed earlier), observe that hypergraph maximal matching and hypergraph MIS cannot be easily compared: in contrast to the situation that we have in graphs, where it is possible to use an algorithm for MIS to solve maximal matching in the same asymptotic runtime, on hypergraphs no reduction of this kind is known. In fact, a priori it is not clear at all which problem is the easier of the two.

In order to better understand our results, it is useful to compare the complexities of hypergraph MIS and hypergraph maximal matching with the complexity of the trivial algorithm that solves these two problems. As discussed above, the trivial algorithm has a runtime of $O(\Delta r + \log^* n)$ rounds, i.e., it has a very low dependency on n , but is quite slow in terms of Δ and r . The question that we study in our work is whether the trivial algorithm is optimal or whether it is possible to improve on it.

Question 1. Is it possible to solve hypergraph maximal matching and/or hypergraph MIS in $O(\log^* n) + o(\Delta r)$ rounds?

Informally, we show the following.

- For hypergraph maximal matching, the trivial algorithm is best possible, unless we spend a much higher dependency on n .
- For hypergraph MIS, it is possible to improve on the trivial algorithm when $r \gg \Delta$.

In the following, we will discuss our results and their relation to the trivial algorithm more formally.

Maximal Matching Lower Bound. As our main result, we show that for the hypergraph maximal matching problem, the trivial algorithm—which is simply a naïve implementation of the sequential greedy algorithm—is best possible.

THEOREM 1.1. *Assume that $\Delta \geq 2$ and $r \geq 2$. Then any deterministic distributed algorithm in the LOCAL model for computing a maximal matching in hypergraphs with maximum degree Δ , rank r , and n nodes requires $\Omega(\min\{\Delta r, \log_{\Delta r} n\})$ rounds.⁴ Any randomized such algorithm requires at least $\Omega(\min\{\Delta r, \log_{\Delta r} \log n\})$ rounds. Moreover, our lower bounds hold already on Δ -regular r -uniform linear hypertrees.*

We remark that, while in general hypergraphs the number of hyperedges can be much larger than the number of nodes, in linear hypertrees (which is the case where our lower bound applies) the sum of the nodes and the hyperedges is linear in n .

Theorem 1.1 shows that, for hypergraph maximal matching, we answer Question 1 negatively: the theorem essentially states that we need to spend $\Omega(\Delta r)$ rounds, unless we spend a much higher dependency on n than

⁴would not change the overall asymptotic runtime.

⁴As lower bounds in several parameters are sometimes difficult to understand, we would also like to provide, as an example, the exact quantification of this statement, which is as follows: There is a constant c such that, for any $\Delta \geq 2$ and $r \geq 2$, and any deterministic algorithm \mathcal{A} , there are infinitely many hypergraphs with maximum degree Δ , rank r , and n nodes on which \mathcal{A} has a runtime of at least $c(\min\{\Delta r, \log_{\Delta r} n\})$ rounds.

the one of the trivial algorithm, i.e., $\Omega(\log_{\Delta r} n)$ for deterministic algorithms and $\Omega(\log_{\Delta r} \log n)$ for randomized ones. In other words, the trivial greedy algorithm is optimal if we insist on having a small dependency on n . Observe that there exist algorithms for hypergraph maximal matching that can beat the $\Omega(\Delta r)$ lower bound that we provide, but they all spend substantially more than \log^* as a function of n , and our results show that this is indeed necessary. For example, the algorithm by Harris [37] has a dependency on Δ that is only polylogarithmic, but it also has a multiplicative $\log n$ dependency.

In fact, as a corollary, we obtain that the complexity of any algorithm for hypergraph maximal matching, when expressed solely as a function of n , must be large. This is known by prior work for $r = 2$, and we get that the same holds even when $r = \Delta$. This can be observed by setting $r = \Delta = \sqrt{\log n / \log \log n}$ in the deterministic lower bound and $r = \Delta = \sqrt{\log \log n / \log \log \log n}$ in the randomized one.

COROLLARY 1.1. *Any deterministic distributed algorithm in the LOCAL model for computing a maximal matching in n -node hypergraphs requires at least $\Omega\left(\frac{\log n}{\log \log n}\right)$ rounds. Any randomized such algorithm requires at least $\Omega\left(\frac{\log \log n}{\log \log \log n}\right)$ rounds. Moreover, our lower bounds hold already on regular uniform linear n -node hypertrees that satisfy $r = \Delta$.*

Hypergraph Coloring. On a high level, the proof of Theorem 1.1 is based on a lower bound proof that we provide for a seemingly unrelated problem, namely, hypergraph coloring. In fact, as a byproduct, we also obtain lower bounds for variants of this problem. There are two natural ways to generalize the standard coloring problem to hypergraphs:

- Color the nodes of the hypergraph with c colors such that, for every hyperedge of rank at least 2, it holds that at least 2 incident nodes have different colors. This is the standard definition of hypergraph coloring.
- Color the nodes of the hypergraph with c colors such that, for every hyperedge, all incident nodes have different colors. This variant is sometimes called *strong coloring* [1], *colorful coloring*, or *rainbow coloring*⁵ [21]. We will refer to this variant as *colorful*.

For the first problem, we show that it cannot be solved fast when $c \leq \Delta$, and for the second one, we show that it cannot be solved fast when $c \leq \Delta(r - 1)$. While these results can also be shown with non-trivial reductions from the known hardness of a different problem called sinkless orientation, the proof that we present for these lower bounds is essential as a building block for the proof of our hypergraph maximal matching lower bound (and, in fact, can be seen as a simplified version of that proof).

THEOREM 1.2. *In the LOCAL model, the Δ -hypergraph coloring and the $(r - 1)\Delta$ -hypergraph colorful coloring require $\Omega(\log_{\Delta r} n)$ rounds for deterministic algorithms and $\Omega(\log_{\Delta r} \log n)$ rounds for randomized ones.*

MIS Upper Bounds. In graphs, it is known that both maximal matching and MIS require a number of rounds that is linear in Δ [5, 7], unless we spend a high dependency on n . Hence, a natural question to ask is whether, like in the case of hypergraph maximal matching, also hypergraph MIS requires $\Omega(\Delta r)$ rounds if we keep the dependency on n as small as possible.

We show that hypergraph MIS behaves differently than hypergraph maximal matching: in fact, Question 1 can be answered affirmatively for a certain parameter range. In particular, we consider the case where $r \gg \Delta$ and design algorithms with a round complexity of the form $O(\log^* n) + f(\Delta, r)$ where $f(\Delta, r) = o(\Delta r)$ if Δ is sufficiently small compared to r . More concretely, we prove the following two results.

THEOREM 1.3. *In the LOCAL model, the MIS problem on n -node hypergraphs of maximum degree Δ and rank r can be solved in $O(\Delta^2 \log r + \Delta \log r \log^* r + \log^* n)$ deterministic rounds.*

THEOREM 1.4. *In the LOCAL model, the MIS problem on n -node hypergraphs of maximum degree Δ and rank r can be solved in $2^{O(\Delta \log \Delta)} \log^* r + O(\log^* n)$ deterministic rounds.*

⁵In the literature, the term rainbow coloring is also used to refer to a different, harder variant of coloring.

Note that the bound in the second theorem is almost independent of the rank r . In particular, Theorem 1.4 implies that for bounded-degree hypergraphs, an MIS can be computed in time $O(\log^* n)$, even if the rank r is not bounded. Observe that, while the dependency on Δ is high, by prior work we know that it has to be at least linear [5, 7].

We further note that the bounds that we provide are incomparable to the upper bounds that were obtained in [51], where the authors tried to optimize the round complexity as a function of n alone. The runtimes of the algorithms in [51] are all at least polylogarithmic in n . In our algorithms, we keep the dependency on n to $O(\log^* n)$ and thus as small as it can be. Our algorithms are faster than the poly log n -time algorithms based on network decomposition in the realistic scenario where Δ and r are much smaller than n .

Open Questions. Our work raises the following natural open question. Is it possible to solve hypergraph MIS in just $O(\Delta + \log^* n)$ rounds, independently of r ? Given the mentioned lower bounds on graphs, such a complexity would be optimal. In this regard, we show that there is a variant of hypergraph coloring that, if solved fast, would provide such an algorithm. This variant is typically called c -unique-maximum coloring [21], and requires to color the nodes of a hypergraph with c colors such that, for each hyperedge e , the maximum color appearing at the nodes incident to e occurs exactly once among the incident nodes. We show that finding an $O(\Delta)$ -unique-maximum coloring in $O(\Delta + \log^* n)$ rounds would allow us to solve hypergraph MIS optimally. While our lower bounds hold for different variants of hypergraph coloring, they do not hold for this variant.

Technical Remarks. In order to prove our lower bounds, we make use of the round elimination framework, which provides a general outline for proving lower bounds in the distributed setting. Round elimination has first been applied in [16] to obtain lower bounds on computing a sinkless orientation and a Δ -coloring of a graph. In 2019, [15] showed that the round elimination framework can be applied to almost any locally checkable problem⁶, at least *in principle*; however to obtain a good lower bound for some concrete locally checkable problem, it is necessary to overcome a certain set of challenges that is individual for each problem or problem family. While for some problems—such as $(\Delta + 1)$ -vertex coloring or $(2\Delta - 1)$ -edge coloring—these challenges seem far beyond the reach of current techniques, they have been overcome for a number of other important problems [5, 18, 8, 6, 7, 4]. Each of these results has been accompanied by a new key idea that succeeded in making the round elimination technique applicable also *in practice* for the considered problem(s); often the key idea unlocked the applicability of round elimination not only for a single problem, but for several problems at once, a whole problem family, or a new setting. In a similar vein, we prove the first substantial lower bounds for important hypergraph problems (that do not simply follow from lower bounds on graphs) by showing how to make these problems susceptible to a round elimination type of approach. We hope that our techniques can serve as a starting point for unlocking the round elimination technique more generally for hypergraph problems.

1.2 Road Map We start, in Section 2, by providing a high-level overview of our lower bound proof, highlighting the challenges and the new ingredients that we bring to the table.

We then proceed, in Section 3, by first defining the model of computing, and then by formally describing the round elimination technique. In this section, we also provide an example of the application of this technique.

In Section 4, we prove lower bounds for hypergraph coloring and some of its variants. Hypergraph coloring seems to be, at first sight, unrelated to hypergraph maximal matching. We show the connections between these two problems in Section 2, where we also explain why the proof that we present for the hardness of hypergraph coloring can be seen as a simplified version of the proof for hypergraph maximal matching.

The hypergraph maximal matching lower bound is presented in Section 5. The complexity of this proof is higher than the one for coloring, and requires some heavy notation. We will start by describing how it is connected to the proof for hypergraph coloring, by informally describing how it can be seen as an extension of the one for hypergraph coloring.

In Section 6, we present upper bounds for the hypergraph MIS problem. We show that, while the trivial algorithm is able to solve the problem in $O(\Delta r + \log^* n)$ rounds, if we allow a larger dependency on Δ we can almost entirely remove the dependency on r . In fact, in hypergraphs of maximum degree $\Delta = O(1)$ and very large rank r , we show that it is possible to solve the hypergraph MIS problem in just $O(\log^* n)$ rounds.

⁶Roughly speaking, a locally checkable problem is a problem that can be defined via local constraints. A standard example of locally checkable problems are proper coloring problems, which can be described by the local constraint that the two endpoints of any edge must receive different colors.

We conclude, in Section 7, by stating some open questions. We show that hypergraph MIS could be solved optimally if we are provided with an $O(\Delta)$ -unique-maximum coloring, and we leave as an open question determining the complexity of this variant of coloring.

2 High-Level Overview of the Lower Bound

Due to the fact that round elimination lower bound proofs are typically technically complex and notation-heavy, which hides the underlying conceptual ideas, we will give a detailed explanation of the round elimination framework, its fundamental issue, approaches of previous work to take care of this issue and our conceptual contributions in the following. While we have to resort to painting a very rough picture in places (as giving all details of the previous approaches would be far beyond the scope of this work), we hope that our high-level overview will make the recent developments regarding distributed lower bounds accessible to a broader audience while highlighting how our approach extends the lower bound tool box.

The Round Elimination Technique. On a high level, round elimination is a technique that can be used to prove lower bounds on the time required to solve (on trees or hypertrees) locally checkable problems, i.e., problems that can be defined by specifying some constraints that each node and hyperedge must satisfy. In the round elimination framework, a problem Π is defined by listing what labelings, or *configurations*, are allowed around the nodes and on the hyperedges. More precisely, each node v has to output some output label for each pair (v, e) such that e is a hyperedge containing v ; the output is considered to be correct if for each node, the configuration of output labels on the pairs with first entry v is contained in an explicitly given list of allowed “node configurations”, and for each hyperedge e , the configuration of output labels on the pairs with second entry e is contained in an explicitly given list of allowed “hyperedge configurations”.

The round elimination technique is a mechanical procedure \mathcal{R} that can be applied to Π to obtain a new problem $\Pi' = \mathcal{R}(\Pi)$ that, under mild assumptions regarding the considered graph class, is exactly one round easier than Π . Suppose we prove that $\Pi' = \mathcal{R}(\Pi)$ cannot be solved in 0 rounds of communication: we obtain that solving Π' requires at least 1 round, and hence that solving Π requires at least 2 rounds. The obtained problem Π' is also always locally checkable, and hence we can apply this technique again to obtain some problem Π'' guaranteed to be exactly one round easier than Π' . If we prove that also this problem cannot be solved in 0 rounds, we obtain a lower bound of 3 rounds for Π .

The Issue with Round Elimination. While the high-level outline given above suggests that obtaining lower bounds via round elimination is an entirely mechanical task, this is unfortunately not the case due to a fundamental issue in the framework. A simple example to explain this issue is the following. Let Π be the problem of 3-coloring a cycle, which is clearly locally checkable. The constraints of this problem can be specified by a list of allowed configurations using 3 output labels (the colors). If we apply round elimination to this problem, we automatically obtain a new problem $\Pi' = \mathcal{R}(\Pi)$ that can be described by using 16 labels. If we apply round elimination again, we get a problem with thousands of labels. While in theory we could continue forever, in practice we cannot, because it is not feasible to compute the next problem (and even less to determine whether it can be solved in 0 rounds or not). In general, the downside of this technique is that the description of Π' can be doubly exponentially larger than the one of Π .

Decomposing a Problem. In [7], a technique has been presented that sometimes allows us to take care of this issue. In the following, we describe this technique for the case of the MIS problem (on graphs). Let $\mathcal{R}^0(\Pi) = \Pi$ and $\mathcal{R}^{i+1}(\Pi) = \mathcal{R}(\mathcal{R}^i(\Pi))$.

It has been observed already in [8] that if we start from Π equal to the MIS problem and apply the round elimination technique iteratively k times, the result $\mathcal{R}^k(\Pi)$ can be essentially decomposed into three parts:

1. the *original* problem, i.e., MIS,
2. a *natural* part of small size, corresponding to the k -coloring problem, and
3. an *unnatural* part whose description is of size roughly equal to a power tower of height k and that cannot be easily understood.

MIS is not the only problem that behaves in this way under round elimination; in fact, many problems admit, after applying round elimination for k steps, a decomposition into the original problem, some natural part that is easy to understand (and often unrelated to the original problem), and some unnatural part of much larger size (where the latter two depend on k).

The main idea of the previous works that use the round elimination technique is to relax the problems obtained at each step with the goal of replacing the unnatural part with something much smaller. In this context, relaxing a problem means adding further configurations to the lists collecting the allowed configurations for the nodes or the hyperedges, at the cost of making the problem potentially easier to solve. Observe that, if a problem is relaxed too much, then it is not possible to obtain strong lower bounds, because we would soon reach a problem that is 0-round solvable. Hence, the typical goal is to find relaxations that make the description of the problem much smaller while still not making the problem much easier. Note that, effectively, the description of a problem can indeed become smaller by *adding* allowed configurations: by allowing more (in a suitable way), a substantial number of allowed configurations can be shown to become useless and hence can be ignored.

Previous Approaches. In order to prove tight lower bounds for ruling sets (and other interesting problems), in [7], the third part in the decomposition of the problems has been handled as follows.

Consider the k -coloring problem, for $k \leq \Delta$. In some sense, this problem is known to be much harder than MIS, since, unlike MIS, it cannot be solved in $O(f(\Delta) + \log^* n)$ rounds for any function f . We call problems that cannot be solved in $O(f(\Delta) + \log^* n)$ rounds for any function f *hard*. If we apply round elimination to the k -coloring problem, we essentially get something very similar to the third part presented above. But for some of the hard problems, we have techniques that allow us to deal with them, that is, finding relaxations of the problem that result in a so-called *fixed point*. This may look like a contradictory statement: it is not possible for a problem to be exactly one round easier than itself. The reason why this is not an actual contradiction is that the condition that makes the round elimination statement hold stops working when the complexity is $\Omega(\log_{\Delta} n)$. In fact, it is possible to prove that a fixed point directly implies a lower bound of $\Omega(\log_{\Delta} n)$ for the problem.

The point of relaxing the k -coloring problem to a fixed point before applying the round elimination technique is the following: if we apply round elimination to the k -coloring problem itself, we obtain a lot of configurations that make the obtained problem hard to understand, while if we apply round elimination to a fixed point relaxation of the k -coloring problem, we obtain the problem itself, and no additional allowed configurations. Observe that, for problems solvable in $O(f(\Delta) + \log^* n)$ rounds, like MIS, fixed point relaxations cannot exist, since they would give a lower bound higher than the upper bound for small values of Δ .

The idea of [7] is to *embed* the fixed point of the k -coloring problem into the sequence obtained from MIS, in the following sense. We start from MIS, and we apply round elimination iteratively. At each step k , we relax the obtained problem by adding all the configurations allowed by the fixed point relaxation of the k -coloring problem. It has been shown in [7] that thereby we obtain a problem sequence in which the number of allowed configurations is always bounded by $O(2^k)$, and hence much better than a power tower of height k . In other words, the third part of the problem (the unnatural part) is replaced by the configurations that we need to add to the k -coloring problem to make it a fixed point, of which there are just $O(2^k)$. The reason why this approach works for MIS seems to be related to the fact that all the unnatural configurations present at step $k + 1$ are generated (by \mathcal{R}) by starting from the k -coloring configurations present at step k .

The take-home message of the approach of [7] is the following: try to decompose the problem into the three parts explained above, and try to find a fixed point relaxation for the natural part; it may then happen that the unnatural part just disappears.

By following this approach, the authors of [7] managed to prove tight lower bounds for MIS, ruling sets, and many other interesting natural problems. Also, they asked if there are other interesting problems for which this technique is applicable (see Open Problem 2 in [7]), and we answer this question affirmatively.

A Colorful Matching. We extend the applicability of the technique to hypergraphs by solving a number of intricate challenges introduced by the more complex setting. In particular, we show that while the explained technique cannot be applied directly, by combining it with a set of new ingredients we obtain substantial lower bounds for the hypergraph maximal matching problem. The good news from the perspective of the general outline is that we can still decompose $\mathcal{R}^k(\Pi)$, where Π is the hypergraph maximal matching problem, into three parts, by making use of k -hypergraph colorful coloring:

1. The *original* problem,
2. a *natural* part: the k -hypergraph colorful coloring problem, and
3. an *unnatural* part, of size roughly equal to a power tower of height k , without any easily discernible useful structure.

By following the approach of [7], in order to prove a lower bound of $\Omega(\Delta r)$, it is necessary to first find a fixed point relaxation for the k -hypergraph colorful coloring problem, for k that can be as large as $\Omega(\Delta r)$.

Challenges and New Ingredients. Two labels ℓ_1 and ℓ_2 are called *equivalent* when, for any allowed configuration containing ℓ_1 , if we replace ℓ_1 with ℓ_2 , we obtain a configuration that is allowed, and for any configuration containing ℓ_2 , if we replace ℓ_2 with ℓ_1 , we also obtain a configuration that is allowed. If two labels are equivalent, we can actually discard all configurations containing one of the two labels, still obtaining a problem that is not harder than the original one. Relaxing a problem in order to make ℓ_1 equivalent to ℓ_2 is called *label merging*, and it is a powerful technique that can be used to reduce the number of labels. In previous works, relaxations are typically done in two ways:

- When proving lower bounds for problems solvable in $O(f(\Delta) + \log^* n)$ rounds, it is typically the case that, by applying round elimination, many new labels are obtained, and to keep the problem small, a lot of label merging is performed.
- When proving lower bounds for hard problems, configurations (possibly containing new labels) are added in order to transform the problem into a fixed point. Label mergings are not performed.

Unfortunately, the latter outline for finding a fixed point relaxation does not seem to work for $\Delta(r-1)$ -hypergraph colorful coloring. Via a new approach, we nevertheless show how to find a relaxation of the $\Delta(r-1)$ -hypergraph colorful coloring problem that is a fixed point, but it comes with two issues: the fixed point does not resemble the $\Delta(r-1)$ -hypergraph colorful coloring problem at all, and the approach of [7], applied to it, does not work.

We obtain the fixed point as follows. We first find a fixed point relaxation for the Δ -hypergraph (non-colorful) coloring problem, and then show that there exists a 0-round algorithm that converts a $\Delta(r-1)$ -hypergraph colorful coloring into a proper Δ -hypergraph coloring. This algorithm can be interpreted in an interesting way: what it does is mapping different colors into a single one, essentially showing that the Δ -hypergraph (non-colorful) coloring problem can be obtained from $\Delta(r-1)$ -hypergraph colorful coloring by performing label merging. This observation allows us to obtain a fixed point for the $\Delta(r-1)$ -hypergraph colorful coloring problem: start from $\Delta(r-1)$ -hypergraph colorful coloring, merge some labels to obtain Δ -hypergraph coloring, then add configurations to this new problem in order to make it a fixed point. We provide a more detailed overview of our approach for obtaining the fixed point in Section 4.

We now turn to explaining the challenges and our solutions for obtaining the hypergraph maximal matching lower bound with the help of the aforementioned fixed point.

Recall the aforementioned decomposition of problems into the original problem, a natural colorful coloring part, and an unnatural part of large description size. In the hypergraph maximal matching problem, by applying round elimination, we essentially get an additional new color in the coloring part per step while the size of the unnatural part grows like a power tower. We want to be able to perform $\Omega(\Delta r)$ steps of round elimination before reaching a 0-round solvable problem, but we cannot let the number of labels grow like a power tower because it would then be infeasible to actually understand the problem sequence. In order to keep the number of labels reasonably small throughout the problem sequence without compromising the lower bound quality, we have to perform relaxations that do not make the problems too easy, as otherwise we would obtain a 0-round solvable problem after $o(\Delta r)$ steps, which would give a worse lower bound.

We would like to use the same approach as [7], i.e., after round elimination step k , we want to add the allowed configurations of the fixed point of the k -hypergraph colorful coloring problem to the obtained problem. It is actually possible to extend the approach of Section 4 to show a fixed point relaxation for the k -hypergraph colorful coloring problem. This fixed point is also a relaxation of the $\lceil k/(r-1) \rceil$ -hypergraph coloring problem. The problem is that this fixed point is not exactly designed with colorful coloring in mind, and this makes it hard to adapt it to the case where colors increase by one at each step. In particular, consider the case where we try to use, as the k -th problem in our relaxed sequence, the problem $\Pi(k)$ containing two parts: the original problem, and the fixed point relaxation of the k -hypergraph colorful coloring problem (as would be the case with the approach from [7]). Note that the second part would contain $\lceil k/(r-1) \rceil$ colors, and these colors would satisfy the constraints of the hypergraph non-colorful coloring problem. By performing round elimination, we obtain a new color, and it does not seem possible to relax the obtained problem to $\Pi(k+1)$. In particular, we would like to relax the coloring part of the obtained problem to the fixed point for $\lceil (k+1)/(r-1) \rceil$ -hypergraph coloring, but it seems that the coloring part of the obtained problem is already easier than that, and that it can only be relaxed to the

fixed point for $(\lceil k/(r-1) \rceil + 1)$ -hypergraph coloring. If instead we use the latter approach (i.e., the approach of relaxing the coloring part of the obtained problem to the fixed point with the smallest number of colors that the coloring part can be relaxed to), the number of colors in the coloring part grows by 1 at each step, meaning that after only $\Delta + 1$ steps of round elimination, we would have $\Delta + 1$ colors. However, the construction of the fixed point is only possible up to Δ colors, which renders also the modified approach infeasible (as already after $\Delta + 1$ steps we would obtain a problem that is 0-round solvable).

In order to make things work, our idea is to not use directly a fixed point for hypergraph coloring, but to build a problem sequence in which, if we take the coloring part of each problem, we obtain a problem sequence that (except for the few first steps) has always Δ colors, and in which the hyperedge constraint gets more relaxed at each step. This coloring part at some point will become actually equal to the Δ -coloring fixed point, but only after $\Omega(\Delta r)$ steps. All this is achieved by merging, at each step, the new color with an existing one, in such a way that we relax the hyperedge constraint as little as possible.

Unfortunately, the challenges mentioned above are just the conceptual issues that we have to take care of—there are a number of technical issues on top of the conceptual ones that make our lower bound proof fairly complicated. One of these technical issues is that even when following our new approach, we obtain a problem sequence that is too relaxed, and hence we cannot obtain a good lower bound in this way. However, if during the relaxation procedure we select one color and treat it slightly different than all the other colors (i.e., amongst all the configurations that we add during the relaxation procedure the ones that contain the selected color have a slightly different structure than those containing other colors), then the relaxations turn out to be sufficiently tight. The choice of the selected color is irrelevant, as long as we do not select the same color in two subsequent relaxation procedures. In general, the hypergraph setting increases the complexity of various crucial ingredients in our proof (such as constructing a suitable problem family, proving that our relaxations indeed lead to the stated problem family by setting up a suitable instance on which we can apply Hall’s marriage theorem, etc.) and therefore of the proof itself—we hope that our contributions (both the conceptual and the technical ones) will help navigating the difficulties of proving further lower bounds for important hypergraph problems. We provide a more detailed overview of our lower bound approach for hypergraph maximal matching in Section 5.

An Easier Proof. We would like to point out that in order to prove a lower bound for hypergraph maximal matching, we could have entirely skipped Section 4, i.e., the proof of the hypergraph coloring lower bound: it is not necessary to prove that the relaxation of the Δ -hypergraph coloring that we provide is indeed a fixed point, because it does not seem possible to make the approach of [7] modular. In other words, we cannot use the lower bound for Δ -hypergraph coloring as a building block for the lower bound proof for hypergraph maximal matching, but we have to prove the latter lower bound from scratch; however, we can make use of some of the structures used in the proof for Δ -hypergraph coloring. In fact, essentially, the lower bound proof for coloring is a much simplified version of the one for hypergraph maximal matching, which has to handle many additional challenges. For example, the number of colors in the aforementioned “natural part” of the problems in the sequence obtained by starting from hypergraph maximal matching grows by 1 at each step, while in the case of Δ -hypergraph coloring it never changes.

3 Definitions and Notation

The upper and lower bounds presented in this work hold on hypergraphs, which have been defined in Section 1. We now provide some additional notation.

We denote with $N_H(v)$ the neighbors of a node v , that is, all those nodes in H , different from v , contained in the hyperedges incident to v .

Each node of a hypergraph is at distance 0 from itself. Let w_0, \dots, w_d be the smallest sequence of nodes of H such that: $w_0 = u$, $w_d = v$, and $w_{i+1} \in N_H(w_i)$ for all $0 \leq i < d$. Then we say that nodes u and v are at distance d in H . We call the r -hop neighborhood of a node u the subhypergraph induced by nodes in H that are at distance at most r from u , where we remove the hyperedges that contain only nodes at distance exactly r from u .

We will omit the subscript “ H ” from the notation if H is clear from the context.

3.1 The LOCAL Model Our lower and upper bounds hold in the classic LOCAL model of distributed computing, introduced by Linial [55]. Given an n -node hypergraph H , each node is equipped with a unique identifier from $\{1, \dots, n^c\}$, for some constant $c \geq 1$. This is a message passing model, where the computation proceeds in synchronous rounds: at each round, nodes of a hypergraph exchange messages with the neighbors and perform

some local computation. In this model, both the local computational power of each node, and the size of the messages, are not bounded. A distributed algorithm that solves a graph problem in the LOCAL model runs at each node in parallel, and the goal is for nodes to produce their local outputs that together form a valid global solution of the desired problem. We say that an algorithm has round complexity T if each node decides its local output and terminates within T communication rounds.

In the randomized version of the LOCAL model, each node is equipped with a private random bit string, and in this context we consider Monte Carlo algorithms, that is, we require to produce a correct solution to the desired problem with high probability, that is, with probability at least $1 - 1/n$.

As typically done in this context, throughout this paper, we assume that each node v of a hypergraph H knows initially its own degree $\deg(v)$, the maximum degree Δ , the maximum rank r , and the total number n of nodes.

Notice that, since in the LOCAL model we do not restrict the bandwidth, a T -round LOCAL algorithm is equivalent to the following: first each node spends T rounds to collect its T -hop neighborhood in H , and then maps each T -hop neighborhood into an output. Sometimes, it is convenient to work on the incidence graph (or bipartite representation) of a hypergraph $H = (V_H, E_H)$, which is a bipartite graph $B = (U_B \cup W_B, E_B)$ where $U_B = V_H$, $W_B = E_H$, and there is an edge between $u \in U_B$ and $v \in W_B$ if and only if the hyperedge of H corresponding to v contains the node of H corresponding to u . Notice that, in the LOCAL model, any T -rounds algorithm that solves a problem in H can clearly be simulated in B in at most $2T$ communication rounds.

The PN Model. While both our upper bound and lower bound results hold in the LOCAL model, for technical reasons, the lower bounds are first shown on the *port numbering* (PN) model, which is weaker than the LOCAL one, and then we lift them for the randomized and deterministic LOCAL model.

The PN model is the same as the LOCAL one (synchronous message passing model, unbounded computational power, unbounded size of messages), with the difference that nodes do not have unique identifiers, and instead each node has an internal ordering of its incident hyperedges. More precisely, incident hyperedges of a node v have a pairwise distinct (port) number in $\{1, \dots, \deg(v)\}$ (assigned arbitrarily). For technical reasons, we use a slightly modified version of the PN model, where in addition we require that incident nodes of a hyperedge have a pairwise distinct number in $\{1, \dots, \text{rank}(e)\}$ (assigned arbitrarily).

3.2 The Automatic Round Elimination Framework Round elimination is a technique that can be used to prove lower bounds in the distributed setting. It has been first used to prove lower bounds for sinkless orientation and Δ -coloring [16]. In its current form, called *automatic* round elimination, it has been introduced by Brandt [15], and since then it has been proved to be extremely useful for proving lower bounds in the LOCAL model.

In this section, we describe the automatic round elimination framework. We start by describing how to encode a problem in this framework (introducing some notation as well), and what should satisfy an output in order to be correct—we illustrate these concepts by showing the concrete example of the encoding of the MIS problem on hypergraphs with maximum rank 2 (that is, on standard graphs).

Encoding of a Problem. For the purpose of showing our lower bounds, it is enough to consider only regular linear hypertrees, that is, hypergraphs such that their incidence graph is a bipartite 2-colored tree, where each white non-leaf node has degree Δ and each black node has degree r . Hence, we show how to encode problems in such case. In this formalism, a problem $\Pi_{\Delta,r}$ is described by a triple $(\Sigma_{\Pi}, \mathcal{N}_{\Pi}, \mathcal{E}_{\Pi})$, where:

- the *alphabet set* Σ_{Π} is a set of allowed labels,
- the *node constraint* \mathcal{N}_{Π} is a set of multisets of size Δ over the alphabet Σ_{Π} ,
- the *hyperedge constraint* \mathcal{E}_{Π} is a set of multisets of size r over the alphabet Σ_{Π} .

Sometimes, instead of *multisets*, we will use *words*. It is just a matter of convenience: writing a word is shorter than writing a multiset. But, the word will still represent a multiset, in the sense that a label can appear many times in a word and the order of appearance of the labels does not matter.

A word of length Δ (resp. r) over the alphabet Σ_{Π} is called *node configuration* (resp. *hyperedge configuration*), and it is *valid* or *allowed* or *it satisfies the constraint* if it is contained in \mathcal{N}_{Π} (resp. \mathcal{E}_{Π}).

The Output. Informally, in this formalism, in order to solve a problem Π on a regular linear hypertree $H = (V, E)$, each node must output a label from Σ_{Π} on each incident hyperedge. More precisely, each (node, incident hyperedge) pair of the set $\{(v, e) \mid v \in V, e \in E, v \in e\}$ must be labeled with an element from the alphabet

Σ_Π . We then say that an output is correct if it satisfies \mathcal{N}_Π and \mathcal{E}_Π . More precisely, let $v \in V$ be a node and let I_v be the set of all hyperedges $e \in E$ such that $v \in e$. Then the word of size Δ described by all the labels given to (v, e) must be in \mathcal{N}_Π . Similarly, let $e \in E$ be a hyperedge and let I_e be the set of all nodes v such that $v \in e$. Then the word of size r described by all the labels given to (v, e) must be in \mathcal{E}_Π .

Example: Encoding of MIS. As an example, let us see how to encode the MIS problem on a Δ -regular graph $G = (V, E)$ in the round elimination framework. We show an example with MIS because, in the round elimination framework, MIS behaves similarly as hypergraph maximal matching. In other words, we want to define Σ , \mathcal{N} , and \mathcal{E} , such that any labeling that satisfies the node and edge constraint results in a maximal independent set, and any maximal independent set can be used to produce such a labeling. The alphabet set is $\Sigma = \{M, P, O\}$, where intuitively M is used to say that a node is in the MIS, P is used to point to a neighbor that is in the MIS, while O stands for “other” and is used to express that somehow it does not matter what happens in that part of the graph (for example, from the perspective of a node v , the label O on an incident half-edge (v, e) indicates that it does not matter whether the neighbor reached through e is in the MIS or not). The node and edge constraints are defined as follows.

\mathcal{N} :	\mathcal{E} :
M^Δ	PM
$PO^{\Delta-1}$	OM
	OO

Nodes in the MIS output the configuration M^Δ . Nodes not in the MIS output the configuration $PO^{\Delta-1}$, where P is used to guarantee the maximality of the solution by pointing to one neighbor that is in the MIS, while O is given on the other incident half-edges that are connect to nodes that may or may not be in the MIS. Since we do not allow two neighboring nodes to be in the MIS, then $MM \notin \mathcal{E}$. Since P must only point to a neighbor in the MIS then we have $PM \in \mathcal{E}$, while $PP, PO \notin \mathcal{E}$. Also, since a node not in the MIS may have more than one neighbor in the MIS, or other neighbors not in the MIS, then we get $OO \in \mathcal{E}$ and $OM \in \mathcal{E}$.

Please observe that the node constraint only specifies what needs to be satisfied by nodes of degree exactly Δ , and this means that leaves are unconstrained. Hence, these constraint do not describe the exact MIS problem, but a similar problem, where only non-leaf nodes need to actually solve MIS. Since we prove lower bounds, this is not an issue: a lower bound for this simplified variant of MIS would imply a lower bound for the standard MIS problem.

3.3 Automatic Round Elimination Technique We now present the automatic round elimination technique, and we dive into some details that are necessary and sufficient for understanding the technical parts of our lower bound proofs. The automatic version of the round elimination technique was introduced by Brandt [15], and later actually implemented by Olivetti [61]. On a high level, if we apply the round elimination technique on a problem Π with complexity T we obtain a (possibly) new problem Π' having complexity $\max\{0, T - 1\}$. Let Π_1 be our problem of interest for which we want to show a lower bound. By applying automatic round elimination iteratively we get a sequence of problems Π_2, Π_3, \dots such that each problem Π_i is at least one round easier than Π_{i-1} , assuming that Π_{i-1} is not already 0-rounds solvable. Hence, if we can apply round elimination for T times before reaching a 0-rounds solvable problem, we get a lower bound of T rounds for our problem of interest Π_1 .

If we take a more fine-grained look at this technique, what actually happens is that, when we apply round elimination on a problem Π , this procedure first gives an *intermediate* problem Π' , and only after applying round elimination on Π' we then get the problem Π'' that is at least one round easier than Π . Hence, let $\mathcal{R}(\cdot)$ be the function that takes in input a problem Π and outputs the intermediate problem Π' , and let $\overline{\mathcal{R}}(\cdot)$ be the function that takes in input an intermediate problem Π' and outputs the problem Π'' . We get that $\Pi' = \mathcal{R}(\Pi)$ and $\Pi'' = \overline{\mathcal{R}}(\Pi') = \overline{\mathcal{R}}(\mathcal{R}(\Pi))$.

Applying round elimination means to first compute $\Pi' = \mathcal{R}(\Pi)$ and then compute $\Pi'' = \overline{\mathcal{R}}(\Pi')$, and Π'' is the problem that is guaranteed to have complexity $\max\{0, T - 1\}$ if Π has complexity T . On a high level, the labels that are used to define Π' are *sets* of labels of Π , and the labels used to define Π'' are *sets* of labels of Π' . Observe that the configurations allowed by the constraints of Π' are still just multisets of labels, but the elements of these multisets are hence now *sets* of labels of Π .

We now present the definition of the function \mathcal{R} . Starting from a problem $\Pi = (\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi)$, the problem

$\mathcal{R}(\Pi = (\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi)) = \Pi' = (\Sigma_{\Pi'}, \mathcal{N}_{\Pi'}, \mathcal{E}_{\Pi'})$ is defined as follows.

- The edge constraint $\mathcal{E}_{\Pi'}$ is defined as follows. Consider a configuration $L_1 \dots L_r$ where for all $1 \leq i \leq r$ it holds that $L_i \in 2^{\Sigma_\Pi} \setminus \{\emptyset\}$, such that for any $(\ell_1, \dots, \ell_r) \in L_1 \times \dots \times L_r$ it holds that a permutation of $\ell_1 \dots \ell_r$ is in \mathcal{E}_Π (note that, by definition of the above configuration, $\ell_i \in \Sigma_\Pi$, for all $1 \leq i \leq r$). Let \mathcal{S} be the set that contains all and only such configurations. We say that any configuration in \mathcal{S} satisfies the universal quantifier. We call a configuration $L_1 \dots L_r \in \mathcal{S}$ non-maximal if there exists a permutation $L'_1 \dots L'_r$ of another configuration in \mathcal{S} such that, $L_i \subseteq L'_i$ for all i , and there exists at least an index i such that $L'_i \subsetneq L_i$. The hyperedge constraint $\mathcal{E}_{\Pi'}$ is defined as the set \mathcal{S} where we remove all non-maximal configurations (in other words, only maximal configurations are kept). We refer to computing \mathcal{S} as applying the universal quantifier.
- $\Sigma_{\Pi'}$ is defined as all subsets of Σ_Π that appear at least once in $\mathcal{E}_{\Pi'}$.
- Consider a configuration $L_1 \dots L_\Delta$ of labels in $\Sigma_{\Pi'}$, such that there exists $(\ell_1, \dots, \ell_\Delta) \in L_1 \times \dots \times L_\Delta$ for which it holds that a permutation of $\ell_1 \dots \ell_\Delta$ is in \mathcal{N}_Π . Let \mathcal{S} be the set of all and only such configurations. We say that any configuration in \mathcal{S} satisfies the existential quantifier. The node constraint $\mathcal{N}_{\Pi'}$ is defined as the set \mathcal{S} . We refer to computing \mathcal{S} as applying the existential quantifier.

The problem $\overline{\mathcal{R}}(\mathcal{R}(\Pi(\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi))) = \overline{\mathcal{R}}(\Pi'(\Sigma_{\Pi'}, \mathcal{N}_{\Pi'}, \mathcal{E}_{\Pi'})) = \Pi''(\Sigma_{\Pi''}, \mathcal{N}_{\Pi''}, \mathcal{E}_{\Pi''})$ is defined similarly. The constraint $\mathcal{N}_{\Pi''}$ is obtained by applying the universal quantifier on $\mathcal{N}_{\Pi'}$, the set of labels $\Sigma_{\Pi''}$ is defined as all subsets of $\Sigma_{\Pi'}$ that appear at least once in $\mathcal{N}_{\Pi''}$, and $\mathcal{E}_{\Pi''}$ is obtained by applying the existential quantifier on $\mathcal{E}_{\Pi'}$.

Example: MIS. Let us see now a concrete example that illustrates the rules described above. Our starting problem is $\Pi = \text{MIS}$ on Δ -regular trees, which we already showed how to encode in this formalism, and we now show how $\Pi' = \mathcal{R}(\Pi)$ looks like.

- We start by computing the edge constraint $\mathcal{E}_{\Pi'}$. In order to do so, it is useful to first compute the set $2^{\Sigma_\Pi} \setminus \{\emptyset\} = \{\{M\}, \{P\}, \{O\}, \{M, P\}, \{M, O\}, \{P, O\}, \{M, P, O\}\}$. All configurations with labels in the above-defined set that satisfy the universal quantifier are the following: $\{P\}\{M\}$, $\{O\}\{M\}$, $\{O\}\{O\}$, $\{O\}\{M, O\}$, $\{M\}\{P, O\}$. Let \mathcal{S} be the set that contains all the above configurations. The set $\mathcal{E}_{\Pi'}$ is then defined as the set \mathcal{S} where we remove all non-maximal configurations. The configuration $\{P\}\{M\}$ is non-maximal because there exists $\{M\}\{P, O\}$; the configurations $\{O\}\{M\}$ and $\{O\}\{O\}$ are non-maximal because there exists $\{O\}\{M, O\}$. Hence, $\mathcal{E}_{\Pi'} = \{\{M\}\{P, O\}, \{O\}\{M, O\}\}$.

– For the sake of readability we rename the sets of labels as follows: $\{M\} \mapsto \boxed{M}$, $\{O\} \mapsto \boxed{O}$, $\{P, O\} \mapsto \boxed{P}$, $\{M, O\} \mapsto \boxed{X}$. Hence $\mathcal{E}_{\Pi'} = \{\boxed{P} \boxed{M}, \boxed{O} \boxed{X}\}$.

- $\Sigma_{\Pi'} = \{\boxed{M}, \boxed{O}, \boxed{P}, \boxed{X}\}$.
- It is possible to check that $\mathcal{N}_{\Pi'}$ is given by all the words generated by the following regular expression: $[\boxed{M} \boxed{X}]^\Delta \mid \boxed{P} [\boxed{O} \boxed{P} \boxed{X}]^{\Delta-1}$.

In order to fully apply round elimination, we would then apply $\overline{\mathcal{R}}$ on Π' , and hence perform the following operations:

- Compute the node constraint $\mathcal{N}_{\Pi''}$ by applying the universal quantifier on $\mathcal{N}_{\Pi'}$.
- Define $\Sigma_{\Pi''}$ as the set of labels appearing in $\mathcal{N}_{\Pi''}$.
- Compute the edge constraint $\mathcal{E}_{\Pi''}$ by applying the existential quantifier on $\mathcal{E}_{\Pi'}$.
- As before, we would then perform a renaming to replace sets with plain labels. This is not necessary, but it improves the readability of the result.

By a suitable renaming of the obtained sets of labels, we would obtain the following problem, which, by round elimination, is exactly one round easier than MIS. The node constraint is given by all the words generated by the regular expression $M^\Delta \mid PO^{\Delta-1} \mid XA^{\Delta-1} \mid C^\Delta$, and the edge constraint is given by all the words generated by the regular expression $[MA][PCAU] \mid [XMCAU]U$.

The Automatic Round Elimination Theorem. In order to prove our lower bounds we will use the following theorem proved in [15].

THEOREM 3.1. ([15], THEOREM 4.3 (REPHRASED)) *Let $T > 0$. Consider a class \mathcal{G} of hypergraphs with girth at least $2T + 2$, and some locally checkable problem Π . Then, there exists an algorithm that solves problem Π 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.*

We notice that the more challenging part of applying the round elimination technique is computing the set of configurations that must satisfy the for all quantifier. For example, in computing the intermediate problem Π' , the challenging part is computing $\mathcal{E}_{\Pi'}$. In fact, an alternative and easy way to compute $\mathcal{N}_{\Pi'}$ is to simply define it as the set that contains all configurations given by regular expressions obtained by taking a configuration in \mathcal{N}_{Π} and replacing each label ℓ with the disjunction of all label sets in $\Sigma_{\Pi'}$ that contain ℓ . One thing that may help in proving lower bounds is that, it is not necessary to show that $\mathcal{E}_{\Pi'}$ contains all *and only* those configurations that satisfy the universal quantifier, but it is enough to show that there are no maximal configurations that satisfy the universal quantifier that are not in $\mathcal{E}_{\Pi'}$. In other words, we can add more allowed node and/or edge configurations to our lower bound sequence of problems and make these problems potentially easier, and if we are able to show a lower bound on these problems, then this lower bound applies to harder problems as well. Hence, we now define the notion of *relaxation of a configuration*, which was already introduced in [8].

DEFINITION 3.1. (RELAXATION) *Let $\mathcal{B} = B_1 \dots B_j$ and $\mathcal{B}' = B'_1 \dots B'_k$ be configurations consisting of subsets of some label space Σ . We say that \mathcal{B} can be relaxed to \mathcal{B}' if $j = k$ and there exists a permutation $\rho: \{1, \dots, j\} \rightarrow \{1, \dots, k\}$ such that for any $1 \leq i \leq j$ we have $B_i \subseteq B'_{\rho(i)}$.*

Let C and C' be two constraints, i.e., two collections of configurations. Then we call C' a relaxation of C if every configuration in C can be relaxed to some configuration in C' . Moreover, for two problems $\Pi_1 = (\Sigma_1, \mathcal{N}_1, \mathcal{E}_1)$ and $\Pi_2 = (\Sigma_2, \mathcal{N}_2, \mathcal{E}_2)$, we call Π_2 a relaxation of Π_1 if \mathcal{N}_2 is a relaxation of \mathcal{N}_1 and \mathcal{E}_2 is a relaxation of \mathcal{E}_1 .

Relation Between Labels. In order to argue about desired properties of our lower bound family of problems, it will be useful to determine a relation between labels. Given a problem Π , let C be either the node or the hyperedge constraint, and let A and B be two labels of Π . We say that A is *at least as strong as B according to C* if, for any configuration $\ell_1 \dots \ell_\delta \in C$ ($\delta = \Delta$ or $\delta = r$) that contains B , there is also a configuration in C that is a permutation of $\ell_1 \dots \ell_\delta$ where B is replaced with A . In other words, consider any configuration $C \in C$ that contains B ; Let C' be the configuration obtained by C where B is replaced with A ; a permutation of C' must be in C . If a label A is at least as strong as another label B , then we say that B is *at least as weak as A* , denoted with $B \leq A$. If A is at least as strong as B and B is not at least as strong as A , then we say that A is *stronger than B* and B is *weaker than A* , denoted with $B < A$. If $A \leq B$ and $B \leq A$, then A and B are *equally strong*.

In our lower bound proof we make use of a *diagram* in order to show the relation between labels according to some constraint C : we call it *node diagram* if it shows the relation between labels according to the node constraint \mathcal{N} , and *hyperedge diagram* (or *edge diagram*, if we are on graphs) if instead the relations are according to \mathcal{E} . In this diagram, there is a directed edge from a label B to a label A if A is at least as strong as B and there is no label $D \notin \{A, B\}$ such that D is stronger than B and weaker than A . See Figure 1 for an example of the diagram of the MIS problem according to its edge constraint. All labels in the diagram that are reachable from a label ℓ are called the *successors* of ℓ . In other words, the successors of ℓ , according to some constraint C , are all those labels that are at least as strong as ℓ according to C .




Figure 1: The edge diagram of the MIS problem. There is an edge from P to O because, in all edge configurations where P appears, that is PM, by replacing P with O we get the constraint OM which is also in \mathcal{E} . There is no strength relation between M and labels P or O.

Additional Notation. Lastly, we introduce some additional notation that is essential and heavily used for describing in a concise way the problem family of our lower bounds. Consider a problem $\Pi = (\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi)$, where $\Sigma_\Pi = \{\ell_1, \dots, \ell_\delta\}$. We denote by $\langle \ell_1, \dots, \ell_\delta \rangle$ the set of all labels that are at least as strong as at least one label in $\{\ell_1, \dots, \ell_\delta\}$, according to either \mathcal{N}_Π or \mathcal{E}_Π . In other words, $\langle \ell_1, \dots, \ell_\delta \rangle$ is the set containing the labels $\ell_1, \dots, \ell_\delta$ along with all the successors of the labels ℓ_i , $1 \leq i \leq \delta$, according to either \mathcal{N}_Π or \mathcal{E}_Π . Let us be a bit more precise regarding when it is that, by using $\langle \cdot \rangle$, we refer to successors of labels according to \mathcal{N}_Π and when instead according to \mathcal{E}_Π . Consider a problem $\Pi = (\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi)$ and let the argument of $\langle \cdot \rangle$ contains labels from Σ_Π . Then, when we compute $\mathcal{R}(\Pi) = \Pi' = (\Sigma_{\Pi'}, \mathcal{N}_{\Pi'}, \mathcal{E}'_{\Pi'})$, the set $\langle \cdot \rangle$ contains successors according to $\mathcal{E}_{\Pi'}$. Suppose now that the argument of $\langle \cdot \rangle$ contains labels from $\Sigma_{\Pi'}$. Then, when we compute $\Pi'' = \overline{\mathcal{R}}(\Pi') = \overline{\mathcal{R}}(\mathcal{R}(\Pi))$, the set $\langle \cdot \rangle$ contains successors according to $\mathcal{N}'_{\Pi'}$. Hence, the relation between labels is always considered according to the constraint on which we apply the universal quantifier. Actually, when we compute $\overline{\mathcal{R}}(\Pi)$ in our lower bound proofs, we will often use the expression of the form $\langle\langle \ell \rangle\rangle$, where $\ell \in \Sigma_\Pi$, that represents a set of sets of labels in Σ_Π , where $\langle \ell \rangle$ is generated according to \mathcal{E}_Π and $\langle\langle \ell \rangle\rangle$ according to $\mathcal{N}'_{\Pi'}$.

We say that a set $S = \{\ell_1, \dots, \ell_\delta\} \subseteq \Sigma_\Pi$ is *right-closed* if it contains also all successors of the labels $\ell_1, \dots, \ell_\delta$ according to \mathcal{E}_Π , that is, if $S = \langle \ell_1, \dots, \ell_\delta \rangle$. For a label ℓ , we denote with *right-closed subset generated by ℓ* the set $\langle \ell \rangle$. Similarly, for a set of labels $\{\ell_1, \dots, \ell_\delta\}$, the right-closed subset generated by that set is $\langle \ell_1, \dots, \ell_\delta \rangle$. Observe that it is equal to $\bigcup_i \langle \ell_i \rangle$. In our proofs, we will make use of the following simple observation from [8].

OBSERVATION 3.1. ([8]) *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 (according to \mathcal{E}_Π). If $\{A_1, \dots, A_p\} \in \Sigma_{\overline{\mathcal{R}}(\Pi)}$, then the set $\{A_1, \dots, A_p\}$ is right-closed (according to \mathcal{N}_Π).*

In order to get our lower bound, we will make use of the following theorem, which has been already proved in [7], and allows us to transform a lower bound obtained through round elimination into a lower bound for the LOCAL model.

THEOREM 3.2. (THEOREM 7.1 OF [7], REPHRASED) *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 Π'_i that is a relaxation of $\mathcal{R}(\Pi_i)$;*
- *Π_{i+1} is a relaxation of $\overline{\mathcal{R}}(\Pi'_i)$;*
- *The number of labels of Π_i , and the ones of Π'_i , are upper bounded by $f(\Delta \cdot r)$.*

Also, assume that Π_t has at most $f(\Delta \cdot r)$ labels and is not 0-round solvable in the deterministic port numbering model. Then, Π_0 requires $\Omega(\min\{t, \log_{\Delta r} n - \log_{\Delta r} \log f(\Delta r)\})$ rounds in the deterministic LOCAL model and $\Omega(\min\{t, \log_{\Delta r} \log n - \log_{\Delta r} \log f(\Delta r)\})$ rounds in the randomized LOCAL model.

For any subset W of nodes in some graph, let $N(W)$ denote the set of nodes that are adjacent to at least one member of W . In our lower bound proofs, we will make use of Hall's marriage theorem, which can be phrased as follows.

THEOREM 3.3. (HALL) *Let $G = (V \cup W, E)$ be a bipartite graph. If $|N(W')| \geq |W'|$ for any $W' \subseteq W$, then there exists a matching saturating W , i.e., a function $f: W \rightarrow V$ such that $f(w) \neq f(w')$ for any $w \neq w'$, and $\{w, f(w)\} \in E$ for any $w \in W$.*

4 Lower Bounds for Hypergraph Colorings

In this section, we prove lower bounds for different variants of hypergraph coloring. We first consider the standard definition of coloring, where on each hyperedge it must hold that at least two incident nodes have different colors. We show how to encode the Δ -hypergraph coloring problem in the round elimination framework, and then we show which additional allowed configurations we have to allow in order to transform it into a fixed point. A problem is a fixed point if, modulo some renaming, $\overline{\mathcal{R}}(\mathcal{R}(\Pi)) = \Pi$, and by prior work we know that if a fixed point problem cannot be solved in 0 rounds in the PN model, then it means that it requires $\Omega(\log_{\Delta r} n)$ rounds for deterministic algorithms in the LOCAL model, and of $\Omega(\log_{\Delta r} \log n)$ rounds for randomized ones. We then show that there is a 0 round algorithm that is able to convert a $\Delta(r - 1)$ -colorful coloring into a Δ -coloring, implying that the lower bound that we show for Δ -hypergraph coloring holds for the colorful variant as well.

4.1 The Δ -Hypergraph Coloring Problem. We now show how to define the hypergraph coloring problem in the round elimination framework. The label set Σ_{Π} contains one label for each color, that is, $\Sigma_{\Pi} = \{1, \dots, \Delta\}$. The node constraint \mathcal{N}_{Π} contains one configuration for each possible color, that is, $\mathcal{N}_{\Pi} = \{c^{\Delta} \mid 1 \leq c \leq \Delta\}$. The hyperedge constraint \mathcal{E}_{Π} contains all words of size r where at least two different colors appear, that is, $\mathcal{E}_{\Pi} = \{c_1 \dots c_r \mid c_i \in \{1, \dots, \Delta\} \wedge c_1 \neq c_2\}$.

We now provide the high level idea on how this problem can be relaxed in order to make it a fixed point. The idea is to allow nodes to output not just colors, but sets of colors. For example, a node could use color 1 and 2 at the same time, but then, on each hyperedge incident to this node, there should be another incident node that is neither of color 1 *nor* of color 2. Additionally, nodes are *rewarded* for using more than one color: if a node uses x colors, then it is allowed to mark up to $x - 1$ incident hyperedges, and on those marked hyperedges the node counts as being uncolored. Observe that if a hyperedge e is marked by at least one node, then the hyperedge constraint is always satisfied on e , because for each color it holds that there is at least one incident node that is *not* colored with that color.

There is no clear intuition why this problem relaxation would be a fixed point, and in fact our proof is essentially a tedious case analysis. But, on a very high level, when applying the universal quantifier, the functions \mathcal{R} and $\overline{\mathcal{R}}$ essentially try to combine the existing configurations in all possible ways that make the forall quantifier satisfied. If we look at $\Pi' = \overline{\mathcal{R}}(\mathcal{R}(\Pi))$, where Π is the original Δ -hypergraph coloring problem, we actually see something very similar to the problem that we just described. The main difference is that the allowed configurations of Π' are a bit more restrictive: in the fixed point relaxations, if a hyperedge is marked by some incident node, then it is always happy, while in Π' this is not always the case. Relaxing Π' to make hyperedges always happy when marked is essentially what gives our fixed point relaxation.

4.2 The Δ -Hypergraph Coloring Fixed Point. We now define a problem Π that is a relaxation of the Δ -hypergraph coloring problem and that we will later show to be a fixed point problem that cannot be solved in 0 rounds.

Let $\mathcal{C} := \{1, \dots, \Delta\}$. We call the elements of \mathcal{C} *colors* from the *color space* \mathcal{C} . In the following, we define $\Pi = (\Sigma_{\Pi}, \mathcal{N}_{\Pi}, \mathcal{E}_{\Pi})$ by specifying Σ_{Π} , \mathcal{N}_{Π} , and \mathcal{E}_{Π} .

The label set Σ_{Π} . The label set Σ_{Π} of Π contains one label $\ell(\mathcal{C})$ for each (possibly empty) color set $\mathcal{C} \subseteq \mathcal{C}$, i.e., $\Sigma_{\Pi} := \{\ell(\mathcal{C}) \mid \mathcal{C} \subseteq \mathcal{C}\}$. In other words, the labels are now sets of colors.

The node constraint \mathcal{N}_{Π} . The node constraint \mathcal{N}_{Π} of Π consists of all configurations of the form $\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \ell(\emptyset)^{|\mathcal{C}|-1}$ for some $\emptyset \neq \mathcal{C} \subseteq \mathcal{C}$, i.e., $\mathcal{N}_{\Pi} := \{\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \ell(\emptyset)^{|\mathcal{C}|-1} \mid \emptyset \neq \mathcal{C} \subseteq \mathcal{C}\}$. In other words, each node outputs a set of colors, and on some of the incident hyperedges the empty set can be used. The intuition is that the empty set corresponds to the marking previously described.

The hyperedge constraint \mathcal{E}_{Π} . The hyperedge constraint \mathcal{E}_{Π} of Π consists of all configurations $\ell(\mathcal{C}_1) \dots \ell(\mathcal{C}_r)$ of labels from Σ_{Π} such that for each color $i \in \mathcal{C}$, there is at least one index $1 \leq j \leq r$ satisfying $i \notin \mathcal{C}_j$.

A Fixed Point. We will prove that Π is a fixed point, that is, under some renaming, $\overline{\mathcal{R}}(\mathcal{R}(\Pi)) = \Pi$.

LEMMA 4.1. *The problem $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ is equivalent to Π .*

Since the proof of this statement is a tedious case analysis, we defer its proof to the end of the section. We now use this statement to derive our lower bounds.

4.3 Proving the Lower Bounds We now prove that Π cannot be solved in 0 rounds. We will then show that this fact, combined with Lemma 4.1, implies strong lower bounds for different hypergraph coloring variants.

LEMMA 4.2. *The problem Π cannot be solved in 0 rounds in the deterministic port numbering model.*

Proof. All nodes of degree Δ that run a 0-round algorithm have the same view (they just know their degree, the size of the graph, and the parameters Δ and r). Hence, any deterministic 0-round algorithm uses the same configuration in \mathcal{N}_{Π} on all nodes. Thus, in order to prove the statement, we consider all possible configurations and we show that they cannot be properly used in a 0-round algorithm. The configurations allowed on the nodes are all those that satisfy $\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \ell(\emptyset)^{|\mathcal{C}|-1}$ for each $\emptyset \neq \mathcal{C} \subseteq \mathcal{C}$. Hence, all nodes must output one of those configuration for the same set \mathcal{C} satisfying $\emptyset \neq \mathcal{C} \subseteq \mathcal{C}$. Since $|\mathcal{C}| \leq \Delta$, then this configuration contains $\ell(\mathcal{C})$ at least once. Since the algorithm is deterministic, then all nodes must output $\ell(\mathcal{C})$ on the same port. W.l.o.g., let this

port be the port number 1. If we connect the port 1 of r nodes to the same hyperedge, then the hyperedge would have the configuration $\ell(\mathcal{C})^r$, which is not in \mathcal{E}_Π . \square

THEOREM 4.1. *The hypergraph Δ -coloring problem requires $\Omega(\log_{\Delta r} n)$ in the LOCAL model for deterministic algorithms, and $\Omega(\log_{\Delta r} \log n)$ for randomized ones.*

Proof. By Lemma 4.1, $\overline{\mathcal{R}}(\mathcal{R}(\Pi)) = \Pi$, and by Lemma 4.2 Π is not 0-round solvable. Hence, we can build an arbitrarily long sequence of problems satisfying Theorem 3.2 (or, in other words, we can take $t = \infty$). Also, note that the number of labels of each problem is bounded by $2^{O(\Delta)}$. Thus, since $\log_{\Delta r} \log f(\Delta) = O(1)$, we obtain a lower bound for Π of $\Omega(\log_{\Delta r} n)$ rounds in the LOCAL model for deterministic algorithms, and $\Omega(\log_{\Delta r} \log n)$ rounds for randomized ones. \square

THEOREM 4.2. *The hypergraph $\Delta(r - 1)$ -colorful coloring problem requires $\Omega(\log_{\Delta r} n)$ in the LOCAL model for deterministic algorithms, and $\Omega(\log_{\Delta r} \log n)$ for randomized ones.*

Proof. We prove that, given a solution for the $\Delta(r - 1)$ -hypergraph coloring problem, we can solve hypergraph Δ -coloring in 0 rounds, implying the claim. We group the $\Delta(r - 1)$ colors of the hypergraph colorful coloring problem into Δ groups G_1, \dots, G_Δ of $r - 1$ colors each, in an arbitrary way. Then, each node of color c outputs $\ell(\{i\})^\Delta$, for i satisfying $c \in G_i$. Observe that the solution clearly satisfies \mathcal{N}_Π . Also, since in any solution for the $\Delta(r - 1)$ -hypergraph coloring problem on each hyperedge there can be at most 1 incident node for each color, then we obtain that each hypergraph has at most $r - 1$ incident nodes with colors from the same group. Hence, the constraint \mathcal{E}_Π is satisfied. \square

4.4 Relations Between Labels In the rest of the section we prove Lemma 4.1. We start by proving a relation between the labels. From the definition of \mathcal{E}_Π , we can infer the following lemma characterizing the strength relations according to \mathcal{E}_Π .

LEMMA 4.3. *Consider any two subsets $\mathcal{C}, \mathcal{C}'$ of \mathfrak{C} . Then $\ell(\mathcal{C}) \leq \ell(\mathcal{C}')$ if and only if $\mathcal{C}' \subseteq \mathcal{C}$.*

Proof. Consider first the case that $\mathcal{C}' \subseteq \mathcal{C}$, and let $\ell(\mathcal{C})\ell(\mathcal{C}_2) \dots \ell(\mathcal{C}_r)$ be an arbitrary configuration from \mathcal{E}_Π . Then it follows directly from the definition of \mathcal{E}_Π that also $\ell(\mathcal{C}')\ell(\mathcal{C}_2) \dots \ell(\mathcal{C}_r)$ is contained in \mathcal{E}_Π . Hence, $\ell(\mathcal{C}) \leq \ell(\mathcal{C}')$, by the definition of strength.

Now consider the other case, namely that $\mathcal{C}' \not\subseteq \mathcal{C}$, and let $i \in \mathcal{C}'$ be a color that is not contained in \mathcal{C} . Consider the configuration $\ell(\mathcal{C})\ell(\mathfrak{C} \setminus \mathcal{C})\ell(\mathfrak{C})^{r-2}$. Since, for any $i' \in \mathfrak{C}$, we have $i' \notin \mathcal{C}$ or $i' \in \mathfrak{C} \setminus \mathcal{C}$, we have $\ell(\mathcal{C})\ell(\mathfrak{C} \setminus \mathcal{C})\ell(\mathfrak{C})^{r-2} \in \mathcal{E}_\Pi$. Now consider the configuration $\ell(\mathcal{C}')\ell(\mathfrak{C} \setminus \mathcal{C})\ell(\mathfrak{C})^{r-2}$ obtained from the above configuration by replacing $\ell(\mathcal{C})$ by $\ell(\mathcal{C}')$. Since $i \in \mathcal{C}'$, $i \in \mathfrak{C} \setminus \mathcal{C}$, and $i \in \mathfrak{C}$, we have $\ell(\mathcal{C}')\ell(\mathfrak{C} \setminus \mathcal{C})\ell(\mathfrak{C})^{r-2} \notin \mathcal{E}_\Pi$. Hence, $\ell(\mathcal{C}) \not\leq \ell(\mathcal{C}')$, by the definition of strength. \square

From Lemma 4.3, we immediately obtain the following corollary, by collecting, for each label $\ell(\mathcal{C}) \in \Sigma_\Pi$, all labels $\ell(\overline{\mathcal{C}}) \in \Sigma_\Pi$ satisfying $\ell(\mathcal{C}) \leq \ell(\overline{\mathcal{C}})$.

COROLLARY 4.1. *For each $\mathcal{C} \subseteq \mathfrak{C}$, we have $\langle \ell(\mathcal{C}) \rangle = \{ \ell(\mathcal{C}') \mid \mathcal{C}' \subseteq \mathcal{C} \}$.*

4.5 Computing $\mathcal{R}(\Pi)$ We now prove that problem $\mathcal{R}(\Pi)$ is defined as in the following lemma.

LEMMA 4.4. *The set $\Sigma_{\mathcal{R}}$ of output labels of $\mathcal{R}(\Pi)$ is given by $\{ \langle \ell(\mathcal{C}) \rangle \mid \mathcal{C} \subseteq \mathfrak{C} \}$.*

The hyperedge constraint $\mathcal{E}_{\mathcal{R}}$ of $\mathcal{R}(\Pi)$ consists of all configurations $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle$ of labels from $\{ \langle \ell(\mathcal{C}) \rangle \mid \mathcal{C} \subseteq \mathfrak{C} \}$ such that for each color $i \in \mathfrak{C}$, there is exactly one index $1 \leq j \leq r$ satisfying $i \notin \mathcal{C}_j$.

The node constraint $\mathcal{N}_{\mathcal{R}}$ of $\mathcal{R}(\Pi)$ consists of all configurations $\mathbf{B}_1 \dots \mathbf{B}_\Delta$ of labels from $\{ \langle \ell(\mathcal{C}) \rangle \mid \mathcal{C} \subseteq \mathfrak{C} \}$ such that there exists a choice $(\mathbf{A}_1, \dots, \mathbf{A}_\Delta) \in \mathbf{B}_1 \times \dots \times \mathbf{B}_\Delta$ satisfying $\mathbf{A}_1 \dots \mathbf{A}_\Delta \in \mathcal{N}_\Pi$.

Proof. We start by showing that the hyperedge constraint $\mathcal{E}_{\mathcal{R}}$ is as given in the lemma. Let \mathcal{E} denote the set of all configurations $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle$ of labels from $\Sigma_{\mathcal{R}}$ such that for each color $i \in \mathfrak{C}$, there is exactly one index $1 \leq j \leq r$ satisfying $i \notin \mathcal{C}_j$. We need to show that $\mathcal{E}_{\mathcal{R}} = \mathcal{E}$.

We first show that $\mathcal{E}_{\mathcal{R}} \subseteq \mathcal{E}$. Let $\mathbf{B}_1 \dots \mathbf{B}_r$ be an arbitrary configuration from $\mathcal{E}_{\mathcal{R}}$. Our first step is to show that each \mathbf{B}_j is of the form $\langle \ell(\mathcal{C}) \rangle$ for some $\mathcal{C} \subseteq \mathfrak{C}$. Assume for a contradiction that there exists some index

$1 \leq j \leq r$ such that $B_j \neq \langle \ell(\mathcal{C}) \rangle$ for each $\mathcal{C} \in \mathfrak{C}$. Let \mathcal{C} be the union of all color sets \mathcal{C}' such that $\ell(\mathcal{C}') \in B_j$. By Corollary 4.1, it follows that $B_j \subseteq \langle \ell(\mathcal{C}) \rangle$, and since $B_j \neq \langle \ell(\mathcal{C}) \rangle$, we know that $\langle \ell(\mathcal{C}) \rangle \setminus B_j$ is nonempty. Observe that, by Observation 3.1, B_j is right-closed (according to \mathcal{E}_Π) as $B_j \in \Sigma_{\mathcal{R}}$ (by the definition of B_j). Hence, if $\ell(\mathcal{C}) \in B_j$, then $\langle \ell(\mathcal{C}) \rangle \subseteq B_j$, yielding a contradiction to the nonemptiness of $\langle \ell(\mathcal{C}) \rangle \setminus B_j$. Thus, $\ell(\mathcal{C}) \notin B_j$.

Now consider the configuration $B_1 \dots B_{j-1} \langle \ell(\mathcal{C}) \rangle B_{j+1} \dots B_r$. Since $B_j \subsetneq \langle \ell(\mathcal{C}) \rangle$, we know that $B_1 \dots B_{j-1} \langle \ell(\mathcal{C}) \rangle B_{j+1} \dots B_r$ and all configurations that $B_1 \dots B_{j-1} \langle \ell(\mathcal{C}) \rangle B_{j+1} \dots B_r$ can be relaxed to are not contained in $\mathcal{E}_{\mathcal{R}}$, by (the maximality condition in) the definition of $\mathcal{E}_{\mathcal{R}}$ (and the fact that $B_1 \dots B_r \in \mathcal{E}_{\mathcal{R}}$). It follows, by the definitions of $\mathcal{E}_{\mathcal{R}}$ and \mathcal{E}_Π , that there exist some choice $(\ell(\mathcal{C}'_1), \dots, \ell(\mathcal{C}'_r)) \in B_1 \times \dots \times B_{j-1} \times \langle \ell(\mathcal{C}) \rangle \times B_{j+1} \times \dots \times B_r$ and some color $i \in \mathfrak{C}$ such that $i \in \mathcal{C}'_{j'}$ for each $1 \leq j' \leq r$. By Corollary 4.1, the fact that $\ell(\mathcal{C}'_{j'}) \in \langle \ell(\mathcal{C}) \rangle$ implies that $\mathcal{C}'_{j'} \subseteq \mathcal{C}$, which, combined with $i \in \mathcal{C}'_{j'}$, yields $i \in \mathcal{C}$. Let $\mathcal{C}^* \subseteq \mathfrak{C}$ be a color set such that $i \in \mathcal{C}^*$ and $\ell(\mathcal{C}^*) \in B_j$. Such a color set \mathcal{C}^* exists due to the definition of \mathcal{C} and the fact that $i \in \mathcal{C}$. Since $i \in \mathcal{C}^*$ and $i \in \mathcal{C}'_{j'}$ for each $1 \leq j' \leq r$ satisfying $j' \neq j$, we obtain $\ell(\mathcal{C}'_1) \dots \ell(\mathcal{C}'_{j-1}) \ell(\mathcal{C}^*) \ell(\mathcal{C}'_{j+1}) \dots \ell(\mathcal{C}'_r) \notin \mathcal{E}_\Pi$, by the definition of \mathcal{E}_Π . But, since $(\ell(\mathcal{C}'_1), \dots, \ell(\mathcal{C}'_{j-1}), \ell(\mathcal{C}^*), \ell(\mathcal{C}'_{j+1}), \dots, \ell(\mathcal{C}'_r)) \in B_1 \times \dots \times B_r$, this implies that $B_1 \dots B_r \notin \mathcal{E}_{\mathcal{R}}$, by the definition of $\mathcal{E}_{\mathcal{R}}$, yielding a contradiction. Hence, each B_j is of the form $\langle \ell(\mathcal{C}) \rangle$ for some $\mathcal{C} \subseteq \mathfrak{C}$.

Now, let $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle$ be an arbitrary configuration from $\mathcal{E}_{\mathcal{R}}$ (justified by the above discussion), and consider an arbitrary color $i \in \mathfrak{C}$. If for each index $1 \leq j \leq r$ we have $i \in \mathcal{C}_j$, then we know that $\ell(\mathcal{C}_1) \dots \ell(\mathcal{C}_r) \notin \mathcal{E}_\Pi$, by the definition of \mathcal{E}_Π ; however, since, by Corollary 4.1, we have $\ell(\mathcal{C}_j) \in \langle \ell(\mathcal{C}_j) \rangle$ for each $1 \leq j \leq r$, it follows that $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle \notin \mathcal{E}_{\mathcal{R}}$ (by the definition of $\mathcal{E}_{\mathcal{R}}$), yielding a contradiction. Hence, we can conclude that there exists at least one index $1 \leq j \leq r$ such that $i \notin \mathcal{C}_j$.

It remains to show that there do not exist two distinct indices $j, j' \in \{1, \dots, r\}$ such that $i \notin \mathcal{C}_j$ and $i \notin \mathcal{C}_{j'}$. For a contradiction, assume that such two indices j, j' exist. Consider the configuration $Y = \langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_{j-1}) \rangle \langle \ell(\mathcal{C}_j \cup \{i\}) \rangle \langle \ell(\mathcal{C}_{j+1}) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle$. Since $i \notin \mathcal{C}_j$, we have $\langle \ell(\mathcal{C}_j \cup \{i\}) \rangle \supsetneq \langle \ell(\mathcal{C}_j) \rangle$, by Corollary 4.1, and it follows by (the maximality condition in) the definition of $\mathcal{E}_{\mathcal{R}}$ (and the fact that $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle \in \mathcal{E}_{\mathcal{R}}$) that $Y \notin \mathcal{E}_{\mathcal{R}}$.

Now consider an arbitrary choice $(\ell(\mathcal{C}'_1), \dots, \ell(\mathcal{C}'_r)) \in \langle \ell(\mathcal{C}_1) \rangle \times \dots \times \langle \ell(\mathcal{C}_{j-1}) \rangle \times \langle \ell(\mathcal{C}_j \cup \{i\}) \rangle \times \langle \ell(\mathcal{C}_{j+1}) \rangle \times \dots \times \langle \ell(\mathcal{C}_r) \rangle$. Observe that $\ell(\mathcal{C}'_j) \in \langle \ell(\mathcal{C}_j \cup \{i\}) \rangle$ implies that $\ell(\mathcal{C}'_j \setminus \{i\}) \in \langle \ell(\mathcal{C}_j) \rangle$, by Corollary 4.1. Thus, $\ell(\mathcal{C}'_1) \dots \ell(\mathcal{C}'_{j-1}) \ell(\mathcal{C}'_j \setminus \{i\}) \ell(\mathcal{C}'_{j+1}) \dots \ell(\mathcal{C}'_r) \in \mathcal{E}_\Pi$, by the definition of $\mathcal{E}_{\mathcal{R}}$ (and the fact that $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle \in \mathcal{E}_{\mathcal{R}}$). It follows, by the definition of \mathcal{E}_Π , that for each $i' \in \mathfrak{C}$ satisfying $i' \neq i$, there is at least one index $1 \leq j'' \leq r$ satisfying $i' \notin \mathcal{C}'_{j''}$. Observe that there also is at least one index $1 \leq j'' \leq r$ satisfying $j'' \neq j$ and $i \notin \mathcal{C}'_{j''}$, due to the assumptions that $j' \neq j$ and $i \notin \mathcal{C}_{j'}$. Hence, we can conclude that also $\ell(\mathcal{C}'_1) \dots \ell(\mathcal{C}'_r) \in \mathcal{E}_\Pi$, by the definition of \mathcal{E}_Π . Since $(\ell(\mathcal{C}'_1), \dots, \ell(\mathcal{C}'_r))$ was chosen arbitrarily from $\langle \ell(\mathcal{C}_1) \rangle \times \dots \times \langle \ell(\mathcal{C}_{j-1}) \rangle \times \langle \ell(\mathcal{C}_j \cup \{i\}) \rangle \times \langle \ell(\mathcal{C}_{j+1}) \rangle \times \dots \times \langle \ell(\mathcal{C}_r) \rangle$, it follows that $Y \in \mathcal{E}_{\mathcal{R}}$, by the definition of $\mathcal{E}_{\mathcal{R}}$. This yields a contradiction to the already established fact that $Y \notin \mathcal{E}_{\mathcal{R}}$. Thus, there is exactly one index $1 \leq j \leq r$ such that $i \notin \mathcal{C}_j$. Since i was chosen arbitrarily from \mathfrak{C} , it follows that $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle \in \mathcal{E}$, by the definition of \mathcal{E} . Since $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle$ was chosen arbitrarily from $\mathcal{E}_{\mathcal{R}}$, this implies in turn that $\mathcal{E}_{\mathcal{R}} \subseteq \mathcal{E}$, as desired.

Now, we show that $\mathcal{E} \subseteq \mathcal{E}_{\mathcal{R}}$. Let $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle$ be an arbitrary configuration from \mathcal{E} , and consider some arbitrary choice $(\ell(\mathcal{C}'_1), \dots, \ell(\mathcal{C}'_r)) \in \langle \ell(\mathcal{C}_1) \rangle \times \dots \times \langle \ell(\mathcal{C}_r) \rangle$. From the definition of \mathcal{E} and Corollary 4.1, it follows that for each color $i \in \mathfrak{C}$, there is at least one index $1 \leq j \leq r$ satisfying $i \notin \mathcal{C}'_j$. Hence $\ell(\mathcal{C}'_1) \dots \ell(\mathcal{C}'_r) \in \mathcal{E}_\Pi$, by the definition of \mathcal{E}_Π . Since $(\ell(\mathcal{C}'_1), \dots, \ell(\mathcal{C}'_r))$ was arbitrarily chosen from $\langle \ell(\mathcal{C}_1) \rangle \times \dots \times \langle \ell(\mathcal{C}_r) \rangle$, this implies, by the definition of $\mathcal{E}_{\mathcal{R}}$, that $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle \in \mathcal{E}_{\mathcal{R}}$ or there is some configuration $B_1 \dots B_r \in \mathcal{E}_{\mathcal{R}}$ such that $\langle \ell(\mathcal{C}_j) \rangle \subseteq B_j$ for all $j \in \{1, \dots, r\}$ and $\langle \ell(\mathcal{C}_j) \rangle \subsetneq B_j$ for at least one $j \in \{1, \dots, r\}$. Assume that such a configuration $B_1 \dots B_r \in \mathcal{E}_{\mathcal{R}}$ exists. By the definition of $B_1 \dots B_r$, the definition of \mathcal{E} , and the fact that $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle \in \mathcal{E}$, there must be some color $i \in \mathfrak{C}$ such that $i \in B_j$ for each $1 \leq j \leq r$. Hence, $B_1 \dots B_r \notin \mathcal{E}$ (by the definition of \mathcal{E}), which contradicts $B_1 \dots B_r \in \mathcal{E}_{\mathcal{R}}$ since $\mathcal{E}_{\mathcal{R}} \subseteq \mathcal{E}$ (as established above). Thus, no such configuration $B_1 \dots B_r \in \mathcal{E}_{\mathcal{R}}$ exists, which implies that $\langle \ell(\mathcal{C}_1) \rangle \dots \langle \ell(\mathcal{C}_r) \rangle \in \mathcal{E}_{\mathcal{R}}$. It follows that $\mathcal{E} \subseteq \mathcal{E}_{\mathcal{R}}$, as desired. Hence, we can conclude that $\mathcal{E}_{\mathcal{R}} = \mathcal{E}$, and it follows that the hyperedge constraint $\mathcal{E}_{\mathcal{R}}$ is as given in the lemma.

Next, we show that $\Sigma_{\mathcal{R}}$ is as given in the lemma. By definition, $\Sigma_{\mathcal{R}}$ consists of those labels that appear in some configuration from $\mathcal{E}_{\mathcal{R}}$. Observe that, for any $\mathcal{C} \subseteq \mathfrak{C}$, the configuration $\langle \ell(\mathcal{C}) \rangle \langle \ell(\mathfrak{C} \setminus \mathcal{C}) \rangle \langle \ell(\mathfrak{C}) \rangle^{r-2}$ is contained in \mathcal{E} , by the definition of \mathcal{E} . Since $\mathcal{E}_{\mathcal{R}} = \mathcal{E}$ (and all labels appearing in some configuration from $\mathcal{E}_{\mathcal{R}}$ are contained in $\{\langle \ell(\mathcal{C}) \rangle \mid \mathcal{C} \subseteq \mathfrak{C}\}$), it follows that $\Sigma_{\mathcal{R}}$ is as given in the lemma.

From the definition of $\mathcal{R}(\cdot)$, it immediately follows that $\mathcal{N}_{\mathcal{R}}$ is as given in the lemma. \square

Similar to before, it will be useful to collect information about the strength of the labels in $\Sigma_{\mathcal{R}}$ (according to $\mathcal{N}_{\mathcal{R}}$) and compute the right-closed subsets generated by each label. We will do so in the following with Lemma 4.5

and Corollary 4.2, which are analogues of Lemma 4.3 and Corollary 4.1 for $\mathcal{R}(\Pi)$ instead of Π .

LEMMA 4.5. Consider any two subsets $\mathcal{C}, \mathcal{C}'$ of \mathfrak{C} . Then $\langle \ell(\mathcal{C}) \rangle \leq \langle \ell(\mathcal{C}') \rangle$ if and only if $\mathcal{C} \subseteq \mathcal{C}'$.

Proof. Consider first the case that $\mathcal{C} \subseteq \mathcal{C}'$. From the definition of $\mathcal{N}_{\mathcal{R}}$, it directly follows that replacing some set in a configuration from $\mathcal{N}_{\mathcal{R}}$ by a superset will result in a configuration that is contained in $\mathcal{N}_{\mathcal{R}}$ as well. Since, by Corollary 4.1, $\mathcal{C} \subseteq \mathcal{C}'$ implies $\langle \ell(\mathcal{C}) \rangle \subseteq \langle \ell(\mathcal{C}') \rangle$, it follows that $\langle \ell(\mathcal{C}) \rangle \leq \langle \ell(\mathcal{C}') \rangle$ (by the definition of strength), as desired.

Now consider the other case, namely that $\mathcal{C} \not\subseteq \mathcal{C}'$, and let $i \in \mathcal{C}$ be a color that is not contained in \mathcal{C}' . In particular, we have $\mathcal{C} \neq \emptyset$. Consider the configuration $\langle \ell(\mathcal{C}) \rangle^{\Delta-|\mathcal{C}|+1} \langle \ell(\emptyset) \rangle^{|\mathcal{C}|-1}$. Since $\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \ell(\emptyset)^{|\mathcal{C}|-1} \in \mathcal{N}_{\Pi}$, we have $\langle \ell(\mathcal{C}) \rangle^{\Delta-|\mathcal{C}|+1} \langle \ell(\emptyset) \rangle^{|\mathcal{C}|-1} \in \mathcal{N}_{\mathcal{R}}$, by the definition of $\mathcal{N}_{\mathcal{R}}$. Now consider the configuration $\langle \ell(\mathcal{C}') \rangle \langle \ell(\mathcal{C}) \rangle^{\Delta-|\mathcal{C}|} \langle \ell(\emptyset) \rangle^{|\mathcal{C}|-1}$ obtained from the above configuration by replacing (one) $\langle \ell(\mathcal{C}) \rangle$ by $\langle \ell(\mathcal{C}') \rangle$. Consider an arbitrary choice $(\ell(\mathcal{C}_1), \dots, \ell(\mathcal{C}_{\Delta})) \in \langle \ell(\mathcal{C}') \rangle \times \langle \ell(\mathcal{C}) \rangle^{\Delta-|\mathcal{C}|} \times \langle \ell(\emptyset) \rangle^{|\mathcal{C}|-1}$. Set $Y := \ell(\mathcal{C}_1) \dots \ell(\mathcal{C}_{\Delta})$. We want to show that $Y \notin \mathcal{N}_{\Pi}$.

For a contradiction, assume that $Y \in \mathcal{N}_{\Pi}$, which implies that there is some color set $\emptyset \neq \mathcal{C}^* \subseteq \mathfrak{C}$ such that $Y = \ell(\mathcal{C}^*)^{\Delta-|\mathcal{C}^*|+1} \ell(\emptyset)^{|\mathcal{C}^*|-1}$. If $|\mathcal{C}^*| < |\mathcal{C}|$, then Y contains strictly more than $\Delta - |\mathcal{C}| + 1$ labels that are not contained in $\langle \ell(\emptyset) \rangle$ (by Corollary 4.1), which yields a contradiction to the definition of Y . Hence, $|\mathcal{C}^*| \geq |\mathcal{C}|$. Consider the case that $\mathcal{C}^* \neq \mathcal{C}$. Then it follows that $\mathcal{C}^* \not\subseteq \mathcal{C}$, which implies $\ell(\mathcal{C}^*) \notin \langle \ell(\mathcal{C}) \rangle$, by Corollary 4.1; it follows that $\ell(\mathcal{C}^*)$ is contained in at most one set from $\langle \ell(\mathcal{C}') \rangle \langle \ell(\mathcal{C}) \rangle^{\Delta-|\mathcal{C}|} \langle \ell(\emptyset) \rangle^{|\mathcal{C}|-1}$, which in turn implies $\Delta - |\mathcal{C}^*| + 1 \leq 1$, by the definition of Y , and therefore $|\mathcal{C}^*| = \Delta$. As $|\mathcal{C}^*| = \Delta$ implies $\mathcal{C}^* = \mathfrak{C}$, but $\ell(\mathfrak{C}) \notin \langle \ell(\mathcal{C}') \rangle \cup \langle \ell(\mathcal{C}) \rangle \cup \langle \ell(\emptyset) \rangle$ (due to the fact that $i \notin \mathcal{C}'$, the fact that $\mathfrak{C} = \mathcal{C}^* \not\subseteq \mathcal{C}$, and Corollary 4.1), we obtain a contradiction to the definition of Y (given the fact that $Y = \ell(\mathcal{C}^*)^{\Delta-|\mathcal{C}^*|+1} \ell(\emptyset)^{|\mathcal{C}^*|-1}$). Hence, we know that $\mathcal{C}^* = \mathcal{C}$. Since $\mathcal{C} \not\subseteq \mathcal{C}'$, it follows, by Corollary 4.1, that $\ell(\mathcal{C}^*) \notin \langle \ell(\mathcal{C}') \rangle$. Thus, $\ell(\mathcal{C}^*)$ is contained in at most $\Delta - |\mathcal{C}| = \Delta - |\mathcal{C}^*|$ sets from $\langle \ell(\mathcal{C}') \rangle \langle \ell(\mathcal{C}) \rangle^{\Delta-|\mathcal{C}|} \langle \ell(\emptyset) \rangle^{|\mathcal{C}|-1}$, which yields a contradiction to the definition of Y (given the fact that $Y = \ell(\mathcal{C}^*)^{\Delta-|\mathcal{C}^*|+1} \ell(\emptyset)^{|\mathcal{C}^*|-1}$). Hence, our assumption was false, and we have $Y \notin \mathcal{N}_{\Pi}$. It follows that $\langle \ell(\mathcal{C}') \rangle \langle \ell(\mathcal{C}) \rangle^{\Delta-|\mathcal{C}|} \langle \ell(\emptyset) \rangle^{|\mathcal{C}|-1} \notin \mathcal{N}_{\mathcal{R}}$, which, by the definition of strength, implies that $\langle \ell(\mathcal{C}) \rangle \not\leq \langle \ell(\mathcal{C}') \rangle$, as desired. \square

From Lemma 4.5 we obtain Corollary 4.2 by simply collecting for any label $L \in \Sigma_{\mathcal{R}}$ the set of all labels L' satisfying $L \leq L'$.

COROLLARY 4.2. For each $\mathcal{C} \subseteq \mathfrak{C}$, we have $\langle \langle \ell(\mathcal{C}) \rangle \rangle = \{ \langle \ell(\mathcal{C}') \rangle \mid \mathcal{C} \subseteq \mathcal{C}' \}$.

4.6 Computing $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ We now proceed by computing $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$. We start with its node constraint.

LEMMA 4.6. The node constraint $\mathcal{N}_{\overline{\mathcal{R}}}$ of $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ consists of all configurations of the form $\langle \langle \ell(\mathcal{C}) \rangle \rangle^{\Delta-|\mathcal{C}|+1} \langle \langle \ell(\emptyset) \rangle \rangle^{|\mathcal{C}|-1}$ for some $\emptyset \neq \mathcal{C} \subseteq \mathfrak{C}$.

Proof. Let \mathcal{N} denote the set of all configurations of the form $\langle \langle \ell(\mathcal{C}) \rangle \rangle^{\Delta-|\mathcal{C}|+1} \langle \langle \ell(\emptyset) \rangle \rangle^{|\mathcal{C}|-1}$ for some $\emptyset \neq \mathcal{C} \subseteq \mathfrak{C}$. We need to show that $\mathcal{N}_{\overline{\mathcal{R}}} = \mathcal{N}$.

We start by showing that each configuration from \mathcal{N} can be relaxed to some configuration from $\mathcal{N}_{\overline{\mathcal{R}}}$. Let $\langle \langle \ell(\mathcal{C}) \rangle \rangle^{\Delta-|\mathcal{C}|+1} \langle \langle \ell(\emptyset) \rangle \rangle^{|\mathcal{C}|-1}$ be an arbitrary configuration from \mathcal{N} , and consider some arbitrary choice $(\langle \ell(\mathcal{C}'_1) \rangle, \dots, \langle \ell(\mathcal{C}'_{\Delta}) \rangle) \in \langle \langle \ell(\mathcal{C}) \rangle \rangle^{\Delta-|\mathcal{C}|+1} \times \langle \langle \ell(\emptyset) \rangle \rangle^{|\mathcal{C}|-1}$. By Corollary 4.2, we know that, for each $1 \leq k \leq \Delta - |\mathcal{C}| + 1$, we have $\mathcal{C} \subseteq \mathcal{C}'_k$. Hence, by Corollary 4.1, we can pick one label from each of the sets in $(\langle \ell(\mathcal{C}'_1) \rangle, \dots, \langle \ell(\mathcal{C}'_{\Delta}) \rangle)$ such that the resulting configuration is $\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \ell(\emptyset)^{|\mathcal{C}|-1}$. Since $\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \ell(\emptyset)^{|\mathcal{C}|-1} \in \mathcal{N}_{\Pi}$, it follows, by the definition of $\mathcal{N}_{\overline{\mathcal{R}}}$ (and $\mathcal{N}_{\mathcal{R}}$), that $\langle \langle \ell(\mathcal{C}) \rangle \rangle^{\Delta-|\mathcal{C}|+1} \langle \langle \ell(\emptyset) \rangle \rangle^{|\mathcal{C}|-1}$ can be relaxed to some configuration from $\mathcal{N}_{\overline{\mathcal{R}}}$. We conclude that each configuration from \mathcal{N} can be relaxed to some configuration from $\mathcal{N}_{\overline{\mathcal{R}}}$. Note that this does not show yet that $\mathcal{N} \subseteq \mathcal{N}_{\overline{\mathcal{R}}}$; we will come back to this goal later.

First, we show that $\mathcal{N}_{\overline{\mathcal{R}}} \subseteq \mathcal{N}$. Let $\mathcal{B} = \mathcal{B}_1 \dots \mathcal{B}_{\Delta}$ be an arbitrary configuration from $\mathcal{N}_{\overline{\mathcal{R}}}$. We claim that \mathcal{B} can be relaxed to some configuration from \mathcal{N} . For a contradiction, assume that the claim is false, i.e., \mathcal{B} cannot be relaxed to any configuration from \mathcal{N} . Observe that \mathcal{B}_k is right-closed for each $1 \leq k \leq \Delta$, by Observation 3.1.

For each $i \in \mathfrak{C}$, define $R_i := \mathfrak{C} \setminus \{i\}$. Moreover, we will make use of the bipartite graph $\overline{G} = (\overline{V} \cup \overline{W}, \overline{E})$ obtained by defining $\overline{V} := \{\mathcal{B}_1, \dots, \mathcal{B}_{\Delta}\}$ and $\overline{W} := \{R_i \mid i \in \mathfrak{C}\}$, and setting \overline{E} to be the set of all edges $\{\mathcal{B}_k, R_i\}$ satisfying $\langle \ell(R_i) \rangle \in \mathcal{B}_k$. Note that, for any distinct k, k' , we consider \mathcal{B}_k and $\mathcal{B}_{k'}$ to be different vertices in \overline{V} even if $\mathcal{B}_k = \mathcal{B}_{k'}$.

Consider any arbitrary subset $\emptyset \neq \mathcal{C} \subseteq \mathfrak{C}$. By our assumption, configuration \mathcal{B} cannot be relaxed to the configuration $\langle\langle \ell(\mathcal{C}) \rangle\rangle^{\Delta-|\mathcal{C}|+1} \langle\langle \ell(\emptyset) \rangle\rangle^{|\mathcal{C}|-1}$, which implies that there are at least $|\mathcal{C}|$ indices $k \in \{1, \dots, \Delta\}$ satisfying $\mathcal{B}_k \not\subseteq \langle\langle \ell(\mathcal{C}) \rangle\rangle$, since $\langle\langle \ell(\emptyset) \rangle\rangle = \Sigma_{\mathcal{R}}$ (by Corollary 4.2). By Corollary 4.2, it follows that there are at least $|\mathcal{C}|$ indices $k \in \{1, \dots, \Delta\}$ such that there exists some $\mathcal{C}' \not\supseteq \mathcal{C}$ satisfying $\langle\langle \ell(\mathcal{C}') \rangle\rangle \in \mathcal{B}_k$. Since, for each $\mathcal{C}' \not\supseteq \mathcal{C}$, there exists some $i \in \mathcal{C}$ satisfying $i \notin \mathcal{C}'$, it follows, by the right-closedness of the \mathcal{B}_k (and Lemma 4.5), that there are at least $|\mathcal{C}|$ indices $k \in \{1, \dots, \Delta\}$ such that there exists some $i \in \mathcal{C}$ satisfying $\langle\langle \ell(R_i) \rangle\rangle \in \mathcal{B}_k$.

Recall the definition of $\overline{\mathcal{G}}$. By the above discussion, we conclude that for any arbitrary subset $\emptyset \neq \mathcal{C} \subseteq \mathfrak{C}$, the vertex set $\{R_i \mid i \in \mathcal{C}\} \subseteq \overline{W}$ has at least $|\mathcal{C}|$ neighbors in \overline{V} . Hence, we can apply Theorem 3.3 (i.e., Hall's marriage theorem) to $\overline{\mathcal{G}}$ and obtain a function $f: \overline{W} \rightarrow \overline{V}$ such that $f(R_i) \neq f(R_{i'})$ for any $i \neq i'$ and $\langle\langle \ell(R_i) \rangle\rangle \in f(R_i)$ for any $i \in \mathfrak{C}$. This implies that there is some choice $(A_1, \dots, A_\Delta) \in \mathcal{B}_1 \times \dots \times \mathcal{B}_\Delta$ such that the configuration $A_1 \dots A_\Delta$ is a permutation of the configuration $\langle\langle \ell(R_1) \rangle\rangle \dots \langle\langle \ell(R_\Delta) \rangle\rangle$. Therefore, by the definition of $\mathcal{N}_{\overline{\mathcal{R}}}$, there exists some choice $(L_1, \dots, L_\Delta) \in \langle\langle \ell(R_1) \rangle\rangle \times \dots \times \langle\langle \ell(R_\Delta) \rangle\rangle$ such that $L_1 \dots L_\Delta \in \mathcal{N}_{\Pi}$. Observe that, for each $\emptyset \neq \mathcal{C} \subseteq \mathfrak{C}$, there are at least $|\mathcal{C}|$ indices $i \in \mathfrak{C}$ (namely all colors i contained in \mathcal{C}) such that $\ell(\mathcal{C}) \notin \langle\langle \ell(R_i) \rangle\rangle$, by Corollary 4.1. It follows that, for each $\emptyset \neq \mathcal{C} \subseteq \mathfrak{C}$, the configuration $\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \ell(\emptyset)^{|\mathcal{C}|-1}$ is not a permutation of $L_1 \dots L_\Delta$. By the definition of \mathcal{N}_{Π} , this yields a contradiction to the fact that $L_1 \dots L_\Delta \in \mathcal{N}_{\Pi}$, and proves the claim.

Hence, each configuration from $\mathcal{N}_{\overline{\mathcal{R}}}$ can be relaxed to some configuration from \mathcal{N} . Recall that, as shown before, each configuration from \mathcal{N} can be relaxed to some configuration from $\mathcal{N}_{\overline{\mathcal{R}}}$. Combining these two insights, we obtain that any arbitrary configuration $\mathcal{B}_1 \in \mathcal{N}_{\overline{\mathcal{R}}}$ can be relaxed to some configuration $\mathcal{B}_2 \in \mathcal{N}$, which in turn can be relaxed to some configuration $\mathcal{B}_3 \in \mathcal{N}_{\overline{\mathcal{R}}}$. From (the maximality condition in) the definition of $\mathcal{N}_{\overline{\mathcal{R}}}$ it follows that $\mathcal{B}_1 = \mathcal{B}_3$, which implies that also $\mathcal{B}_1 = \mathcal{B}_2$. Hence, $\mathcal{B}_1 \in \mathcal{N}$, which implies that $\mathcal{N}_{\overline{\mathcal{R}}} \subseteq \mathcal{N}$, as desired.

We conclude the proof by showing that also $\mathcal{N} \subseteq \mathcal{N}_{\overline{\mathcal{R}}}$. We start by observing that for any two nonempty subsets $\mathcal{C}, \mathcal{C}' \subseteq \mathfrak{C}$ satisfying $\mathcal{C} \neq \mathcal{C}'$, the configuration $\langle\langle \ell(\mathcal{C}) \rangle\rangle^{\Delta-|\mathcal{C}|+1} \langle\langle \ell(\emptyset) \rangle\rangle^{|\mathcal{C}|-1}$ cannot be relaxed to $\langle\langle \ell(\mathcal{C}') \rangle\rangle^{\Delta-|\mathcal{C}'|+1} \langle\langle \ell(\emptyset) \rangle\rangle^{|\mathcal{C}'|-1}$. If $\mathcal{C}' \subseteq \mathcal{C}$, then $|\mathcal{C}'| < |\mathcal{C}|$, which implies that there are strictly fewer than $|\mathcal{C}| - 1$ sets in $\langle\langle \ell(\mathcal{C}') \rangle\rangle^{\Delta-|\mathcal{C}'|+1} \langle\langle \ell(\emptyset) \rangle\rangle^{|\mathcal{C}'|-1}$ that are a superset of $\langle\langle \ell(\emptyset) \rangle\rangle$, since $\langle\langle \ell(\mathcal{C}') \rangle\rangle \not\supseteq \langle\langle \ell(\emptyset) \rangle\rangle$, by Corollary 4.2. If $\mathcal{C}' \not\subseteq \mathcal{C}$, then $\langle\langle \ell(\mathcal{C}) \rangle\rangle \not\subseteq \langle\langle \ell(\mathcal{C}') \rangle\rangle$ (by Corollary 4.2), which implies that no set in $\langle\langle \ell(\mathcal{C}) \rangle\rangle^{\Delta-|\mathcal{C}|+1} \langle\langle \ell(\emptyset) \rangle\rangle^{|\mathcal{C}|-1}$ is a subset of $\langle\langle \ell(\mathcal{C}') \rangle\rangle$, as also $\langle\langle \ell(\emptyset) \rangle\rangle \not\subseteq \langle\langle \ell(\mathcal{C}') \rangle\rangle$ (as seen above).

Hence, no configuration from \mathcal{N} can be relaxed to a different configuration from \mathcal{N} . Recall (again) that each configuration from \mathcal{N} can be relaxed to some configuration from $\mathcal{N}_{\overline{\mathcal{R}}}$ and each configuration from $\mathcal{N}_{\overline{\mathcal{R}}}$ can be relaxed to some configuration from \mathcal{N} . It follows that if there is some configuration from \mathcal{N} that is not contained in $\mathcal{N}_{\overline{\mathcal{R}}}$, then it can be relaxed to a different configuration from \mathcal{N} , by first relaxing it to a (necessarily different) configuration from $\mathcal{N}_{\overline{\mathcal{R}}}$ and then relaxing the obtained configuration to some configuration from \mathcal{N} . As no configuration from \mathcal{N} can be relaxed to a different configuration from \mathcal{N} (as shown above), we conclude that each configuration from \mathcal{N} is also contained in $\mathcal{N}_{\overline{\mathcal{R}}}$. Thus, $\mathcal{N} \subseteq \mathcal{N}_{\overline{\mathcal{R}}}$, as desired. It follows that $\mathcal{N}_{\overline{\mathcal{R}}} = \mathcal{N}$. \square

We now describe what is the label set $\Sigma_{\overline{\mathcal{R}}}$ of the problem $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$, and we then compute $\mathcal{E}_{\overline{\mathcal{R}}}$.

LEMMA 4.7. *The set $\Sigma_{\overline{\mathcal{R}}}$ of output labels of $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ is given by $\{\langle\langle \ell(\mathcal{C}) \rangle\rangle \mid \mathcal{C} \in \mathfrak{C}\}$*

Moreover, the hyperedge constraint $\mathcal{E}_{\overline{\mathcal{R}}}$ of $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ consists of all configurations $\langle\langle \ell(\mathcal{C}_1) \rangle\rangle \dots \langle\langle \ell(\mathcal{C}_r) \rangle\rangle$ of labels from $\Sigma_{\overline{\mathcal{R}}}$ such that for each color $i \in \mathfrak{C}$, there is at least one index $1 \leq j \leq r$ satisfying $i \notin \mathcal{C}_j$.

Proof. From Lemma 4.6 (and the fact that $|\mathfrak{C}| = \Delta \geq 2$), it follows directly that $\Sigma_{\overline{\mathcal{R}}}$ is as given in the lemma, by the definition of $\Sigma_{\overline{\mathcal{R}}}$. Hence, what remains is to show that the hyperedge constraint $\mathcal{E}_{\overline{\mathcal{R}}}$ is as given in the lemma. Let \mathcal{E} denote the set of all configurations $\langle\langle \ell(\mathcal{C}_1) \rangle\rangle \dots \langle\langle \ell(\mathcal{C}_r) \rangle\rangle$ of labels from $\Sigma_{\overline{\mathcal{R}}}$ such that for each color $i \in \mathfrak{C}$, there is at least one index $1 \leq j \leq r$ satisfying $i \notin \mathcal{C}_j$. We need to show that $\mathcal{E}_{\overline{\mathcal{R}}} = \mathcal{E}$.

We first show that $\mathcal{E} \subseteq \mathcal{E}_{\overline{\mathcal{R}}}$. Let $\langle\langle \ell(\mathcal{C}_1) \rangle\rangle \dots \langle\langle \ell(\mathcal{C}_r) \rangle\rangle$ be an arbitrary configuration from \mathcal{E} . Since for each color $i \in \mathfrak{C}$, there is at least one index $1 \leq j \leq r$ satisfying $i \notin \mathcal{C}_j$ (by the definition of \mathcal{E}), there also exists a collection $\mathcal{C}'_1, \dots, \mathcal{C}'_r$ of subsets of \mathfrak{C} such that, for each $1 \leq j \leq r$, we have $\mathcal{C}_j \subseteq \mathcal{C}'_j$ and, for each $i \in \mathfrak{C}$, there is exactly one index $1 \leq j \leq r$ satisfying $i \notin \mathcal{C}'_j$. By Corollary 4.2, we know that $(\langle\langle \ell(\mathcal{C}'_1) \rangle\rangle, \dots, \langle\langle \ell(\mathcal{C}'_r) \rangle\rangle) \in \langle\langle \ell(\mathcal{C}_1) \rangle\rangle \times \dots \times \langle\langle \ell(\mathcal{C}_r) \rangle\rangle$; since $\langle\langle \ell(\mathcal{C}'_1) \rangle\rangle \dots \langle\langle \ell(\mathcal{C}'_r) \rangle\rangle \in \mathcal{E}_{\overline{\mathcal{R}}}$ (by the definition of the \mathcal{C}'_j), it follows that $\langle\langle \ell(\mathcal{C}_1) \rangle\rangle \dots \langle\langle \ell(\mathcal{C}_r) \rangle\rangle \in \mathcal{E}_{\overline{\mathcal{R}}}$, by the definition of $\mathcal{E}_{\overline{\mathcal{R}}}$. Hence, $\mathcal{E} \subseteq \mathcal{E}_{\overline{\mathcal{R}}}$, as desired.

Now, we show that $\mathcal{E}_{\overline{\mathcal{R}}} \subseteq \mathcal{E}$. Let $\langle\langle \ell(\mathcal{C}_1) \rangle\rangle \dots \langle\langle \ell(\mathcal{C}_r) \rangle\rangle$ be an arbitrary configuration from $\mathcal{E}_{\overline{\mathcal{R}}}$. By the definitions of $\mathcal{E}_{\overline{\mathcal{R}}}$ and $\mathcal{E}_{\mathcal{R}}$, we know that there is some choice $(\langle\langle \ell(\mathcal{C}'_1) \rangle\rangle, \dots, \langle\langle \ell(\mathcal{C}'_r) \rangle\rangle) \in \langle\langle \ell(\mathcal{C}_1) \rangle\rangle \times \dots \times \langle\langle \ell(\mathcal{C}_r) \rangle\rangle$ such that, for each color $i \in \mathfrak{C}$, there is exactly one index $1 \leq j \leq r$ satisfying $i \notin \mathcal{C}'_j$. Since, for each $1 \leq j \leq r$, $\langle\langle \ell(\mathcal{C}'_j) \rangle\rangle \in \langle\langle \ell(\mathcal{C}_j) \rangle\rangle$ implies $\mathcal{C}_j \subseteq \mathcal{C}'_j$ (by Corollary 4.2), it follows that, for each color $i \in \mathfrak{C}$, there is at least one index $1 \leq j \leq r$

satisfying $i \notin C_j$. This implies that $\langle\langle\ell(C_1)\rangle\rangle \dots \langle\langle\ell(C_r)\rangle\rangle \in \mathcal{E}$, by the definition of \mathcal{E} . Hence, $\mathcal{E}_{\overline{\mathcal{R}}} \subseteq \mathcal{E}$, and we obtain $\mathcal{E}_{\overline{\mathcal{R}}} = \mathcal{E}$, as desired. \square

4.7 Renaming We now show that, if we rename the labels of $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ correctly, we obtain that $\overline{\mathcal{R}}(\mathcal{R}(\Pi)) = \Pi$.

LEMMA 4.1. *The problem $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ is equivalent to Π .*

Proof. Consider the following renaming:

$$\langle\langle\ell(C)\rangle\rangle \text{ for each } \emptyset \neq C \subseteq \mathfrak{C} \quad \rightarrow \quad \ell(C)$$

Observe that, under this renaming, $\mathcal{N}_{\overline{\mathcal{R}}}$ becomes equal to \mathcal{N}_{Π} , and $\mathcal{E}_{\overline{\mathcal{R}}}$ becomes equal to \mathcal{E}_{Π} . \square

5 Lower Bound for Hypergraph MM

In this section, we prove an $\Omega(\min\{\Delta r, \log_{\Delta r} n\})$ -round deterministic and an $\Omega(\min\{\Delta r, \log_{\Delta r} \log n\})$ -round randomized lower bound for the problem of computing a maximal matching on hypergraphs.

A Sequence of Problems. In order to prove a lower bound for hypergraph coloring, it was enough to provide a single problem Π that is a relaxation of hypergraph coloring and prove that $\overline{\mathcal{R}}(\mathcal{R}(\Pi)) = \Pi$, that is, that Π is a fixed point. This approach can only work for “hard” problems, that is, problems that cannot be solved in $O(f(\Delta, r) + \log^* n)$ for any function f . Instead, for problems solvable with this runtime, like hypergraph maximal matching, we have to follow a different approach.

The idea is to design a sequence of problems Π_0, Π_1, \dots such that Π_0 is essentially a relaxed version of hypergraph MM, and any other problem Π_i in the sequence is a relaxation of $\overline{\mathcal{R}}(\mathcal{R}(\Pi_{i-1}))$. Roughly speaking, this implies that each problem in the sequence can be solved at least one round faster than the previous problem—now, all we have to do is to prove that it takes $\Omega(\Delta r)$ steps in the sequence to reach a 0-round solvable problem. While this is a simplified outline that disregards certain technicalities (we, e.g., have to ensure that the sizes of the label sets of the problems in the sequence do not grow too fast), there are known techniques in the round elimination framework that take care of these technicalities, i.e., the real challenge lies in designing the aforementioned sequence and proving that it satisfies the desired properties.

A Parametrized Family of Problems. In order to be able to prove these properties, it is useful if the problems in the sequence come from some parameterized problem family. In Section 5.2, we describe such a problem family that we will use to prove the lower bound for hypergraph MM—all problems in the desired problem sequence will come from this family. For technical reasons, we will assume throughout Section 5 that $\Delta \geq 3$ and $r \geq 3$. Note that for $\Delta = 2$, hypergraph MM is essentially⁷ the same as MM (on graphs), while for $r = 2$, hypergraph MM is essentially the same as MIS (on graphs). In particular, the lower bounds from [5] directly imply an $\Omega(\min\{r, \log_r n\})$ -round deterministic and an $\Omega(\min\{r, \log_r \log n\})$ -round randomized lower bound for hypergraph MM with $\Delta = 2$, and an $\Omega(\min\{\Delta, \log_{\Delta} n\})$ -round deterministic and an $\Omega(\min\{\Delta, \log_{\Delta} \log n\})$ -round randomized lower bound for hypergraph MM with $r = 2$.⁸

A Summary of our Goals. We proceed as follows. The proof presented for hypergraph coloring can be seen as a simplified version of the proof presented in this section. Hence, we start in Section 5.1 by highlighting the analogies, and the differences, between the hypergraph maximal matching lower bound proof and the hypergraph coloring lower bound proof.

Then, in Section 5.2, we describe the problem family, characterized by two parameters. We will prove that, for each problem Π in the family satisfying some conditions on these parameters, a relaxation of $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ is also in the family. The proof of this fact is going to be a tedious case analysis, and we will defer it to the end of the section.

⁷Note that $\Delta = 2$ allows for nodes of degree 1. Hence by interpreting each node in the hypergraph MM problem as a (hyper)edge and each hyperedge as a node, in the case of $\Delta = 2$ we have a problem that is at least as hard as MM on graphs. It is also easy to see that these additional hyperedges (= nodes in the hypergraph MM problem) of rank 1 do not really make the MM problem harder as they can simply be taken care of in one round of computation after executing a proper MM algorithm on all rank-2 hyperedges. Analogously, in the case of $r = 2$, hypergraph MM formally allows hyperedges of rank 1, but these do not make the problem (asymptotically) harder as they can be taken care of in one round of computation after executing an MIS algorithm.

⁸To obtain these lower bounds already on trees, we additionally require [7] for the case $r = 2$.

In Section 5.3, we will then show that, for certain values of the parameters, Π cannot be solved in 0 rounds in the port numbering model. We will then put things together, by showing that we can construct a sequence of problems Π_0, \dots, Π_t , satisfying that:

- All problems are in the family and not 0-round solvable;
- Each problem Π_{i+1} is a relaxation of $\overline{\mathcal{R}}(\mathcal{R}(\Pi_i))$;
- The length of the sequence t is in $\Omega(\Delta r)$;
- The first problem is a relaxation of the hypergraph MM problem.

We will then get our claimed lower bound by applying Theorem 3.2.

5.1 Analogies and Differences with the Hypergraph Coloring Lower Bound

Recap of the Coloring Relaxations. As already mentioned while discussing the hypergraph coloring lower bound, there is some intuition on why the relaxations that we applied on hypergraph coloring are able to give a fixed point. The idea, there, was that, by computing $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$, for Π equal to the hypergraph coloring problem, we obtain a problem where nodes are allowed to use multiple colors at once, and they are rewarded for using more colors. The reward is that they can mark some amount of incident hyperedges, where the amount of marked hyperedges depends on the amount of colors, and marked hyperedges need to satisfy more relaxed constraints. In $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ (which we did not even present, as it would result in a very unnatural and hard to describe problem) there are many different ways to mark the edges, that give different guarantees. In the relaxation that we presented, we essentially relaxed the constraints so that there is a single possible way to mark the hyperedges.

The Behavior of Hypergraph Maximal Matchng. Informally, by applying round elimination on hypergraph maximal matching for k times, what we obtain is essentially a problem that can be decomposed into three parts: the original problem, a hypergraph k -colorful coloring, and some unnatural part of much larger size. This can be confirmed experimentally, but since the unnatural part seems to have a size equal to a power tower of height k , it makes unfeasible to actually give an explicit form of it.

Our Problem Sequence. What we would like to do, is to relax the coloring part of the problem, in order to make the unnatural part disappear. This is essentially the approach used in [7]. Unfortunately, as explained in Section 2, if we try to do this, we cannot prevent the colors from growing by 1 at each step, and hence, while we would like to obtain a sequence of problems of length $\Omega(\Delta r)$, after only $\Delta + 1$ steps of round elimination we would have $\Delta + 1$ colors, but the fixed point only tolerates up to Δ colors, and in fact we would obtain a problem that is 0 round solvable.

Hence, this suggests that, in our problem sequence, we cannot let the colors just grow to $\Omega(\Delta r)$, and that we have to keep them bounded to at most Δ (for technical reasons, we will actually restrict them to $\Delta - 1$). How can we perform $\Omega(\Delta r)$ steps of round elimination and keep the colors bounded to at most Δ , if at each step the obtained problem allows the node to use one additional color? The idea is to perform simplifications that *reduce* the amount of colors: this seems contradictory, how can we make the problem easier if nodes are allowed to use less colors? In order to achieve this, we remove the colors at the cost of relaxing the requirements of the remaining colors from the point of view of the hyperedge constraint. In particular, a problem in our family is described as a vector of length at most $\Delta - 1$, where each position represents a color, and for each color we specify how many nodes incident on the same hyperedge are allowed to use that color. The idea is that we can get rid of a color at the cost of increasing some values in this vector.

Comparison with Hypergraph Coloring. Summarizing, in order to obtain a fixed point for hypergraph coloring, we had to relax the problem to allow nodes to use sets of colors, and reward nodes for using more colors by allowing nodes to mark hyperedges, such that marked hyperedges are always happy.

In the case of hypergraph maximal matching, we do something very similar. A problem in our family is essentially defined very similarly as hypergraph coloring, and the differences are the following:

- We have at most $\Delta - 1$ colors;
- We also allow the original configurations allowed by hypergraph maximal matching;
- There is a vector describing how hard it is to use a color on a hyperedge.

Unfortunately, since now the labels of the coloring part and the matching part can mix in the allowed configurations in nontrivial ways, it is not possible to give proper intuition behind some configurations allowed by the hyperedge constraint.

What we are going to prove is that if we take a problem Π in this family, we get that $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$ can be relaxed to a different problem of the family, where some values in the vector increase. For technical reasons, this is not the full description of the problems in our family: in the problem sequence that we define, at each step, we have to take one of the $\Delta - 1$ colors and treat it differently (and the unfamiliar configuration $D^{\Delta-1}X$ present in the node constraint presented later is related to it).

5.2 The Problem Family Each problem $\Pi(z, s)$ in the family is characterized by two parameters z, s , where $z = (z_1, \dots, z_k)$ is a vector of $2 \leq k \leq \Delta - 1$ nonnegative integers $z_i \leq r - 1$ and $s \in \{1, \dots, k\}$. We call k the *length* of z , denoted by $\text{len}(z)$. Intuitively, some labels of $\Pi(z, s)$ can be seen as sets of colors, and the length of z tells us the number of colors, and for each color i , parameter z_i describes how often i (or a color set containing i) can appear in the same hyperedge configuration. Parameter s singles out one of the $\text{len}(z)$ colors that behaves a bit differently than the other colors. In the following, we describe $\Pi(z, s)$ formally.

The Label Set. To describe the space of colors discussed above, define $\mathcal{C}(z) := \{1, \dots, \text{len}(z)\}$. The set $\Sigma(z, s)$ of output labels of $\Pi(z, s)$ is given by $\Sigma(z, s) = \{D, M, P, U, X\} \cup \{\ell(\mathcal{C}) \mid \emptyset \neq \mathcal{C} \subseteq \mathcal{C}(z)\}$ (where we use the expression $\ell(\mathcal{C})$ instead of \mathcal{C} to clearly distinguish between color sets and labels).

The Node Constraint. We denote the node constraint of $\Pi(z, s)$ by $\mathcal{N}(z, s)$. It is given by the following configurations.

- M^Δ
- $P \cup \Delta - 1$
- $D^{\Delta-1}X$
- $\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \cup |\mathcal{C}|-1$ for each $\emptyset \neq \mathcal{C} \subseteq \mathcal{C}(z)$

The Hyperedge Constraint. We denote the hyperedge constraint of $\Pi(z, s)$ by $\mathcal{E}(z, s)$. It is given by all configurations $L_1 \dots L_r$ satisfying at least one of the following three conditions.

1. There is some index $1 \leq j \leq r$ such that $L_j = M$ and $L_{j'} \notin \{D, M\}$ for each $j' \neq j$.
2. There are two distinct indices $j, j' \in \{1, \dots, r\}$ such that $L_j = X$, $L_{j'}$ is arbitrary, and $L_{j''} \notin \{D, M\}$ for each $j'' \in \{1, \dots, r\} \setminus \{j, j'\}$.
3. All of the following properties hold.
 - (a) $L_j \neq P$ for each $1 \leq j \leq r$.
 - (b) There is at most one index $j \in \{1, \dots, r\}$ such that $L_j \in \{D, M\}$.
 - (c) There are at most z_s indices $j \in \{1, \dots, r\}$ such that $L_j = D$ or $L_j = \ell(\mathcal{C})$ for some color set \mathcal{C} containing color s .
 - (d) For each $i \in \mathcal{C}(z)$ satisfying $i \neq s$, there are at most z_i indices $j \in \{1, \dots, r\}$ such that $L_j = \ell(\mathcal{C})$ for some color set \mathcal{C} containing color i .

The Relation between the Problems. We will prove that the problems in the defined family are related in the following way.

LEMMA 5.1. *Let $z = (z_1, \dots, z_k)$ be a vector of $2 \leq k \leq \Delta - 1$ nonnegative integers $z_i \leq r - 1$ and let $s \in \{1, \dots, k\}$. Let q be an integer satisfying $q \neq s$, $q \in \{1, \dots, k\}$, and $z_q \leq r - 2$. Then, the problem $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$ can be relaxed to $\Pi(z', q)$, for $z' = (z'_1, \dots, z'_k)$, where $k = \text{len}(z)$, $z'_i = z_i + 1$ if $i = q$, and $z'_i = z_i$ otherwise.*

Since the proof of this statement is a tedious case analysis, we defer its proof to the end of the section. We now use this statement to derive our lower bounds.

5.3 Proving the Lower Bounds In this section, we prove our lower bounds. We start by showing that all problems in our problem family cannot be solved in 0 rounds in the port numbering model.

LEMMA 5.2. *Let $z = (z_1, \dots, z_k)$ be a vector of $2 \leq k \leq \Delta - 1$ nonnegative integers $z_i \leq r - 1$, and let $s \in \{1, \dots, k\}$. The problem $\Pi(z, s)$ cannot be solved in 0 rounds in the deterministic port numbering model.*

Proof. All nodes that run a 0-round algorithm have the same view (they just know their degree, the size of the graph, and the parameters Δ and r). Hence, any deterministic 0-round algorithm uses the same configuration in $\mathcal{N}(z, s)$ on all nodes. Thus, in order to prove the statement, we consider all possible configurations and we show that they cannot be properly used in a 0-round algorithm.

- M^Δ : consider r nodes connected to the same hyperedge. They would all output M on such hyperedge, that would hence have the configuration $M^r \notin \mathcal{E}(z, s)$.
- $PU^{\Delta-1}$: since the algorithm is deterministic, then all nodes must output P on the same port. W.l.o.g., let this port be the port number 1. If we connect the port 1 of r nodes to the same hyperedge, then the hyperedge would have the configuration $P^r \notin \mathcal{E}(z, s)$.
- $D^{\Delta-1}X$: as before, we can obtain a hyperedge labeled with the configuration $D^r \notin \mathcal{E}(z, s)$.
- $\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \cup^{|\mathcal{C}|-1}$ for each $\emptyset \neq \mathcal{C} \subseteq \mathcal{C}(z)$: since $|\mathcal{C}| \leq \Delta - 1$, then this configuration contains $\ell(\mathcal{C})$ at least twice. Hence, we can obtain a hyperedge labeled with the configuration $\ell(\mathcal{C})^r$, which is not in $\mathcal{E}(z, s)$ by the assumption that $z_i \leq r - 1$ for all $i \in \mathcal{C}$.

□

We are now ready to prove our main result of this section.

THEOREM 5.1. *Assume that $\Delta \geq 2$ and $r \geq 2$. Then any deterministic distributed algorithm in the LOCAL model for computing a maximal matching in hypergraphs with maximum degree Δ , rank r , and n nodes requires $\Omega(\min\{\Delta r, \log_{\Delta r} n\})$ rounds.⁹ Any randomized such algorithm requires at least $\Omega(\min\{\Delta r, \log_{\Delta r} \log n\})$ rounds. Moreover, our lower bounds hold already on Δ -regular r -uniform linear hypertrees.*

Proof. We show that we can build a problem sequence satisfying Theorem 3.2, such that the asymptotic complexity of Π_0 is not higher than the one of the hypergraph maximal matching problem.

Consider a problem sequence $\Pi_0, \Pi_1, \dots, \Pi_{(\Delta-1)(r-2)}$ such that $\Pi_0 = \Pi((1, 1, \dots, 1), \Delta - 1)$, $\Pi_{(\Delta-1)(r-2)} = \Pi((r-1, r-1, \dots, r-1), \Delta - 1)$, and for any two subsequent problems Π_j, Π_{j+1} , there exist vectors z, z' of length $\Delta - 1$ and colors $s, s' \in \{1, \dots, \Delta - 1\}$ such that

- $\Pi_j = \Pi(z, s)$ and $\Pi_{j+1} = \Pi(z', s')$,
- $s \neq s'$,
- $z'_{s'} = z_{s'} + 1$, and
- $z'_i = z_i$ for all $1 \leq i \leq \Delta - 1$ satisfying $i \neq s'$.

Clearly, such a problem sequence exists as we can go from the length- $(\Delta - 1)$ vector $(1, 1, \dots, 1)$ to the length- $(\Delta - 1)$ vector $(r - 1, r - 1, \dots, r - 1)$ in $(\Delta - 1)(r - 2)$ steps of increasing a single entry (that is a different entry than in the previous step) by 1 (such that we choose to increase the $(\Delta - 1)$ th entry in the last but not the first step).

Observe that, each problem of the sequence is not 0-round solvable in the port numbering model by Lemma 5.2. Also, the number of labels of each problem is bounded by $2^{O(\Delta r)}$. Also, by Lemma 5.1, each problem Π_{i+1} is a relaxation of $\overline{\mathcal{R}}(\mathcal{R}(\Pi_i))$, and also the intermediate problem, $\mathcal{R}(\Pi_i)$, has a number of labels bounded by $2^{O(\Delta r)}$. Hence, Theorem 3.2 applies, and since $\log_{\Delta r} \log f(\Delta r) = O(1)$, we obtain a lower bound for Π_0 of

⁹As lower bounds in several parameters are sometimes difficult to understand, we would also like to provide, as an example, the exact quantification of this statement, which is as follows: There is a constant c such that, for any $\Delta \geq 2$ and $r \geq 2$, and any deterministic algorithm \mathcal{A} , there are infinitely many hypergraphs with maximum degree Δ , rank r , and n nodes on which \mathcal{A} has a runtime of at least $c(\min\{\Delta r, \log_{\Delta r} n\})$ rounds.

$\Omega(\min\{\Delta \cdot r, \log_{\Delta r} n\})$ in the LOCAL model for deterministic algorithms, and $\Omega(\min\{\Delta \cdot r, \log_{\Delta r} \log n\})$ for randomized ones.

We now prove that, given an algorithm A with complexity T for solving hypergraph maximal matching, we can turn it into an algorithm with complexity $O(T)$ for solving Π_0 , implying the claim. In order to solve Π_0 on some hypergraph H , we start by simulating A in the hypergraph H' obtained by reversing the role of nodes and hyperedges of H . This can be performed in $O(T)$ rounds. Then, given a solution for the hypergraph maximal matching problem on H' , we solve Π_0 as follows:

- Nodes of H (that is, hyperedges of H') that are in the matching, output M^Δ .
- Each node v not in the matching must have at least one neighboring hyperedge h that is incident to some other node u in the matching. Node v outputs P on h and U on all the other incident hyperedges.

Observe that, on the nodes, we only use configurations M^Δ and $PU^{\Delta-1}$, which are allowed by Π_0 . Then, on each hyperedge, we obtain that there is at most one M , and that a P is present only if there is also an M , while all the other labels are U . Hence, also on the hyperedges we obtain only configurations allowed by Π_0 . \square

5.4 Relations Between Labels In the rest of the section we prove Lemma 5.1. We start by proving a relation between the labels. From the description of the hyperedge constraint, we can extract the strength relations w.r.t. $\mathcal{E}(z, s)$ by simply checking for any two labels $L, L' \in \Sigma(z, s)$ whether every hyperedge configuration containing L remains a configuration in $\mathcal{E}(z, s)$ if we replace L by L' . An example is shown in Figure 2.




Figure 2: The edge diagram of the problems in the family satisfying $s = 2$ and $\text{len}(z) = 3$. Numbers correspond to color sets, and $\ell(\cdot)$ is omitted for clarity.

LEMMA 5.3. *The following collection lists all strength relations (according to $\mathcal{E}(z, s)$) between distinct labels in $\Sigma(z, s)$.*

- $P < \ell(C)$ for each $\emptyset \neq C \subseteq \mathcal{C}(z)$
- $P < U$
- $\ell(C) < \ell(C')$ for any two $C, C' \in 2^{\mathcal{C}(z)} \setminus \{\emptyset\}$ satisfying $C' \subsetneq C$
- $\ell(C) < U$ for each $\emptyset \neq C \subseteq \mathcal{C}(z)$
- $D < M$
- $D < \ell(\{s\})$
- $D < U$
- $L < X$ for each $L \in \Sigma(z, s) \setminus \{X\}$

Proof. We start by comparing the strength of X with the strength of all other labels. The definition of $\mathcal{E}(z, s)$ implies that in any configuration $L_1 \dots L_r \in \mathcal{E}(z, s)$ there is at most one L_j with $L_j \in \{D, M\}$. Hence, replacing an arbitrary L_j with X will yield a configuration from $\mathcal{E}(z, s)$, by Condition 2, and we obtain $L \leq X$ for each

$L \in \Sigma(z, s) \setminus \{X\}$. To show that $X \not\leq L$ for each $L \in \Sigma(z, s) \setminus \{X\}$, it suffices to observe that $XMP^{r-2} \in \mathcal{E}(z, s)$, which implies $X \not\leq M$ (as replacing X with M does not yield a configuration from $\mathcal{E}(z, s)$), and $XP^{r-1} \in \mathcal{E}(z, s)$, which implies $X \not\leq L$ for each $L \in \Sigma(z, s) \setminus \{M, X\}$.

Next, we compare the strength of P with the strength of the labels from $\Sigma(z, s) \setminus \{P, X\}$. From the definition of $\mathcal{E}(z, s)$, we can see that the only configurations containing P are of the form $PML_1 \dots L_{r-2}$ or $PXL_1 \dots L_{r-2}$ where, in the former case, $L_j \notin \{D, M\}$ for each $1 \leq j \leq r-2$, and, in the latter case, $L_j \notin \{D, M\}$ for each $2 \leq j \leq r-2$. In either case, replacing P with some arbitrary label from $\Sigma(z, s) \setminus \{D, M\}$ yields a configuration contained in $\mathcal{E}(z, s)$, by Conditions 1 and 2. It follows that $P \leq L$ for each $L \in \Sigma(z, s) \setminus \{D, M\}$. Moreover, the configuration $MP^{r-1} \in \mathcal{E}(z, s)$ certifies that $P \not\leq D$ and $P \not\leq M$ since neither of the two configurations MDP^{r-2} and M^2P^{r-2} is contained in $\mathcal{E}(z, s)$, by Conditions 1, 2 and 3b. In order to show that $L \not\leq P$ for each $L \in \Sigma(z, s) \setminus \{P, X\}$, it suffices to observe that $LU^{r-1} \in \mathcal{E}(z, s)$ for each $L \in \Sigma(z, s) \setminus \{P, X\}$ (by Condition 3), while $PU^{r-1} \notin \mathcal{E}(z, s)$ (by Conditions 1, 2 and 3a).

We continue by comparing the strength of D and the strength of M with each other and with the strength of the labels from $\Sigma(z, s) \setminus \{D, M, P, X\}$. As already observed above, in any configuration $L_1 \dots L_r \in \mathcal{E}(z, s)$ there is at most one L_j with $L_j \in \{D, M\}$. Hence, we obtain $DMU^{r-2} \notin \mathcal{E}(z, s)$ and $M^2U^{r-2} \notin \mathcal{E}(z, s)$. Since, by Condition 1, $LMU^{r-2} \in \mathcal{E}(z, s)$ for each $L \in \Sigma(z, s) \setminus \{D, M\}$, it follows that $L \not\leq L'$ for each $L \in \Sigma(z, s) \setminus \{D, M\}$ and $L' \in \{D, M\}$. Moreover, since $MP^{r-1} \in \mathcal{E}(z, s)$ by Condition 1, but $LP^{r-1} \notin \mathcal{E}(z, s)$ for each $L \in \Sigma(z, s) \setminus \{M, P, X\}$ by Conditions 1, 2 and 3a, we obtain $M \not\leq L$ for each $L \in \Sigma(z, s) \setminus \{M, P, X\}$.

To show that $D \not\leq L$ for each $L \in \Sigma(z, s) \setminus \{D, M, P, U, X, \ell(\{s\})\}$, assume that $\Sigma(z, s) \setminus \{D, M, P, U, X, \ell(\{s\})\} \neq \emptyset$ (otherwise we are done) and consider some arbitrary $L \in \Sigma(z, s) \setminus \{D, M, P, U, X, \ell(\{s\})\}$. By the definition of $\Sigma(z, s)$, it follows that $L = \ell(\mathcal{C})$ for some $\mathcal{C} \in 2^{\mathcal{C}(z)} \setminus \{\emptyset, \{s\}\}$. Hence, there exists some $i \in \mathcal{C}(z)$ satisfying $i \neq s$ such that $i \in \mathcal{C}$; consider such an i . Consider the configuration $D\ell(\{i\})^{z_i}U^{r-z_i-1}$ (which is well-defined as the definition of problem $\Pi(z, s)$ specifies that $z_i \leq r-1$). By Condition 3, it is contained in $\mathcal{E}(z, s)$ whereas the configuration $\ell(\mathcal{C})\ell(\{i\})^{z_i}U^{r-z_i-1}$ is not contained in $\mathcal{E}(z, s)$, by Conditions 1, 2 and 3d. Therefore, $D \not\leq \ell(\mathcal{C})$, and we obtain that $D \not\leq L$ for each $L \in \Sigma(z, s) \setminus \{D, M, P, U, X, \ell(\{s\})\}$.

To show that $D \leq L$ for each $L \in \{M, U, \ell(\{s\})\}$, it suffices to observe that if some configuration $L_1 \dots L_r$ that contains D satisfies one of the six conditions (1, 2, 3a, 3b, 3c, 3d), then the same configuration where D is replaced by L satisfies the same condition, for each $L \in \{M, U, \ell(\{s\})\}$.

It remains to compare the strength of the labels in $\Sigma(z, s) \setminus \{D, M, P, U, X\}$ with each other and with the strength of U . Similarly to before, it is straightforward to check that if some configuration $L_1 \dots L_r$ that contains some label $\ell(\mathcal{C})$ satisfies one of the six conditions (1, 2, 3a, 3b, 3c, 3d), then the same configuration where \mathcal{C} is replaced by U satisfies the same condition (irrespective of the choice of \mathcal{C}). Hence, $L \leq U$ for each $L \in \Sigma(z, s) \setminus \{D, M, P, U, X\}$.

In order to show that $U \not\leq L$ for each $L \in \Sigma(z, s) \setminus \{D, M, P, U, X\}$, consider an arbitrary label $L \in \Sigma(z, s) \setminus \{D, M, P, U, X\}$. By the definition of $\Sigma(z, s)$, we have $L = \ell(\mathcal{C})$ for some $\emptyset \neq \mathcal{C} \subseteq \mathcal{C}(z)$. Let i be some arbitrary color in \mathcal{C} , and consider the configuration $\ell(\{i\})^{z_i}U^{r-z_i}$ (which contains at least one U since $z_i \leq r-1$). By Condition 3, this configuration is contained in $\mathcal{E}(z, s)$ whereas the configuration $\ell(\mathcal{C})\ell(\{i\})^{z_i}U^{r-z_i-1}$ (obtained by replacing one U by $\ell(\mathcal{C})$) is not contained in $\mathcal{E}(z, s)$, by Conditions 1, 2 and one of 3c, 3d (depending on whether $i = s$ or $i \neq s$). Hence, $U \not\leq L$ for each $L \in \Sigma(z, s) \setminus \{D, M, P, U, X\}$.

Finally, consider two distinct labels $L, L' \in \Sigma(z, s) \setminus \{D, M, P, U, X\}$. As above, we know that $L = \ell(\mathcal{C})$ and $L' = \ell(\mathcal{C}')$ for some $\mathcal{C}, \mathcal{C}' \in 2^{\mathcal{C}(z)} \setminus \{\emptyset\}$ satisfying $\mathcal{C} \neq \mathcal{C}'$. Consider first the case that $\mathcal{C} \not\subseteq \mathcal{C}'$, which implies that there exists some color $i \in \mathcal{C} \setminus \mathcal{C}'$. Similarly to before, we can observe that $\ell(\mathcal{C}')\ell(\{i\})^{z_i}U^{r-z_i-1} \in \mathcal{E}(z, s)$ by Condition 3 and the fact that $i \notin \mathcal{C}'$, whereas $\ell(\mathcal{C})\ell(\{i\})^{z_i}U^{r-z_i-1} \notin \mathcal{E}(z, s)$ by Condition 3c or 3d (depending on whether $i = s$ or $i \neq s$). Hence, $\ell(\mathcal{C}') \not\leq \ell(\mathcal{C})$.

Now, consider the other case, namely that $\mathcal{C} \subsetneq \mathcal{C}'$. Again, we observe that if some configuration $L_1 \dots L_r$ that contains \mathcal{C}' satisfies one of the six conditions (1, 2, 3a, 3b, 3c, 3d), then the same configuration where the \mathcal{C}' is replaced by \mathcal{C} satisfies the same condition. Thus, $\ell(\mathcal{C}') \leq \ell(\mathcal{C})$. By symmetry, we obtain $\ell(\mathcal{C}) \not\leq \ell(\mathcal{C}')$ if $\mathcal{C}' \not\subseteq \mathcal{C}$, and $\ell(\mathcal{C}) \leq \ell(\mathcal{C}')$ if $\mathcal{C}' \subsetneq \mathcal{C}$.

The above discussion shows that the strength relations according to $\mathcal{E}(z, s)$ are exactly those listed in the lemma. \square

Using Lemma 5.3 we can compute the set $\langle L \rangle$ for each label $L \in \Sigma(z, s)$.

COROLLARY 5.1. *We have*

- $\langle X \rangle = \{X\}$,

- $\langle M \rangle = \{M, X\}$,
- $\langle U \rangle = \{U, X\}$,
- $\langle D \rangle = \{D, M, U, X, \ell(\{s\})\}$,
- $\langle P \rangle = \{P, U, X\} \cup \{\ell(C) \mid \emptyset \neq C \subseteq \mathfrak{C}(z)\}$, and
- $\langle \ell(C) \rangle = \{U, X\} \cup \{\ell(C') \mid \emptyset \neq C' \subseteq C\}$ for each $\emptyset \neq C \subseteq \mathfrak{C}(z)$.

5.5 Computing $\mathcal{R}(\Pi(z, s))$ After defining our problem family and collecting some basic facts about it, the next step is to examine how the problems from this family behave under applying $\mathcal{R}(\cdot)$. Specifically, in Lemma 5.4, we compute $\mathcal{R}(\Pi(z, s))$. To this end, for any $L \in \Sigma(z, s) \setminus \{D, M\}$, define

$$\langle L \rangle' := \begin{cases} \langle L, D \rangle & \text{if } L \leq \ell(\{s\}), \\ \langle L, M \rangle & \text{if } L \not\leq \ell(\{s\}). \end{cases}$$

LEMMA 5.4. *The set $\Sigma_{\mathcal{R}}$ of output labels of $\mathcal{R}(\Pi(z, s))$ is given by*

$$\{\langle L \rangle, \langle L \rangle' \mid L \in \Sigma(z, s) \setminus \{D, M\}\}.$$

Moreover, $\langle L \rangle \neq \langle L' \rangle'$ for any $L, L' \in \Sigma(z, s) \setminus \{D, M\}$.

The hyperedge constraint $\mathcal{E}_{\mathcal{R}}$ of $\mathcal{R}(\Pi(z, s))$ consists of all configurations $\langle L_1 \rangle' \langle L_2 \rangle \langle L_3 \rangle \dots \langle L_r \rangle$ satisfying $L_j \in \Sigma(z, s) \setminus \{D, M\}$ for all $1 \leq j \leq r$ and at least one of the following two conditions.

1. There exists some index $1 \leq j \leq r$ such that $L_j = X$ and $L_{j'} = P$ for each $j' \neq j$.
2. Both of the following properties hold.
 - (a) $L_j \notin \{P, X\}$ for each $1 \leq j \leq r$.
 - (b) For each $i \in \mathfrak{C}(z)$, there are exactly z_i indices $j \in \{1, \dots, r\}$ such that $L_j = \ell(C)$ for some color set C containing color i .

The node constraint $\mathcal{N}_{\mathcal{R}}$ of $\mathcal{R}(\Pi(z, s))$ consists of all configurations $B_1 \dots B_{\Delta}$ of labels from $\{\langle L \rangle, \langle L \rangle' \mid L \in \Sigma(z, s) \setminus \{D, M\}\}$ such that there exists a choice $(A_1, \dots, A_{\Delta}) \in B_1 \times \dots \times B_{\Delta}$ satisfying $A_1 \dots A_{\Delta} \in \mathcal{N}(z, s)$.

Proof. Let \mathcal{E} denote the set of all configurations $\langle L_1 \rangle' \langle L_2 \rangle \langle L_3 \rangle \dots \langle L_r \rangle$ with labels from $\Sigma(z, s) \setminus \{D, M\}$ that satisfy the two conditions given in the lemma. Recall Definition 3.1. We start by showing that the hyperedge constraint $\mathcal{E}_{\mathcal{R}}$ of $\mathcal{R}(\Pi(z, s))$ is equal to \mathcal{E} . To this end, by the definition of $\mathcal{E}_{\mathcal{R}}$, it suffices to show that the following three conditions hold.

- (Z1) For each configuration $B_1 \dots B_r \in \mathcal{E}$ and each choice $(A_1, \dots, A_r) \in B_1 \times \dots \times B_r$, we have $A_1 \dots A_r \in \mathcal{E}(z, s)$.
- (Z2) For each configuration $B_1 \dots B_r$ with labels from $2^{\Sigma(z, s)} \setminus \{\emptyset\}$ that cannot be relaxed to any configuration from \mathcal{E} , there exists a choice $(A_1, \dots, A_r) \in B_1 \times \dots \times B_r$ satisfying $A_1 \dots A_r \notin \mathcal{E}(z, s)$.
- (Z3) If some configuration $B_1 \dots B_r \in \mathcal{E}$ can be relaxed to some configuration $\overline{B}_1 \dots \overline{B}_r \in \mathcal{E}$, then $B_1 \dots B_r$ and $\overline{B}_1 \dots \overline{B}_r$ are identical (up to permutation).

We first show that (Z1) holds. Consider an arbitrary configuration $Y = \langle L_1 \rangle' \langle L_2 \rangle \langle L_3 \rangle \dots \langle L_r \rangle$ from \mathcal{E} . If Y satisfies Condition 1 in Lemma 5.4, then $Y = \langle X, M \rangle \langle P \rangle^{r-1} = \langle M \rangle \langle P \rangle^{r-1}$ or $Y = \langle P, D \rangle \langle X \rangle \langle P \rangle^{r-2}$ (where we use¹⁰ Corollary 5.1 and the fact that $X \not\leq \ell(\{s\})$ and $P \leq \ell(\{s\})$, by Lemma 5.3). By the definition of $\langle \cdot \rangle$, it follows that Y can be obtained from at least one member of $\{MP^{r-1}, XP^{r-1}, DXP^{r-2}\}$ by repeatedly choosing some label in the configuration and replacing it by a stronger one. Since, by Conditions 1 and 2 in the definition of $\mathcal{E}(z, s)$, we know that MP^{r-1} , XP^{r-1} and DXP^{r-2} are all contained in $\mathcal{E}(z, s)$, we can conclude, by the definition of strength, that (Z1) is satisfied for each configuration Y satisfying Condition 1 in Lemma 5.4.

¹⁰Recall also that all permutations of a configuration are considered to be the same configuration since, formally, they are multisets.

If, on the other hand, Y satisfies Condition 2 in Lemma 5.4, then, in order to show that (Z1) holds, it suffices to show that $DL_2 \dots L_r \in \mathcal{E}(z, s)$ if $L_1 \leq \ell(\{s\})$, $ML_2 \dots L_r \in \mathcal{E}(z, s)$ if $L_1 \not\leq \ell(\{s\})$, and $L_1 \dots L_r \in \mathcal{E}(z, s)$, by the same argumentation as in the previous case. Each of these three configurations (under the respective condition) satisfies Conditions 3a, 3b, and 3d in the definition of $\mathcal{E}(z, s)$, by Condition 2a in Lemma 5.4, the fact that $L_j \in \Sigma(z, s) \setminus \{D, M\}$ for all $2 \leq j \leq r$, and Condition 2b in Lemma 5.4, respectively. Moreover, we claim that it follows from Condition 2b (and Condition 2a) in Lemma 5.4 that each of these three configurations (under the respective condition) satisfies also Condition 3c in the definition of $\mathcal{E}(z, s)$. For $ML_2 \dots L_r$ and $L_1 \dots L_r$ this is immediate; for $DL_2 \dots L_r$ under the condition $L_1 \leq \ell(\{s\})$, the argumentation is a bit more involved: Observe that $L_1 \leq \ell(\{s\})$, together with Condition 2a and Lemma 5.3, implies that $L_1 = \ell(C)$ for some C containing color s . By Condition 2b in Lemma 5.4, it follows that there are exactly $z_s - 1$ indices $j \in \{2, \dots, r\}$ such that $L_j = \ell(C')$ for some C' containing color s , which in turn implies that $DL_2 \dots L_r$ satisfies Condition 3c in the definition of $\mathcal{E}(z, s)$, proving the claim. Hence, all of $DL_2 \dots L_r$, $ML_2 \dots L_r$, and $L_1 \dots L_r$ (under the mentioned respective conditions) satisfy Condition 3 in the definition of $\mathcal{E}(z, s)$, which implies that they are contained in $\mathcal{E}(z, s)$. As shown, (Z1) follows.

Next, we prove (Z2). Let $Y = B_1 \dots B_r$ be a configuration with labels from $2^{\Sigma(z,s)} \setminus \{\emptyset\}$ that cannot be relaxed to any configuration from \mathcal{E} . For a contradiction, assume that (Z2) does not hold, which implies that for each choice $(A_1, \dots, A_r) \in B_1 \times \dots \times B_r$ we have $A_1 \dots A_r \in \mathcal{E}(z, s)$. Observe that adding to some set B_j all labels $L \in \Sigma(z, s) \setminus B_j$ that satisfy $L' \leq L$ for some $L' \in B_j$ does not change whether there exists a choice $(A_1, \dots, A_r) \in B_1 \times \dots \times B_r$ satisfying $A_1 \dots A_r \notin \mathcal{E}(z, s)$, by the definition of strength. As adding labels to the sets B_j also cannot make Y relaxable to some configuration it could not be relaxed to before the addition, we can (and will) therefore assume that B_j is right-closed for each $1 \leq j \leq r$. We will now collect some properties of Y that we derive from the above knowledge and assumptions regarding Y .

From the definition of $\mathcal{E}(z, s)$ and the assumption that each choice $(A_1, \dots, A_r) \in B_1 \times \dots \times B_r$ satisfies $A_1 \dots A_r \in \mathcal{E}(z, s)$, it follows directly that there is at most one index $1 \leq j \leq r$ such that $M \in B_j$. W.l.o.g., let 1 be this index (if it exists), i.e., we have

$$(5.1) \quad M \notin B_j \text{ for all } 2 \leq j \leq r,$$

which, by Corollary 5.1 and the right-closedness of the B_j , implies

$$(5.2) \quad B_j \subseteq \langle P \rangle \text{ for all } 2 \leq j \leq r.$$

Now, consider the configuration $\langle M \rangle \langle P \rangle^{r-1} = \langle X, M \rangle \langle P \rangle^{r-1} = \langle X \rangle' \langle P \rangle^{r-1}$, which is contained in \mathcal{E} , by Condition 1 in Lemma 5.4. Since Y cannot be relaxed to any configuration from \mathcal{E} , we obtain $B_1 \not\subseteq \langle M \rangle$, by (5.2).

Moreover, if $B_j = \langle X \rangle$ for some $2 \leq j \leq r$, then $B_j \subseteq \langle X \rangle$, which, combined with $B_1 \subseteq \Sigma(z, s) = \langle P \rangle'$ (which follows from Lemma 5.3 and Corollary 5.1) and $B_{j'} \subseteq \langle P \rangle$ for all $j' \in \{2, \dots, r\} \setminus \{j\}$ (which follows from (5.2)), would imply that Y can be relaxed to $\langle P \rangle' \langle X \rangle \langle P \rangle^{r-2} \in \mathcal{E}$, yielding a contradiction. Hence, $B_j \neq \langle X \rangle$ for all $2 \leq j \leq r$, which, together with (5.1), Corollary 5.1, and the already established fact that $B_1 \not\subseteq \langle M \rangle$, implies

$$(5.3) \quad B_j \not\subseteq \{M, X\} \text{ for all } 1 \leq j \leq r.$$

It follows that if $P \in B_j$ for some $1 \leq j \leq r$, then there exists a choice $(A_1, \dots, A_r) \in B_1 \times \dots \times B_r$ such that $A_j = P$ and $A_{j'} \notin \{M, X\}$ for all $j' \neq j$, which would imply that $A_1 \dots A_r \notin \mathcal{E}(z, s)$ (by Conditions 1, 2 and 3a in the definition of $\mathcal{E}(z, s)$), yielding a contradiction. Hence,

$$(5.4) \quad P \notin B_j \text{ for all } 1 \leq j \leq r.$$

While, so far, we only made use of Conditions 1, 2 and 3a in the definition of $\mathcal{E}(z, s)$ and the fact that Y cannot be relaxed to any configuration satisfying Condition 1 in Lemma 5.4, we will now also take advantage of the other conditions in the definition of $\mathcal{E}(z, s)$ and Lemma 5.4.

For each $1 \leq j \leq r$ let S_j be the set of all colors $i \in \mathcal{C}(z)$ satisfying $\ell(\{i\}) \in B_j$. We claim that

$$(5.5) \quad \text{for each } i \in \mathcal{C}(z) \text{ there are at most } z_i \text{ indices } j \in \{1, \dots, r\} \text{ such that } i \in S_j.$$

If this was not true for some $i \in \mathcal{C}(z)$, then, by (5.3) and the definition of the S_j , it would be possible to pick one label from each B_j such that the resulting configuration is $\ell(\{i\})^{z_i+1} L_{z_i+2} \dots L_r$ where $L_{j'} \notin \{M, X\}$ for each

$z_i + 2 \leq j' \leq r$. By Conditions 1, 2, 3c and 3d in the definition of $\mathcal{E}(z, s)$, this configuration is not contained in $\mathcal{E}(z, s)$, yielding a contradiction and proving the claim.

Consider B_j for some arbitrary index $2 \leq j \leq r$. By the right-closedness of B_j , Corollary 5.1, and the definition of S_j , we know that $\ell(C) \notin B_j$ for each $C \not\subseteq S_j$. By (5.2), (5.4) and Corollary 5.1, it follows that

$$(5.6) \quad B_j \subseteq \langle \ell(S_j) \rangle \text{ for all } 2 \leq j \leq r,$$

where we set $\ell(\emptyset) := U$. Now consider B_1 . Using an analogous argumentation we obtain that $\ell(C) \notin B_1$ for each $C \not\subseteq S_1$ and, if $s \notin S_1$, additionally that $D \notin B_1$. By (5.4), Corollary 5.1, and the observation that $s \notin S_1$ if and only if $\ell(S_1) \not\subseteq \ell(\{s\})$, it follows that

$$(5.7) \quad B_1 \subseteq \langle \ell(S_1) \rangle'.$$

Let S'_1, \dots, S'_r be a collection of subsets of $\mathcal{C}(z)$ such that $S_j \subseteq S'_j$ for all $1 \leq j \leq r$ and for each $i \in \mathcal{C}(z)$ there are exactly z_i indices $j \in \{1, \dots, r\}$ such that $i \in S'_j$. Such a collection exists by (5.5) and the fact that $z_i \leq r - 1$. By Corollary 5.1 and the fact that $S_j \subseteq S'_j$ for each $1 \leq j \leq r$, we know that $\langle \ell(S_j) \rangle \subseteq \langle \ell(S'_j) \rangle$ and $\langle \ell(S_j) \rangle' \subseteq \langle \ell(S'_j) \rangle'$, for each $1 \leq j \leq r$. By (5.6) and (5.7), it follows that $B_1 \subseteq \langle \ell(S'_1) \rangle'$ and $B_j \subseteq \langle \ell(S'_j) \rangle'$ for each $2 \leq j \leq r$, which in turn implies that Y can be relaxed to the configuration $\langle \ell(S'_1) \rangle' \langle \ell(S'_2) \rangle' \langle \ell(S'_3) \rangle' \dots \langle \ell(S'_r) \rangle'$. As $i \in S'_j$ if and only if $\ell(S'_j) = \ell(C)$ for some color set C containing color i , the definition of the S'_j implies that $\langle \ell(S'_1) \rangle' \langle \ell(S'_2) \rangle' \langle \ell(S'_3) \rangle' \dots \langle \ell(S'_r) \rangle'$ satisfies Condition 2 in Lemma 5.4. Since also $\ell(S'_j) \in \Sigma(z, s) \setminus \{D, M\}$, we obtain $\langle \ell(S'_1) \rangle' \langle \ell(S'_2) \rangle' \langle \ell(S'_3) \rangle' \dots \langle \ell(S'_r) \rangle' \in \mathcal{E}_{\mathcal{R}}$, yielding a contradiction to the fact that Y cannot be relaxed to any configuration from \mathcal{E} . Hence, (Z2) follows.

Now, we prove (Z3). For a contradiction, assume that (Z3) does not hold. Then, there exist two configurations $Y = B_1 \dots B_r \in \mathcal{E}$ and $\bar{Y} = \bar{B}_1 \dots \bar{B}_r \in \mathcal{E}$ such that $B_j \subseteq \bar{B}_j$ for each $1 \leq j \leq r$ and there exists some index j satisfying $B_j \subsetneq \bar{B}_j$. W.l.o.g., assume that 1 is such an index, i.e., we have $B_1 \subsetneq \bar{B}_1$. Observe that \bar{B}_1 is right-closed, by the definition of \mathcal{E} .

Consider first the case that Y satisfies Condition 1 in Lemma 5.4. Since $\langle P \rangle$ and $\langle P \rangle'$ both contain P , and $P \notin \langle L \rangle$ as well as $P \notin \langle L \rangle'$ for each $L \in \Sigma(z, s) \setminus \{P\}$, it follows that also \bar{Y} satisfies Condition 1, by Condition 2a. Observe further that $\langle X \rangle \subsetneq \langle X \rangle' \subsetneq \langle P \rangle'$, $\langle X \rangle \subsetneq \langle P \rangle \subsetneq \langle P \rangle'$, $\langle P \rangle \not\subseteq \langle X \rangle'$, and $\langle X \rangle' \not\subseteq \langle P \rangle$, by Corollary 5.1. Hence, it holds for all $1 \leq j \leq r$ that if $B_j \in \{\langle P \rangle, \langle P \rangle'\}$, then $\bar{B}_j \in \{\langle P \rangle, \langle P \rangle'\}$, and if $B_j \in \{\langle X \rangle', \langle P \rangle'\}$, then $\bar{B}_j \in \{\langle X \rangle', \langle P \rangle'\}$. As, by Condition 1, both Y and \bar{Y} contain some element from $\{\langle X \rangle', \langle P \rangle'\}$ in exactly one position and some element from $\{\langle P \rangle, \langle P \rangle'\}$ in exactly $r - 1$ positions, it follows that $B_j = \bar{B}_j$ for each $1 \leq j \leq r$, yielding a contradiction to $B_1 \subsetneq \bar{B}_1$.

Consider now the other case, namely that Y satisfies Condition 2 in Lemma 5.4. We start by observing that sets of the form $\langle L \rangle$ for some $L \in \Sigma(z, s) \setminus \{D, M\}$ do not contain the label M whereas sets of the form $\langle L \rangle'$ for some $L \in \Sigma(z, s) \setminus \{D, M\}$ contain the label M . In particular, Y and \bar{Y} contain a set of the form $\langle L \rangle'$ for some $L \in \Sigma(z, s) \setminus \{D, M\}$ in exactly one position, and the index specifying the position is the same for Y and \bar{Y} .

For each $1 \leq j \leq r$, let $\underline{L}_j \in \Sigma(z, s)$ be the (unique) label satisfying $B_j = \langle \underline{L}_j \rangle$ or $B_j = \langle \underline{L}_j \rangle'$, and $\bar{\underline{L}}_j \in \Sigma(z, s)$ the (unique) label satisfying $\bar{B}_j = \langle \bar{\underline{L}}_j \rangle$ or $\bar{B}_j = \langle \bar{\underline{L}}_j \rangle'$. Since $B_1 \subsetneq \bar{B}_1$, we have $\bar{\underline{L}}_1 < \underline{L}_1$, by Lemma 5.3, Corollary 5.1 and the definition of $\langle \cdot \rangle'$. Analogously, from $B_j \subseteq \bar{B}_j$ we obtain $\bar{\underline{L}}_j \leq \underline{L}_j$ for each $2 \leq j \leq r$. Observe that \bar{Y} cannot satisfy Condition 1 in Lemma 5.4 as otherwise $\langle X \rangle$ or $\langle X \rangle'$ is contained in \bar{Y} whereas no subset of $\langle X \rangle$ or $\langle X \rangle'$ is contained in Y (by Condition 2a and Corollary 5.1). Hence, also \bar{Y} satisfies Condition 2. From Condition 2a it follows that both \underline{L}_1 and $\bar{\underline{L}}_1$ are contained in $\Sigma(z, s) \setminus \{D, M, P, X\}$, which, together with $\bar{\underline{L}}_1 < \underline{L}_1$ and Lemma 5.3, implies that $\bar{\underline{L}}_1 = \ell(\bar{C})$ for some $\emptyset \neq \bar{C} \subseteq \mathcal{C}(z)$ while $\underline{L}_1 = U$ or $\underline{L}_1 = \ell(C)$ for some $\emptyset \neq C \subsetneq \bar{C}$. In particular, there is some color $i \in \bar{C}$ such that there is no $C' \subseteq \mathcal{C}(z)$ containing i and satisfying $\underline{L}_1 = \ell(C')$. Moreover, for each index $2 \leq j \leq r$ such that $\underline{L}_j = \ell(C_j)$ for some color set C_j containing color i , we also have that $\bar{\underline{L}}_j = \ell(\bar{C}_j)$ for some color set \bar{C}_j containing color i , by $\bar{\underline{L}}_j \leq \underline{L}_j$, Lemma 5.3, and the fact that \underline{L}_j and $\bar{\underline{L}}_j$ are contained in $\Sigma(z, s) \setminus \{D, M, P, X\}$. It follows that the number of indices $1 \leq j \leq r$ such that $\bar{\underline{L}}_j = \ell(\bar{C}_j)$ for some color set \bar{C}_j containing color i is strictly larger than the number of indices $1 \leq j' \leq r$ such that $\underline{L}_{j'} = \ell(C_{j'})$ for some color set $C_{j'}$ containing color i . This yields a contradiction to Condition 2b and proves that (Z3) holds. It follows that $\mathcal{E}_{\mathcal{R}} = \mathcal{E}$, i.e., $\mathcal{E}_{\mathcal{R}}$ is indeed as specified in the lemma.

The set $\Sigma_{\mathcal{R}}$ of output labels of $\mathcal{R}(\Pi(z, s))$ is precisely the set of all labels appearing in at least one configuration in $\mathcal{E}_{\mathcal{R}}$. From Condition 1 in Lemma 5.4, it follows that $\{\langle L \rangle, \langle L \rangle' \mid L \in \{P, X\}\} \subseteq \Sigma_{\mathcal{R}}$. Now, consider some (not necessarily nonempty) arbitrary color set $C \subseteq \mathcal{C}(z)$ and, for each color $i \in \mathcal{C}(z)$, define $z'_i := z_i - 1$ if

$i \in \mathcal{C}$, and $z'_i := z_i$ if $i \notin \mathcal{C}$. Consider the configuration $Y = \langle \ell(\mathcal{C}_1) \rangle' \langle \ell(\mathcal{C}_2) \rangle \langle \ell(\mathcal{C}_3) \rangle \dots \langle \ell(\mathcal{C}_{r-1}) \rangle \langle \ell(\mathcal{C}) \rangle$ defined by $\mathcal{C}_j := \{i \in \mathcal{C}(z) \mid j \leq z'_i\}$ for each $1 \leq j \leq r-1$, where we set $\ell(\emptyset) := U$. By construction (and the fact that $z_i \leq r-1$ for each $i \in \mathcal{C}(z)$), configuration Y satisfies Condition 2, which implies that $\{\langle L \rangle, \langle L \rangle' \mid L \in \Sigma(z, s) \setminus \{D, M, P, X\}\}$. It follows that $\Sigma_{\mathcal{R}}$ is as specified in the lemma. The claimed property that $\langle L \rangle \neq \langle L \rangle'$ for any $L, L' \in \Sigma(z, s) \setminus \{D, M\}$ follows from the fact that $\langle L \rangle'$ contains M for any $L' \in \Sigma(z, s) \setminus \{D, M\}$ while $\langle L \rangle$ does not contain M for any $L \in \Sigma(z, s) \setminus \{D, M\}$, by Lemma 5.3.

The provided characterization of $\mathcal{N}_{\mathcal{R}}$ directly follows from the definition of $\mathcal{R}(\cdot)$. \square

Set $\Sigma_{\mathcal{R}}^0 := \{\langle L \rangle \mid L \in \Sigma(z, s) \setminus \{D, M\}\}$ and $\Sigma_{\mathcal{R}}^1 := \{\langle L \rangle' \mid L \in \Sigma(z, s) \setminus \{D, M\}\}$. By Lemma 5.4, it follows that $\Sigma_{\mathcal{R}}^0$ and $\Sigma_{\mathcal{R}}^1$ are disjoint and $\Sigma_{\mathcal{R}}^0 \cup \Sigma_{\mathcal{R}}^1 = \Sigma_{\mathcal{R}}$.

Similar to before, it will be useful to collect information about the strength of the labels in $\Sigma_{\mathcal{R}}$ and compute the right-closed subsets generated by each label. We will do so in the following with Lemma 5.5 and Corollary 5.2, which are analogues of Lemma 5.3 and Corollary 5.1 for $\mathcal{R}(\Pi(z, s))$ instead of $\Pi(z, s)$.




Figure 3: The node diagram of $\mathcal{R}(\Pi(z, s))$ for $\text{len}(z) = 3$. Numbers correspond to color sets, and $\ell(\cdot)$ is omitted for clarity.

LEMMA 5.5. For any $f(\cdot), \bar{f}(\cdot) \in \{\langle \cdot \rangle, \langle \cdot \rangle'\}$ and any $L, \bar{L} \in \Sigma(z, s) \setminus \{D, M\}$, we have $f(L) \leq \bar{f}(\bar{L})$ if and only if the following two conditions are satisfied.

1. $f(\cdot) = \langle \cdot \rangle$ or $\bar{f}(\cdot) = \langle \cdot \rangle'$ (or both).
2. At least one of the following conditions holds.
 - (a) $L = X$
 - (b) $\bar{L} = P$
 - (c) $L = U$ and $\bar{L} \neq X$
 - (d) $L = \ell(\mathcal{C})$ and $\bar{L} = \ell(\bar{\mathcal{C}})$ for some $\mathcal{C}, \bar{\mathcal{C}} \in 2^{\mathcal{C}(z)} \setminus \{\emptyset\}$ satisfying $\mathcal{C} \subseteq \bar{\mathcal{C}}$

Moreover, two labels from $\Sigma_{\mathcal{R}}$ are equally strong if and only if they are identical.

Proof. Observe that, by the definition of $\mathcal{N}_{\mathcal{R}}$, it holds for any $B, \bar{B} \in \Sigma_{\mathcal{R}}$ with $B \subseteq \bar{B}$ that for any configuration $B B_2 \dots B_{\Delta} \in \mathcal{N}_{\mathcal{R}}$, we also have $\bar{B} B_2 \dots B_{\Delta} \in \mathcal{N}_{\mathcal{R}}$. Hence, for any $B, \bar{B} \in \Sigma_{\mathcal{R}}$ with $B \subseteq \bar{B}$ we have $B \leq \bar{B}$, by the definition of strength. By Corollary 5.1, it follows that $\langle L \rangle \leq \langle \bar{L} \rangle$ for any L, \bar{L} satisfying Condition 2 in

the lemma. Moreover, since $\langle L \rangle' \in \{\langle L \rangle \cup \{M\}, \langle L \rangle \cup \{D, M\}\}$ for any $L \in \Sigma(z, s) \setminus \{D, M\}$, we also obtain that $\langle L \rangle \leq \langle \bar{L} \rangle'$ for any L, \bar{L} satisfying Condition 2. Furthermore, note that for any $L, \bar{L} \in \Sigma(z, s) \setminus \{D, M\}$, it holds that if $\langle L \rangle' = \langle L \rangle \cup \{D, M\}$ and $\langle \bar{L} \rangle' = \langle \bar{L} \rangle \cup \{M\}$, then $L \leq \ell(\{s\})$ and $\bar{L} \not\leq \ell(\{s\})$ (by the definition of $\langle \cdot \rangle'$), which implies $\langle L \rangle \not\leq \langle \bar{L} \rangle$ (by the definition of $\langle \cdot \rangle$). Since, by Corollary 5.1, any L, \bar{L} satisfying Condition 2 satisfy $\langle L \rangle \leq \langle \bar{L} \rangle$, we also obtain that $\langle L \rangle' \leq \langle \bar{L} \rangle'$ for any L, \bar{L} satisfying Condition 2, analogously to above. Hence, we conclude that if Conditions 1 and 2 hold, then we have $f(L) \leq \bar{f}(\bar{L})$, as desired. What remains is to show that $f(L) \leq \bar{f}(\bar{L})$ does not hold if at least one of the two conditions is violated.

Consider first the case that Condition 1 is violated, i.e., consider $\langle L \rangle', \langle \bar{L} \rangle'$ for some arbitrary $L, \bar{L} \in \Sigma(z, s) \setminus \{D, M\}$. We need to show that $\langle L \rangle' \not\leq \langle \bar{L} \rangle'$. Consider the configuration $\langle L \rangle' (\langle X \rangle')^{\Delta-1}$. Since $M^\Delta \in \mathcal{N}(z, s)$ and, by the definition of $\langle \cdot \rangle'$, we have $M \in \langle L \rangle'$ and $M \in \langle X \rangle'$, it holds that $\langle L \rangle' (\langle X \rangle')^{\Delta-1} \in \mathcal{N}_{\mathcal{R}}$, by the definition of $\mathcal{N}_{\mathcal{R}}$. Now, consider the configuration $\langle \bar{L} \rangle' (\langle X \rangle')^{\Delta-1}$ obtained from the above configuration by replacing $\langle L \rangle'$ by $\langle \bar{L} \rangle'$. Observe that we have $\langle X \rangle = \{M, X\}$ and the only configurations in $\mathcal{N}(z, s)$ containing M or X are M^Δ and $D^{\Delta-1}X$. Since, by Corollary 5.1, neither M nor D is contained in $\langle \bar{L} \rangle'$, it follows that $\langle \bar{L} \rangle' (\langle X \rangle')^{\Delta-1} \notin \mathcal{N}_{\mathcal{R}}$, by the definition of $\mathcal{N}_{\mathcal{R}}$. By the definition of strength, it follows that $\langle L \rangle' \not\leq \langle \bar{L} \rangle'$, as desired.

Now, consider the other case, i.e., consider some arbitrary $f(\cdot), \bar{f}(\cdot) \in \{\langle \cdot \rangle, \langle \cdot \rangle'\}$ and $L, \bar{L} \in \Sigma(z, s) \setminus \{D, M\}$ violating Condition 2. We need to show that $f(L) \not\leq \bar{f}(\bar{L})$. Note that Condition 2 is violated if and only if all of Conditions 2a, 2b, 2c, and 2d are violated. We now go through the different cases for L . Note that $L \neq X$, due to the violation of Condition 2a.

If $L = U$, consider the configuration $f(L) \langle P \rangle \langle U \rangle^{\Delta-2}$. Observe that, by Corollary 5.1, U is contained in $\langle U \rangle$ and $f(L)$, and P is contained in $\langle P \rangle$. Since $PU^{\Delta-1} \in \mathcal{N}(z, s)$, it follows that $f(L) \langle P \rangle \langle U \rangle^{\Delta-2} \in \mathcal{N}_{\mathcal{R}}$. Further, observe that the violation of Condition 2c implies that $\bar{L} = X$, and consider the configuration $\bar{f}(X) \langle P \rangle \langle U \rangle^{\Delta-2}$ obtained from the above configuration by replacing $f(L)$ by $\bar{f}(\bar{L})$. We have $\bar{f}(X) \in \{\{X\}, \{M, X\}\}$, the only configurations in $\mathcal{N}(z, s)$ containing M or X are M^Δ and $D^{\Delta-1}X$, and both M and D are not contained in $\langle P \rangle$, by Corollary 5.1. It follows that $\bar{f}(X) \langle P \rangle \langle U \rangle^{\Delta-2} \notin \mathcal{N}_{\mathcal{R}}$. Hence, $f(L) \not\leq \bar{f}(\bar{L})$, as desired.

If $L = \ell(\mathcal{C})$ for some $\emptyset \neq \mathcal{C} \subseteq \mathcal{C}(z)$, consider the configuration $f(L) \langle L \rangle^{\Delta-|\mathcal{C}|} \langle U \rangle^{|\mathcal{C}|-1}$. Observe that, by Corollary 5.1, $\ell(\mathcal{C})$ is contained in $\langle L \rangle$ and $f(L)$, and U is contained in $\langle U \rangle$. Since $\ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} U^{|\mathcal{C}|-1} \in \mathcal{N}(z, s)$, it follows that $f(L) \langle L \rangle^{\Delta-|\mathcal{C}|} \langle U \rangle^{|\mathcal{C}|-1} \in \mathcal{N}_{\mathcal{R}}$. Further, observe that the violation of Conditions 2b and 2d implies that $\bar{L} \in \{U, X\}$ or $\bar{L} = \ell(\bar{\mathcal{C}})$ for some $\emptyset \neq \bar{\mathcal{C}} \subseteq \mathcal{C}(z)$ satisfying $\mathcal{C} \not\subseteq \bar{\mathcal{C}}$. Consider the configuration $Y := \bar{f}(\bar{L}) \langle \ell(\mathcal{C}) \rangle^{\Delta-|\mathcal{C}|} \langle U \rangle^{|\mathcal{C}|-1}$ obtained from the above configuration by replacing $f(L)$ by $\bar{f}(\bar{L})$, and assume for a contradiction that $Y \in \mathcal{N}_{\mathcal{R}}$. Since $\text{len}(z) \leq \Delta - 1$, we have $|\mathcal{C}| \leq |\mathcal{C}(z)| \leq \Delta - 1$, which implies that $\Delta - |\mathcal{C}| \geq 1$, i.e., Y contains $\langle \ell(\mathcal{C}) \rangle$. Observe also that our assumption for this section, $\Delta \geq 3$, implies that there are at least two positions in Y with a label from $\{\langle \ell(\mathcal{C}) \rangle, \langle U \rangle\}$. Since, by Corollary 5.1 and our insights about L , we have $M \notin \langle \ell(\mathcal{C}) \rangle$, $P \notin \bar{f}(\bar{L}) \cup \langle \ell(\mathcal{C}) \rangle \cup \langle U \rangle$, and $D \notin \langle \ell(\mathcal{C}) \rangle \cup \langle U \rangle$, it follows that there is no choice $(A_1, \dots, A_\Delta) \in \bar{f}(\bar{L}) \times \langle \ell(\mathcal{C}) \rangle^{\Delta-|\mathcal{C}|} \times \langle U \rangle^{|\mathcal{C}|-1}$ such that $A_1 \dots A_\Delta \in \{M^\Delta, PU^{\Delta-1}, D^{\Delta-1}X\}$. By the definition of $\mathcal{N}(z, s)$ and the fact that $Y \in \mathcal{N}_{\mathcal{R}}$, this implies that there is some $\emptyset \neq \mathcal{C}^* \subseteq \mathcal{C}(z)$ such that $\ell(\mathcal{C}^*)$ is contained in at least $\Delta - |\mathcal{C}^*| + 1$ of the Δ sets in Y . Since $\ell(\mathcal{C}^*) \notin \langle U \rangle$ (by Corollary 5.1), we have $\Delta - |\mathcal{C}^*| + 1 \leq \Delta - |\mathcal{C}| + 1$, which implies $|\mathcal{C}^*| \geq |\mathcal{C}|$. If $\mathcal{C}^* \neq \mathcal{C}$, then it follows that $\mathcal{C}^* \not\subseteq \mathcal{C}$, which implies $\ell(\mathcal{C}^*) \notin \langle \ell(\mathcal{C}) \rangle$, by Corollary 5.1; it follows that $\ell(\mathcal{C}^*)$ is contained in at most one of the Δ sets in Y , which implies $|\mathcal{C}^*| \geq \Delta$, yielding a contradiction to the fact that $|\mathcal{C}^*| \leq |\mathcal{C}(z)| \leq \Delta - 1$. Hence, we know that $\mathcal{C}^* = \mathcal{C}$. Recall that we have $\bar{L} \in \{U, X\}$ or $\bar{L} = \ell(\bar{\mathcal{C}})$ for some $\emptyset \neq \bar{\mathcal{C}} \subseteq \mathcal{C}(z)$ satisfying $\mathcal{C} \not\subseteq \bar{\mathcal{C}}$, and observe that in the latter case we have $\ell(\mathcal{C}) \notin \langle \ell(\mathcal{C}^*) \rangle$ and $\ell(\mathcal{C}) \notin \langle \ell(\mathcal{C}^*) \rangle'$, by Corollary 5.1 and the definition of $\langle \cdot \rangle'$. Since $\ell(\mathcal{C})$ is also not contained in any of $\langle U \rangle, \langle U \rangle', \langle X \rangle$, and $\langle X \rangle'$, it follows that $\ell(\mathcal{C}) \notin \bar{f}(\bar{L})$. Moreover, since $\mathcal{C}^* = \mathcal{C}$, we have $\ell(\mathcal{C}^*) \notin \bar{f}(\bar{L})$. By the definition of Y , it follows that $\ell(\mathcal{C}^*)$ is contained in at most $\Delta - |\mathcal{C}| = \Delta - |\mathcal{C}^*|$ of the Δ sets in Y . This contradicts the fact established above that $\ell(\mathcal{C}^*)$ is contained in at least $\Delta - |\mathcal{C}^*| + 1$ of the Δ sets in Y . Hence, we have $Y \notin \mathcal{N}_{\mathcal{R}}$, which implies that $f(L) \not\leq \bar{f}(\bar{L})$, as desired.

If $L = P$, consider the configuration $f(L) \langle U \rangle^{\Delta-1}$. Observe that, by Corollary 5.1, P is contained in $f(L)$ and U is contained in $\langle U \rangle$. Since $PU^{\Delta-1} \in \mathcal{N}(z, s)$, it follows that $f(L) \langle U \rangle^{\Delta-1} \in \mathcal{N}_{\mathcal{R}}$. Further, observe that the violation of Condition 2b implies that $\bar{L} \neq P$, and consider the configuration $\bar{f}(\bar{L}) \langle U \rangle^{\Delta-1}$ obtained from the above configuration by replacing $f(L)$ by $\bar{f}(\bar{L})$. We have $\langle U \rangle = \{U, X\}$, the only configurations in $\mathcal{N}(z, s)$ containing some element from $\{U, X\}$ in at least $\Delta - 1$ positions is $PU^{\Delta-1}$ (due to $\Delta \geq 3$ and $|\mathcal{C}(z)| = \text{len}(z) \leq \Delta - 1$), and P is contained neither in $\bar{f}(\bar{L})$ nor in $\langle U \rangle$, by Corollary 5.1. It follows that $\bar{f}(\bar{L}) \langle U \rangle^{\Delta-1} \notin \mathcal{N}_{\mathcal{R}}$. Hence, $f(L) \not\leq \bar{f}(\bar{L})$, as desired.

This covers all possible cases for $L \in \Sigma(z, s) \setminus \{D, M\}$ and concludes the second case. Hence, the lemma

statement follows. Note that the fact that two labels from $\Sigma_{\mathcal{R}}$ are equally strong if and only if they are identical follows from the fact that for any two distinct labels $B, B' \in \Sigma_{\mathcal{R}}$ we have $B \not\leq B'$ or $B' \not\leq B$.

□

From Lemma 5.5 we obtain Corollary 5.2 by simply collecting for any label $L \in \Sigma_{\mathcal{R}}$ the set of all labels L' satisfying $L \leq L'$.

COROLLARY 5.2. *We have*

- $\langle\langle P \rangle\rangle = \{\langle P \rangle, \langle P \rangle'\}$,
- $\langle\langle U \rangle\rangle = \{\langle L \rangle, \langle L \rangle' \mid L \in \{P, U\} \cup \{\ell(C) \mid \emptyset \neq C \subseteq \mathfrak{C}(z)\}\}$,
- $\langle\langle X \rangle\rangle = \Sigma_{\mathcal{R}}$,
- $\langle\langle \ell(C) \rangle\rangle = \{\langle L \rangle, \langle L \rangle' \mid L \in \{P\} \cup \{\ell(C') \mid C \subseteq C' \subseteq \mathfrak{C}(z)\}\}$ for each $\emptyset \neq C \subseteq \mathfrak{C}(z)$, and
- $\langle\langle L \rangle'\rangle = \langle\langle L \rangle\rangle \cap \Sigma_{\mathcal{R}}^1$, for each $L \in \Sigma(z, s) \setminus \{D, M\}$.

5.6 Computing $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$ After computing and analyzing $\mathcal{R}(\Pi(z, s))$, we now turn our attention to the next step, computing $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$. However, instead of computing $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$ exactly, we instead show that $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$ can be relaxed to some problem Π_q^* (without computing $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$ explicitly). We will define Π_q^* so that by just renaming the labels in the set $\Sigma_q^* := \Sigma_{\Pi_q^*}$ suitably we obtain some problem of the form $\Pi(z', s')$ with new parameters z', s' (compared to $\Pi(z, s)$). As relaxing a problem, by definition, does not increase the complexity of the problem, this approach will yield a proof that $\Pi(z', s')$ can be solved at least one round faster than $\Pi(z, s)$; analyzing how the parameters change from $\Pi(z, s)$ to $\Pi(z', s')$ will then enable us to prove the desired lower bound. We start by defining Π_q^* , which depends on some parameter q that can be chosen arbitrarily from $\mathfrak{C}(z) \setminus \{s\}$. Note that such a parameter q exists since, by definition, $\text{len}(z) \geq 2$.

DEFINITION 5.1. *Fix some parameter $q \in \mathfrak{C}(z) \setminus \{s\}$ and define $\ell(\emptyset) := U$. The label set Σ_q^* of Π_q^* is defined by*

$$\Sigma_q^* := \{\langle\langle P \rangle\rangle, \langle\langle U \rangle\rangle, \langle\langle U \rangle'\rangle, \langle\langle X \rangle\rangle, \langle\langle X \rangle'\rangle\} \cup \mathcal{F}_q^- \cup \mathcal{F}_q^+$$

where

$$\begin{aligned} \mathcal{F}_q^- &:= \{\langle\langle \ell(C) \rangle\rangle \mid \emptyset \neq C \subseteq \mathfrak{C}(z), q \notin C\} \quad \text{and} \\ \mathcal{F}_q^+ &:= \{\langle\langle \ell(C), \ell(C \setminus \{q\}) \rangle\rangle' \mid \emptyset \neq C \subseteq \mathfrak{C}(z), q \in C\}. \end{aligned}$$

The node constraint \mathcal{N}_q^* of Π_q^* consists of the following configurations.

- $\langle\langle X \rangle'\rangle^\Delta$
- $\langle\langle P \rangle\rangle \langle\langle U \rangle\rangle^{\Delta-1}$
- $\langle\langle U \rangle'\rangle^{\Delta-1} \langle\langle X \rangle\rangle$
- $\langle\langle \ell(C) \rangle\rangle^{\Delta-|C|+1} \langle\langle U \rangle\rangle^{|C|-1}$ for each $\emptyset \neq C \subseteq \mathfrak{C}(z)$ satisfying $q \notin C$
- $\langle\langle \ell(C), \ell(C \setminus \{q\}) \rangle\rangle'^{\Delta-|C|+1} \langle\langle U \rangle\rangle^{|C|-1}$ for each $\emptyset \neq C \subseteq \mathfrak{C}(z)$ satisfying $q \in C$

The hyperedge constraint \mathcal{E}_q^* of Π_q^* consists of all configurations $Q_1 \dots Q_r$ with $Q_j \in \Sigma_q^*$ for all $1 \leq j \leq r$ such that there exists a choice $(B_1, \dots, B_r) \in Q_1 \times \dots \times Q_r$ satisfying $B_1 \dots B_r \in \mathcal{E}_{\mathcal{R}}$.

Now we are ready to prove the desired relation between $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$ and Π_q^* . We start by relating the node configurations of the two problems.

LEMMA 5.6. *Each node configuration of $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$ can be relaxed to some configuration from \mathcal{N}_q^* .*

Proof. For a contradiction, suppose that the lemma does not hold, i.e., there exists some configuration $\mathcal{Q} = Q_1 \dots Q_\Delta$ in the node constraint of $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$ that cannot be relaxed to any configuration from \mathcal{N}_q^* . By the definition of $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$, this implies that for any choice $(B_1, \dots, B_\Delta) \in Q_1 \times \dots \times Q_\Delta$ there exists some choice $(A_1, \dots, A_\Delta) \in B_1 \times \dots \times B_\Delta$ such that $A_1 \dots A_\Delta \in \mathcal{N}(z, s)$. Moreover, as \mathcal{Q} is contained in the node constraint of $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$, we know that Q_k is right-closed for each $1 \leq k \leq \Delta$, by Observation 3.1.

We distinguish two cases based on whether there exists some Q_k containing $\langle X \rangle'$. Consider first the case that $\langle X \rangle' \in Q_k$ for some $1 \leq k \leq \Delta$. W.l.o.g., assume that $k = 1$, i.e., we have $\langle X \rangle' \in Q_1$. We claim that this implies $\langle P \rangle \notin Q_k$ for each $2 \leq k \leq \Delta$: If there exists some $2 \leq k \leq \Delta$ satisfying $\langle P \rangle \in Q_k$, then there exists some choice $(B_1, \dots, B_\Delta) \in Q_1 \times \dots \times Q_\Delta$ with $B_1 = \langle X \rangle'$ and $B_k = \langle P \rangle$. It follows by Lemma 5.3 and Corollary 5.1 that for every choice $(A_1, \dots, A_\Delta) \in B_1 \times \dots \times B_\Delta$ we have $A_1 \in \{M, X\}$ and $A_k \notin \{D, M\}$; now the claim follows by observing that every configuration in $\mathcal{N}(z, s)$ containing M requires all other labels in the configuration to be M as well and every configuration in $\mathcal{N}(z, s)$ containing X requires all other labels in the configuration to be D .

Combining the claim with the right-closedness of the Q_k , we obtain that $Q_k \subseteq \langle \langle X \rangle' \rangle$ for each $2 \leq k \leq \Delta$, by Lemma 5.5 and Corollary 5.2. As, by definition, \mathcal{Q} cannot be relaxed to $\langle \langle X \rangle' \rangle^\Delta$, it follows that $Q_1 \not\subseteq \langle \langle X \rangle' \rangle$. Since, by Lemma 5.5, we have $L \leq \langle P \rangle$ for each $L \in \Sigma_{\mathcal{R}} \setminus \langle \langle X \rangle' \rangle$, we conclude that $\langle P \rangle \in Q_1$, by the right-closedness of Q_1 . This implies that $\langle X \rangle' \notin Q_k$ for each $2 \leq k \leq \Delta$, as otherwise there would (again) exist some choice $(B_1, \dots, B_\Delta) \in Q_1 \times \dots \times Q_\Delta$ containing $\langle P \rangle$ and $\langle X \rangle'$, which leads to a contradiction as already seen above. From $Q_k \subseteq \langle \langle X \rangle' \rangle$ and $\langle X \rangle' \notin Q_k$ for each $2 \leq k \leq \Delta$, we infer that $Q_k \subseteq \langle \langle U \rangle' \rangle$ for each $2 \leq k \leq \Delta$, by Corollary 5.2. Since, by Corollary 5.2, $Q_1 \subseteq \Sigma_{\mathcal{R}} = \langle \langle X \rangle' \rangle$, it follows that \mathcal{Q} can be relaxed to the configuration $\langle \langle U \rangle' \rangle^{\Delta-1} \langle \langle X \rangle' \rangle \in \mathcal{N}_q^*$, yielding the desired contradiction.

Now consider the other case, namely that $\langle X \rangle' \notin Q_k$ for each $1 \leq k \leq \Delta$. Observe that this implies

$$(5.8) \quad Q_k \subseteq \langle \langle U \rangle' \rangle \text{ for each } 1 \leq k \leq \Delta,$$

by Corollary 5.2 and the fact that $\langle \langle X \rangle' \rangle < \langle \langle X \rangle' \rangle$ (which follows from Lemma 5.5). It follows that

$$(5.9) \quad \langle \ell(\mathcal{C}(z)) \rangle' \in Q_k \text{ for each } 1 \leq k \leq \Delta,$$

as otherwise there would be some $1 \leq k \leq \Delta$ satisfying $Q_k \subseteq \langle P \rangle$ (by Lemma 5.5, Corollary 5.2 and the right-closedness of the Q_k), which in turn would imply that \mathcal{Q} can be relaxed to $\langle \langle P \rangle \rangle \langle \langle U \rangle' \rangle^{\Delta-1} \in \mathcal{N}_q^*$, yielding a contradiction.

While, so far, we used (only) that \mathcal{Q} cannot be relaxed to any of the configurations $\langle \langle X \rangle' \rangle^\Delta$, $\langle \langle P \rangle \rangle \langle \langle U \rangle' \rangle^{\Delta-1}$, and $\langle \langle U \rangle' \rangle^{\Delta-1} \langle \langle X \rangle' \rangle$, we will now also take advantage of the fact that \mathcal{Q} cannot be relaxed to the remaining (conceptually more complex) configurations from \mathcal{N}_q^* . For each $i \in \mathcal{C}(z)$, define $R_i := \mathcal{C}(z) \setminus \{i\}$. Moreover, we will make use of the bipartite graph $\overline{G} = (\overline{V} \cup \overline{W}, \overline{E})$ obtained by defining $\overline{V} := \{Q_1, \dots, Q_\Delta\}$ and $\overline{W} := \{R_i \mid i \in \mathcal{C}(z)\}$, and setting \overline{E} to be the set of all edges $\{Q_k, R_i\}$ satisfying

- $\langle \ell(R_i) \rangle' \in Q_k$ if $i \neq q$, and
- $\langle \ell(R_i) \rangle \in Q_k$ if $i = q$.

Note that, for any distinct k, k' , we consider Q_k and $Q_{k'}$ to be different vertices in \overline{V} even if $Q_k = Q_{k'}$.

Consider any arbitrary subset $\emptyset \neq \mathcal{C} \subseteq \mathcal{C}(z)$ and recall that $\ell(\emptyset) := U$. We distinguish two cases, based on whether \mathcal{C} contains q .

If $q \notin \mathcal{C}$, then, by the definition of \mathcal{N}_q^* , configuration \mathcal{Q} cannot be relaxed to the configuration $\langle \langle \ell(\mathcal{C}) \rangle \rangle^{\Delta-|\mathcal{C}|+1} \langle \langle U \rangle \rangle^{|\mathcal{C}|-1}$, which, by (5.8), implies that there are at least $|\mathcal{C}|$ indices $k \in \{1, \dots, \Delta\}$ satisfying $Q_k \not\subseteq \langle \langle \ell(\mathcal{C}) \rangle \rangle$. Observe that for each label $L \in \langle \langle U \rangle \rangle \setminus \langle \langle \ell(\mathcal{C}) \rangle \rangle$, we have $L \in \{U, \langle U \rangle'\}$ or $L \in \{\langle \ell(\overline{C}) \rangle, \langle \ell(\overline{C}) \rangle'\}$ for some $\overline{C} \not\supseteq \mathcal{C}$, by Corollary 5.2. By Lemma 5.5, this implies that for each label $L \in \langle \langle U \rangle \rangle \setminus \langle \langle \ell(\mathcal{C}) \rangle \rangle$ there exists some $i \in \mathcal{C}$ satisfying $L \leq \langle \ell(R_i) \rangle'$. By (5.8) and the right-closedness of the Q_k , it follows that there are at least $|\mathcal{C}|$ indices $k \in \{1, \dots, \Delta\}$ such that there exists some $i \in \mathcal{C}$ satisfying $\langle \ell(R_i) \rangle' \in Q_k$.

Similarly, if $q \in \mathcal{C}$, then \mathcal{Q} cannot be relaxed to the configuration $\langle \langle \ell(\mathcal{C}) \rangle, \langle \ell(\mathcal{C} \setminus \{q\}) \rangle \rangle^{\Delta-|\mathcal{C}|+1} \langle \langle U \rangle \rangle^{|\mathcal{C}|-1}$, which, by (5.8), implies that there are at least $|\mathcal{C}|$ indices $k \in \{1, \dots, \Delta\}$ satisfying $Q_k \not\subseteq \langle \langle \ell(\mathcal{C}) \rangle, \langle \ell(\mathcal{C} \setminus \{q\}) \rangle \rangle$. Observe that for each label $L \in \langle \langle U \rangle \rangle \setminus \langle \langle \ell(\mathcal{C}) \rangle, \langle \ell(\mathcal{C} \setminus \{q\}) \rangle \rangle$, we have $L \in \{U, \langle U \rangle'\}$, $L = \langle \ell(\overline{C}) \rangle$ for some $\overline{C} \not\supseteq \mathcal{C}$, or $L = \langle \ell(\overline{C}) \rangle'$ for some $\overline{C} \not\supseteq \mathcal{C} \setminus \{q\}$, by Lemma 5.5 and Corollary 5.2. By Lemma 5.5, this implies¹¹ that for each

¹¹Note that the implication is also correct in the special case $\mathcal{C} = \{q\}$.

label $L \in \langle\langle U \rangle\rangle \setminus \langle\langle \ell(C), \ell(C \setminus \{q\}) \rangle\rangle'$ we have $L \leq \langle\ell(R_q)\rangle$ or there exists some $i \in C \setminus \{q\}$ satisfying $L \leq \langle\ell(R_i)\rangle'$. By (5.8) and the right-closedness of the Q_k , it follows that there are at least $|C|$ indices $k \in \{1, \dots, \Delta\}$ such that $\langle\ell(R_q)\rangle \in Q_k$ or there exists some $i \in C$ satisfying $\langle\ell(R_i)\rangle' \in Q_k$.

Recall the definition of \overline{G} . By the above discussion, we conclude that for any arbitrary subset $\emptyset \neq C \subseteq \mathfrak{C}(z)$, the vertex set $\{R_i \mid i \in C\} \subseteq \overline{W}$ has at least $|C|$ neighbors in \overline{V} . Hence, we can apply Theorem 3.3 (i.e., Hall's marriage theorem) to \overline{G} and obtain a function $f: \overline{W} \rightarrow \overline{V}$ such that $f(R_i) \neq f(R_{i'})$ for any $i \neq i'$, $\langle\ell(R_q)\rangle \in f(R_q)$, and $\langle\ell(R_i)\rangle' \in f(R_i)$ for any $i \neq q$. This implies that there is some choice $(B_1, \dots, B_\Delta) \in Q_1 \times \dots \times Q_\Delta$ that contains $\langle\ell(R_q)\rangle$ as well as $\langle\ell(R_i)\rangle'$ for all $i \in \mathfrak{C}(z) \setminus \{q\}$. By (5.9), it follows that there is some choice $(B_1, \dots, B_\Delta) \in Q_1 \times \dots \times Q_\Delta$ such that $Y = B_1 \dots B_\Delta$ is a permutation of $\langle\ell(R_1)\rangle' \dots \langle\ell(R_{q-1})\rangle' \langle\ell(R_q)\rangle \langle\ell(R_{q+1})\rangle' \dots \langle\ell(R_{\text{len}(z)})\rangle' (\langle\ell(\mathfrak{C}(z))\rangle')^{\Delta - \text{len}(z)}$.

Consider an arbitrary choice $(A_1, \dots, A_\Delta) \in B_1 \times \dots \times B_\Delta$, and set $\mathcal{A} := A_1 \dots A_\Delta$. Using the aforementioned characterization of Y together with Lemma 5.3 and Corollary 5.1, we collect some properties of \mathcal{A} in the following. Since $P \notin B_k$ for each $1 \leq k \leq \Delta$, we have $\mathcal{A} \neq P U^{\Delta-1}$. Since $M \notin \langle\ell(R_q)\rangle$, we have $\mathcal{A} \neq M^\Delta$. Since $D \notin \langle\ell(R_q)\rangle$, $D \notin \langle\ell(R_s)\rangle'$, and $s \neq q$ (by the definition of Π_q^*), we have $\mathcal{A} \neq D^{\Delta-1} X$. Finally, for each $\emptyset \neq C \subseteq \mathfrak{C}(z)$, we have $\ell(C) \notin \langle\ell(R_i)\rangle'$ for each $i \in C \setminus \{q\}$ and, if $q \in C$, additionally $\ell(C) \notin \langle\ell(R_q)\rangle$, which implies that $\mathcal{A} \neq \ell(C)^{\Delta - |C| + 1} \cup^{|C|-1}$.

By the definition of $\mathcal{N}(z, s)$, it follows that for any choice $(A_1, \dots, A_\Delta) \in B_1 \times \dots \times B_\Delta$, we have $A_1 \dots A_\Delta \notin \mathcal{N}(z, s)$, yielding a contradiction to the fact that \mathcal{Q} is contained in the node constraint of $\mathcal{R}(\mathcal{R}(\Pi(z, s)))$ and concluding the proof. \square

Next, we provide a useful characterization of \mathcal{E}_q^* .

LEMMA 5.7. *The hyperedge constraint \mathcal{E}_q^* of Π_q^* consists of all configurations $L_1 \dots L_r$ (with labels from Σ_q^*) satisfying at least one of the following three conditions.*

1. *There is some index $1 \leq j \leq r$ such that $L_j = \langle\langle X \rangle\rangle'$ and $L_{j'} \notin \{\langle\langle U \rangle\rangle', \langle\langle X \rangle\rangle'\}$ for each $j' \neq j$.*
2. *There are two distinct indices $j, j' \in \{1, \dots, r\}$ such that $L_j = \langle\langle X \rangle\rangle'$, $L_{j'}$ is arbitrary, and $L_{j''} \notin \{\langle\langle U \rangle\rangle', \langle\langle X \rangle\rangle'\}$ for each $j'' \in \{1, \dots, r\} \setminus \{j, j'\}$.*
3. *All of the following properties hold.*
 - (a) $L_j \neq \langle\langle P \rangle\rangle'$ for each $1 \leq j \leq r$.
 - (b) *There is at most one index $j \in \{1, \dots, r\}$ such that $L_j \in \{\langle\langle U \rangle\rangle', \langle\langle X \rangle\rangle'\}$.*
 - (c) *There are at most $z_q + 1$ indices $j \in \{1, \dots, r\}$ such that $L_j = \langle\langle U \rangle\rangle'$ or $L_j = \langle\langle \ell(C), \ell(C \setminus \{q\}) \rangle\rangle'$ for some color set C containing color q .*
 - (d) *For each $i \in \mathfrak{C}(z)$ satisfying $i \neq q$, there are at most z_i indices $j \in \{1, \dots, r\}$ such that $L_j = \langle\langle \ell(C) \rangle\rangle'$ for some color set C containing color i .*

Proof. We start by showing that any configuration (with labels from Σ_q^*) satisfying at least one of the three conditions stated in the lemma is contained in \mathcal{E}_q^* . Let $\mathcal{Q} = Q_1 \dots Q_r$ be an arbitrary configuration satisfying at least one of the conditions. If \mathcal{Q} satisfies Condition 1, then, by Corollary 5.2, we know that some permutation of $(\langle\langle X \rangle\rangle', \langle\langle P \rangle\rangle', \dots, \langle\langle P \rangle\rangle')$ is contained in $Q_1 \times \dots \times Q_r$. Since Condition 1 in Lemma 5.4 implies that $\langle\langle X \rangle\rangle' \langle\langle P \rangle\rangle'^{r-1} \in \mathcal{E}_{\mathcal{R}}$, we obtain $\mathcal{Q} \in \mathcal{E}_q^*$, by the definition of \mathcal{E}_q^* . If \mathcal{Q} satisfies Condition 2, then, by Corollary 5.2, we know that some permutation of $(\langle\langle X \rangle\rangle', \langle\langle P \rangle\rangle', \dots, \langle\langle P \rangle\rangle')$ is contained in $Q_1 \times \dots \times Q_r$. Again, Condition 1 in Lemma 5.4 implies that $\langle\langle X \rangle\rangle' \langle\langle P \rangle\rangle' \langle\langle P \rangle\rangle'^{r-2} \in \mathcal{E}_{\mathcal{R}}$, and we obtain $\mathcal{Q} \in \mathcal{E}_q^*$.

Hence, consider the case that \mathcal{Q} satisfies Condition 3. If there is some $1 \leq j \leq r$ satisfying $Q_j = \langle\langle X \rangle\rangle'$, then, by Condition 3b, \mathcal{Q} satisfies also Condition 1 and we are done. Thus, assume in the following that $Q_j \neq \langle\langle X \rangle\rangle'$ for all $1 \leq j \leq r$. Choose some $(B_1, \dots, B_r) \in Q_1 \times \dots \times Q_r$ as follows. If there is some $1 \leq j \leq r$ satisfying $Q_j = \langle\langle U \rangle\rangle'$, then, for any $1 \leq j' \leq r$,

- if $Q_{j'} = \langle\langle X \rangle\rangle'$ choose $B_{j'} = \langle U \rangle$,
- if $Q_{j'} = \langle\langle L \rangle\rangle'$ for some $L \in \{U\} \cup \{\ell(C) \mid q \notin C \subseteq \mathfrak{C}(z)\}$, choose $B_{j'} = \langle L \rangle$,
- if $Q_{j'} = \langle\langle U \rangle\rangle'$, choose $B_{j'} = \langle U \rangle'$, and

- if $Q_{j'} = \langle \langle \ell(\mathcal{C}) \rangle, \langle \ell(\mathcal{C} \setminus \{q\}) \rangle' \rangle$ for some $\mathcal{C} \subseteq \mathfrak{C}(z)$ containing q , choose $B_{j'} = \langle \ell(\mathcal{C}) \rangle$.

If there is no $1 \leq j \leq r$ satisfying $Q_j = \langle \langle U \rangle' \rangle$, then choose the $B_{j'}$ as above, with the only difference that if there is at least one j'' with $Q_{j''} = \langle \langle \ell(\mathcal{C}) \rangle, \langle \ell(\mathcal{C} \setminus \{q\}) \rangle' \rangle$ for some $\mathcal{C} \subseteq \mathfrak{C}(z)$ containing q , then choose one such j'' arbitrarily, and set $B_{j''} = \langle \ell(\mathcal{C} \setminus \{q\}) \rangle'$ for this j'' . By Corollary 5.2 and the definition of $\langle \cdot \rangle$, each chosen $B_{j'}$ is indeed contained in $Q_{j'}$; by the fact that $Q_{j'} \notin \{ \langle \langle P \rangle \rangle, \langle \langle X \rangle' \rangle \}$ for all $1 \leq j' \leq r$ (due to Condition 3a and the above discussion) and the definition of Σ_q^* , it is also guaranteed that we chose some $B_{j'}$ for each $1 \leq j' \leq r$.

The configuration $\mathcal{B} := B_1 \dots B_r$ is not necessarily contained in \mathcal{E}_q^* ; however, we will show that with some additional changes (that preserve containment of the B_k in the respective Q_k), we can transform \mathcal{B} into a configuration that is contained in \mathcal{E}_q^* . Before explaining this transformation, we collect some properties of \mathcal{B} in the following. From the definition of \mathcal{B} , it follows that

$$(5.10) \quad B_j \notin \{ \langle P \rangle, \langle P \rangle', \langle X \rangle, \langle X \rangle' \} \text{ for each } 1 \leq j \leq r.$$

From Condition 3b and the definition of $\Sigma_{\mathcal{R}}^1$, it follows that

$$(5.11) \quad \text{there is at most one index } 1 \leq j \leq r \text{ such that } B_j \in \Sigma_{\mathcal{R}}^1.$$

From Condition 3c, it follows that

$$(5.12) \quad \text{there are at most } z_q \text{ indices } j \in \{1, \dots, r\} \text{ such that } B_j \in \{ \langle \ell(\mathcal{C}) \rangle, \langle \ell(\mathcal{C}) \rangle' \mid q \in \mathcal{C} \subseteq \mathfrak{C}(z) \}.$$

From Condition 3d, it follows that

$$(5.13) \quad \begin{aligned} &\text{for each } 1 \leq i \leq z \text{ satisfying } i \neq q, \text{ there are at most } z_i \text{ indices} \\ &j \in \{1, \dots, r\} \text{ such that } B_j \in \{ \langle \ell(\mathcal{C}) \rangle, \langle \ell(\mathcal{C}) \rangle' \mid i \in \mathcal{C} \subseteq \mathfrak{C}(z) \}. \end{aligned}$$

Now, we will perform the aforementioned transformation on \mathcal{B} . Observe that, by interpreting U as $\ell(\emptyset)$, we have $B_j \in \{ \langle \ell(\mathcal{C}) \rangle, \langle \ell(\mathcal{C}) \rangle' \mid \mathcal{C} \subseteq \mathfrak{C}(z) \}$ for each $1 \leq j \leq r$, by (5.10) and the definition of Σ_q^* . Now, iterate through the colors in $\mathfrak{C}(z)$. When processing color $i \in \mathfrak{C}(z)$, repeat the following until there are exactly z_i indices $j \in \{1, \dots, r\}$ such that $B_j \in \{ \langle \ell(\mathcal{C}) \rangle, \langle \ell(\mathcal{C}) \rangle' \mid i \in \mathcal{C} \subseteq \mathfrak{C}(z) \}$: choose some B_j satisfying $B_j = f(\ell(\mathcal{C}))$ for some $f \in \{ \langle \cdot \rangle, \langle \cdot \rangle' \}$ and some \mathcal{C} that does not contain i , and replace it with $f(\ell(\mathcal{C} \cup \{i\}))$. This is possible due to (5.12), (5.13), and the fact that $z_i \leq r - 1$ for each $i \in \mathfrak{C}(z)$. Note that each step in the transformation preserves the correctness of (5.10), (5.11), (5.12), and (5.13). Observe also that in each replacement performed during the transformation we replace some label with a stronger label, by Lemma 5.5. By the right-closedness of all labels in Σ_q^* (which implies that all Q_j are right-closed), it follows that the transformation preserves that $B_j \in Q_j$ for each $1 \leq j \leq r$.

From the above construction and discussion, we conclude that there is some $(B_1, \dots, B_r) \in Q_1 \times \dots \times Q_r$ such that $B_1 \dots B_r$ is a permutation of some configuration $\langle L_1 \rangle' \langle L_2 \rangle \langle L_3 \rangle \dots \langle L_r \rangle$ satisfying $L_j \in \Sigma(z, s) \setminus \{D, M\}$ for all $1 \leq j \leq r$ and Condition 2 in Lemma 5.4, i.e., such that $B_1 \dots B_r \in \mathcal{E}_q^*$. This concludes the proof that any configuration satisfying at least one of the three conditions stated in the lemma is contained in \mathcal{E}_q^* .

Now we will show the other direction, i.e., that any configuration from \mathcal{E}_q^* satisfies at least one of the three conditions stated in the lemma. We will do so by showing the contrapositive, i.e., that any configuration (with labels from Σ_q^*) that violates all three conditions is not contained in \mathcal{E}_q^* . Let $\mathcal{Q} = Q_1 \dots Q_r$ be an arbitrary configuration violating all three conditions, i.e., \mathcal{Q} violates Conditions 1 and 2, and at least one of Conditions 3a, 3b, 3c, and 3d.

We will first make use of the fact that \mathcal{Q} violates Conditions 1 and 2. If there is some $1 \leq j \leq r$ such that $Q_j = \langle \langle X \rangle' \rangle$, then the fact that \mathcal{Q} violates Condition 1 implies that there is some $j' \neq j$ satisfying $Q_{j'} \in \{ \langle \langle U \rangle' \rangle, \langle \langle X \rangle' \rangle \}$, which in turn implies that for any choice $(B_1, \dots, B_r) \in Q_1 \times \dots \times Q_r$, we have $B_j, B_{j'} \in \Sigma_{\mathcal{R}}^1$, by Corollary 5.2. Since, by definition, each configuration from $\mathcal{E}_{\mathcal{R}}$ contains at most one label from $\Sigma_{\mathcal{R}}^1$, it follows, by the definition of \mathcal{E}_q^* , that $\mathcal{Q} \notin \mathcal{E}_q^*$, and we are done. Similarly, if there is some $1 \leq j \leq r$ such that $Q_j = \langle \langle X \rangle \rangle$, then the fact that \mathcal{Q} violates Condition 2 implies that there are two distinct indices $j', j'' \in \{1, \dots, r\} \setminus \{j\}$ such that $Q_{j'}, Q_{j''} \in \{ \langle \langle U \rangle' \rangle, \langle \langle X \rangle' \rangle \}$, which analogously to above implies that $\mathcal{Q} \notin \mathcal{E}_q^*$, and we are done. Hence, assume in the following that $Q_j \in \Sigma_q^* \setminus \{ \langle \langle X \rangle \rangle, \langle \langle X \rangle' \rangle \}$, which by Corollary 5.2 and the definition of $\langle \cdot \rangle$ implies that

$$(5.14) \quad Q_j \subseteq \Sigma_{\mathcal{R}} \setminus \{ \langle X \rangle, \langle X \rangle' \} \text{ for each } 1 \leq j \leq r.$$

Now, we will make use of the fact that \mathcal{Q} violates at least one of Conditions 3a, 3b, 3c, and 3d. Consider first the case that \mathcal{Q} violates Condition 3a, i.e., there is some $1 \leq j \leq r$ satisfying $\mathcal{Q}_j = \langle\langle P \rangle\rangle$. It follows, by Corollary 5.2 and (5.14), that for any choice $(B_1, \dots, B_r) \in Q_1 \times \dots \times Q_r$, we have $B_j \in \{\langle P \rangle, \langle P' \rangle\}$ for some $1 \leq j \leq r$ and $B_{j'} \notin \{\langle X \rangle, \langle X' \rangle\}$ for each $1 \leq j' \leq r$. By Conditions 1 and 2a in Lemma 5.4, it follows that $B_1 \dots B_r \notin \mathcal{E}_{\mathcal{R}}$ for all $(B_1, \dots, B_r) \in Q_1 \times \dots \times Q_r$, which in turn implies $\mathcal{Q} \notin \mathcal{E}_q^*$, as desired.

Now, consider the case that \mathcal{Q} violates Condition 3b, i.e., there are at least two indices $j, j' \in \{1, \dots, r\}$ such that $\mathcal{Q}_j \in \{\langle\langle U \rangle\rangle, \langle\langle X \rangle\rangle\}$. As already seen above, this implies $\mathcal{Q} \notin \mathcal{E}_q^*$, as desired.

Next, consider the case that \mathcal{Q} violates Condition 3c, i.e., there are at least $z_q + 2$ indices $j \in \{1, \dots, r\}$ such that $\mathcal{Q}_j = \langle\langle U \rangle\rangle$ or $\mathcal{Q}_j = \langle\langle \ell(C) \rangle\rangle, \langle\langle \ell(C \setminus \{q\}) \rangle\rangle$ for some color set C containing color q . Consider an arbitrary choice $(B_1, \dots, B_r) \in Q_1 \times \dots \times Q_r$. If there exists some index $1 \leq j \leq r$ satisfying $B_j \in \{\langle P \rangle, \langle P' \rangle\}$, then (5.14) implies that $B_1 \dots B_r \notin \mathcal{E}_{\mathcal{R}}$, by Conditions 1 and 2a in Lemma 5.4; it follows that $\mathcal{Q} \notin \mathcal{E}_q^*$, and we are done. Hence, assume in the following that $B_j \notin \{\langle P \rangle, \langle P' \rangle\}$ for each $1 \leq j \leq r$. Since, by definition, $\mathcal{E}_{\mathcal{R}}$ contains at most one label from $\Sigma_{\mathcal{R}}^1$, it follows, by Lemma 5.5 and Corollary 5.2, that there are at least $(z_q + 2) - 1$ indices $j \in \{1, \dots, r\}$ such that $B_j = \langle\ell(C)\rangle$ for some color set $C \subseteq \mathcal{C}(z)$ that is a superset of a set $C' \subseteq \mathcal{C}(z)$ containing color q . In other words, there are at least $z_q + 1$ indices $j \in \{1, \dots, r\}$ such that $B_j = \langle\ell(C)\rangle$ for some color set $C \subseteq \mathcal{C}(z)$ containing color q . Combining this conclusion with the fact that, by (5.14), $B_{j'} \notin \{\langle X \rangle, \langle X' \rangle\}$ for each $1 \leq j' \leq r$, we obtain that $B_1 \dots B_r \notin \mathcal{E}_{\mathcal{R}}$, by Conditions 1 and 2b in Lemma 5.4. It follows that $\mathcal{Q} \notin \mathcal{E}_q^*$, as desired.

Finally, consider the case that \mathcal{Q} violates Condition 3d, i.e., there exists some color $i \in \mathcal{C}(z)$ satisfying $i \neq q$ such that there are at least $z_i + 1$ indices $j \in \{1, \dots, r\}$ such that $\mathcal{Q}_j = \langle\langle \ell(C) \rangle\rangle$ for some color set C containing color i . Consider an arbitrary choice $(B_1, \dots, B_r) \in Q_1 \times \dots \times Q_r$. Analogously to the previous case, we obtain that the existence of an index $1 \leq j \leq r$ satisfying $B_j \in \{\langle P \rangle, \langle P' \rangle\}$ implies that $\mathcal{Q} \notin \mathcal{E}_q^*$. On the other hand, the nonexistence of such an index implies, by Corollary 5.2, that there are at least $z_i + 1$ indices $j \in \{1, \dots, r\}$ such that $B_j \in \{\langle\ell(C)\rangle, \langle\ell(C)'\rangle\}$ for some color set $C \subseteq \mathcal{C}(z)$ containing color i ; by (5.14) and Conditions 1 and 2b in Lemma 5.4, it follows that $B_1 \dots B_r \notin \mathcal{E}_{\mathcal{R}}$, which in turn implies that $\mathcal{Q} \notin \mathcal{E}_q^*$, as desired. \square

5.7 Renaming We now show that Π_q^* can be relaxed to $\Pi(z', s')$ for some parameters z' and s' .

LEMMA 5.1. *Let $z = (z_1, \dots, z_k)$ be a vector of $2 \leq k \leq \Delta - 1$ nonnegative integers $z_i \leq r - 1$ and let $s \in \{1, \dots, k\}$. Let q be an integer satisfying $q \neq s$, $q \in \{1, \dots, k\}$, and $z_q \leq r - 2$. Then, the problem $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$ can be relaxed to $\Pi(z', q)$, for $z' = (z'_1, \dots, z'_k)$, where $k = \text{len}(z)$, $z'_i = z_i + 1$ if $i = q$, and $z'_i = z_i$ otherwise.*

Proof. We show that Π_q^* , which we proved to be a relaxation of $\overline{\mathcal{R}}(\mathcal{R}(\Pi(z, s)))$, is equivalent to $\Pi(z', q)$, if we rename the labels correctly. Consider the following renaming:

$$\begin{aligned}
 \langle\langle X \rangle\rangle &\rightarrow M \\
 \langle\langle P \rangle\rangle &\rightarrow P \\
 \langle\langle U \rangle\rangle &\rightarrow U \\
 \langle\langle U \rangle\rangle' &\rightarrow D \\
 \langle\langle X \rangle\rangle &\rightarrow X \\
 \langle\langle \ell(C) \rangle\rangle \text{ for each } \emptyset \neq C \subseteq \mathcal{C}(z) \text{ satisfying } q \notin C &\rightarrow \ell(C) \\
 \langle\langle \ell(C) \rangle\rangle, \langle\langle \ell(C \setminus \{q\}) \rangle\rangle' \text{ for each } \emptyset \neq C \subseteq \mathcal{C}(z) \text{ satisfying } q \in C &\rightarrow \ell(C)
 \end{aligned}$$

Observe that, under this renaming, \mathcal{N}_q^* , as defined in Definition 5.1, becomes equal to $\mathcal{N}(z', q)$, and \mathcal{E}_q^* , as characterized in Lemma 5.7, becomes equal to $\mathcal{E}(z', q)$. \square

6 Deterministic Upper Bounds for Hypergraph MIS

In this section, we present deterministic algorithms for solving the maximal independent set problem on hypergraphs of maximum degree Δ and rank r . We start with a straightforward algorithm that solves the problem in $O(\Delta r + \log^* n)$ rounds. We then describe an algorithm that has a quadratic dependency on Δ , but only a logarithmic dependency on r , achieving a time complexity of roughly $O(\Delta^2 \log r + \log^* n)$ rounds. Finally, we present an algorithm that, perhaps surprisingly, almost does not depend at all on r , and that runs in $O(f(\Delta) \cdot \log^*(\Delta r) + \log^* n)$ deterministic rounds, for some function f .

Straightforward Algorithm. We start by describing a simple algorithm that solves the MIS problem on hypergraphs in $O(\Delta r + \log^* n)$ deterministic rounds. The algorithm works as follows: we first compute a coloring of the nodes using $O(\Delta r)$ colors; then, we go through color classes and add a node to the MIS if and only if it does not have any incident hyperedge e with $\text{rank}(e) - 1$ incident nodes already in the MIS.

It is easy to see that, for computing the coloring, we can compute a distance-2 $(\bar{\Delta} + 1)$ -coloring on the incidence graph G of H , where $\bar{\Delta} \leq \Delta r$ is the maximum degree in the power graph G^2 . This can be done in $\tilde{O}(\sqrt{\Delta r} + \log^* n)$ rounds [59, 27]. Now we can go through color classes and safely add a node (of H) in the set if, by doing so, we still satisfy the desired constraints of the problem, and this requires $O(\Delta r)$ rounds. Hence, in total, we get a time complexity of $O(\Delta r + \log^* n)$ rounds. Therefore, we obtain the following theorem.

THEOREM 6.1. *The MIS problem on hypergraphs can be solved in $O(\Delta r + \log^* n)$ deterministic rounds on hypergraphs with maximum degree Δ and rank r .*

6.1 Slow-in- Δ Algorithm We now introduce an algorithm that solves the MIS problem on hypergraphs in a time that has a slow dependency on Δ , but only a logarithmic dependency on r . More precisely, we prove the following.

THEOREM 6.2. *In the LOCAL model, the MIS problem on n -node hypergraphs of maximum degree Δ and rank r can be solved in $O(\Delta^2 \log r + \Delta \log r \log^* r + \log^* n)$ deterministic rounds.*

In this algorithm, each node is either *active* or *inactive*. The algorithm terminates when all nodes are inactive. At the beginning, each node is in the active state. The algorithm is as follows.

1. Compute an $O((\Delta r)^2)$ coloring of the nodes of H .
2. Let H' be the hypergraph induced by active nodes. Active nodes construct a virtual graph G by splitting each hyperedge e of H' into $\lceil \text{rank}_{H'}(e)/2 \rceil$ virtual hyperedges of rank at most 2.¹²
3. Color nodes of G with $O(\Delta)$ colors.
4. In the hypergraph H' , go through color classes, and while processing a node v , do the following.
 - (a) If v has at least one incident hyperedge e having $\text{rank}_{H'}(e) - 1$ nodes already in the MIS, v considers itself outside the MIS and becomes inactive.
 - (b) If, for each hyperedge e incident to v , there exists another node u incident to e such that $\text{color}(v) \neq \text{color}(u)$ and u is not in the MIS (either because it got processed and decided to not enter the MIS, or because it has not been processed yet), then v enters the MIS and becomes inactive.
 - (c) Otherwise, v does nothing and remains active.
5. Repeat from Item 2 until all nodes are inactive.

Correctness and Time Complexity. Nodes can spend $O(\log^* n)$ rounds to compute the initial coloring in Item 1 [55].

Nodes can create the virtual graph G as in Item 2 in constant time. Then, the coloring described in Item 3 can be done in $\tilde{O}(\sqrt{\Delta} + \log^*(\Delta r))$ rounds, by exploiting the precomputed initial coloring [59, 27]. Notice that, while this is a proper coloring of the virtual graph, this results into a defective coloring of the hypergraph H' , and this defective coloring has some desired properties. In fact, since in G we have a proper coloring, then every pair of nodes connected to the same virtual edge must have different colors, implying that, for each hyperedge e of H' , at most $\lceil \text{rank}_{H'}(e)/2 \rceil$ incident nodes have the same color.

In Item 4 we process nodes going through color classes, and based on some conditions that can be verified in constant time, nodes decide what to do: enter or not the MIS and become inactive, or be still undecided and remain active. Hence, since we have $O(\Delta)$ many colors, this part can be executed in $O(\Delta)$ rounds. While

¹²This is a standard procedure, where nodes can easily agree on which nodes are connected to a given virtual edge: for example, there is a virtual edge created from the hyperedge e that connects the two nodes that, among all nodes incident to e , are the ones with the smallest ID; then there is another virtual edge that connects the second pair of nodes with smallest ID, and so on.

the coloring in the virtual graph G is a proper one, the coloring in H' is not a proper coloring, and therefore neighboring nodes that have the same color get processed at the same time, but this is not an issue. In fact, Item 4a and Item 4b guarantee that, for each hyperedge e , at most $\text{rank}(e) - 1$ incident nodes enter the MIS. Also, Item 4b guarantees that, as long as it is safe to enter the MIS, a node will do so. All nodes that remain still active in Item 4c are basically nodes that still have to decide whether to join or not the MIS. A node v remains active if both of the following conditions hold:

- for each incident hyperedge e it holds that at most $\text{rank}_{H'}(e) - 2$ nodes incident to e are in the MIS (otherwise Item 4a would apply);
- there must exist a hyperedge e incident to v such that all active nodes incident to e have the same color as v (otherwise Item 4b would apply).

Since, as discussed above, for each hyperedge at most half of its incident nodes have the same color, this means that, in the hypergraph induced by active nodes, for each node it holds that there exists at least one incident hyperedge that in H' has rank x and in the hypergraph induced by active nodes of the next phase has rank at most $\lceil x/2 \rceil$. Hence, for each node, in each phase, the rank of one incident hyperedge at least halves. This means that, after repeating this process for $O(\Delta \log r)$ many times, we get that all nodes are inactive, achieving a total runtime of $O(\Delta \log r(\Delta + \log^*(\Delta r)) + \log^* n) = O(\Delta \log r(\Delta + \log^* r) + \log^* n)$, and hence proving Theorem 1.3.

6.2 (Almost) Independent-of- r Algorithm In this section, we present an algorithm for solving the MIS problem on hypergraphs, the running time of which almost does not depend at all on the rank r . We will first start by presenting an algorithm that, assuming we are given an $O((\Delta r)^2)$ -coloring of the nodes, computes an MIS in $O(\log^*(\Delta r))$ rounds, in the case where the maximum degree of the nodes is 2. Then we show how to generalize these ideas and present an algorithm for the general case of $\Delta > 2$.

6.2.1 The Case $\Delta = 2$. Let H be a hypergraph with maximum degree 2 and rank r , where we are given in input an $O((\Delta r)^2)$ -coloring of the nodes (we will later get rid of this assumption and actually compute this coloring, but for now, assume we have it for free). We assume that all nodes have degree exactly 2. In fact, nodes of degree 1 can be easily handled: at the beginning we remove them from the hypergraph, then we solve the problem, and then we add all of them in the solution (except for the case in which a hyperedge is only incident to such nodes, in which case we leave one out).

Since nodes in H have degree 2, we can see them as edges, and create a virtual graph G with maximum degree r , where nodes are the hyperedges of H , and there is an edge between two nodes in G if and only if the hyperedge that they correspond to share a node in H . Note that G may have parallel edges (if H is not a linear hypergraph). Observe that any T -rounds algorithm that is designed to run in G can be simulated by nodes in H in $O(T)$ -rounds. In the following, we show an algorithm that computes a subset S of the edges of G and then show that this set is an MIS for H . On a high level, an edge e of G will result outside the set S if and only if there exists one endpoint v of e such that it has already $\deg(v) - 1$ incident edges in S . Initially, let $S = \emptyset$. The algorithm works as follows.

1. Discard parallel edges. That is, for any pair of nodes $\{u, v\}$ connected by at least one edge, we keep exactly one edge.
2. Compute a $(2, 2)$ -ruling edge set of G (i.e., a $(2, 2)$ -ruling set of the line graph of G).
3. Each node v marks itself with its distance to the nearest edge in the ruling set, resulting with nodes marked with a number in $\{0, 1, 2\}$. More precisely:
 - if v is an endpoint of a ruling-set edge, then v is marked with 0;
 - otherwise, if there is a neighbor of v that is an endpoint of a ruling-set edge, then v is marked with 1;
 - otherwise, v is marked with 2 (note that, since we have a $(2, 2)$ -ruling set of the edges, there is no node at distance larger than 2 from a ruling-set edge, and nodes marked 2 form an independent set).
4. Each node marked with 2 proposes to an arbitrary neighbor (which is marked with 1) to add the edge between them in the ruling set.

5. Each node marked with 1 that receives at least one proposal from a 2-marked node, accepts exactly one of them and rejects the others (choosing arbitrarily).
6. Add to the ruling set the edges over which the proposals are accepted (we still obtain a $(2, 2)$ -ruling set).
7. Nodes recompute the distances to the nearest ruling-set edge (note that each node now gets marked with either 0 or 1).
8. Put back removed parallel edges. None of them is added to the ruling set. Observe that we still have a $(2, 2)$ -ruling set and the distances of the nodes from edges of the ruling set does not change.
9. The set S contains the edges selected according to the following rules.
 - (a) Add to S all edges not in the ruling set.
 - (b) Each 1-marked node removes from S one incident edge towards a neighbor marked 0, breaking ties arbitrarily (note that there must exist at least one such edge).
 - (c) For each edge in the ruling set, add it to S if and only if both endpoints have at least two incident edges not in S .

Correctness and Time Complexity. We now show that the set S constructed by the above algorithm satisfies the following claim.

CLAIM 6.1. *Each node u has at most $\deg(u) - 1$ incident edges in S , and for each edge not in S it holds that at least one of its endpoints has all other incident edges in S .*

Firstly, we argue that, if we prove Claim 6.1, then it directly implies that our solution computes an MIS in H . Recall that, by construction, nodes of G correspond to the hyperedges of H , and edges of G correspond to the nodes of H . Hence, if the above claim is true, it means that the algorithm returns a subset of nodes of H such that, for each hyperedge e , at most $\text{rank}(e) - 1$ incident nodes are selected, and for each non-selected node v , there exists a hyperedge e incident to v that has $\text{rank}(e) - 1$ incident selected nodes. This, by definition, means that the set S in G corresponds to an MIS in H . Hence, in order to show the correctness, it is enough to prove the above claim, which we do in the following.

A $(2, 2)$ -ruling set of the edges (Item 2) can be computed using the algorithm of [47], which runs in $O(\log^*(\Delta r))$ rounds if we are given in input an $O((\Delta r)^2)$ -coloring. Let X be the set of all nodes that are an endpoint of an edge in the ruling set. Nodes can spend $O(1)$ rounds and mark themselves with the distance from the nearest node in X (Item 3). It is easy to see that a $(2, 2)$ -ruling set of the edges of G implies that

- each node in G is at distance at most 2 from a node in X , and
- nodes marked with 2 form an independent set (that is, all their neighbors are marked with 1).

In Item 4, Item 5, and Item 6, we modify the ruling set obtained in Item 2 such that each node is at distance at most 1 from a node in X . For this, nodes marked with 2 propose to a 1-marked neighbor to put the edge between them into the ruling set, and 1-marked nodes accept only one proposal (if they receive any). Note that, since edges between 1-marked and 2-marked nodes are not incident to the edges of the ruling set, and since 1-marked nodes accept only one proposal, then we still obtain a $(2, 2)$ -ruling set. Let u be the node marked 2 that performed a proposal to its neighbor v marked with 1. If the proposal of u gets accepted, then the edge (u, v) enters the ruling set and nodes u and v will be marked with 0 in Item 7. Otherwise, if the proposal gets rejected, it means that v accepted another proposal from another 2-marked neighbor, meaning that, in Item 7, node v will be marked with 0 and hence node u will be marked with 1. All these operations can be done in $O(1)$ rounds. Therefore, we remain with a $(2, 2)$ -ruling set of the edges where each node is either marked with 0 or with 1.

Now, in Item 8, we put back the parallel edges removed in Item 1, and we put none of them in the ruling set. Observe that we still have a $(2, 2)$ -ruling set, and the distances of the nodes from the nearest ruling set edge does not change.

Then, we start constructing our set S according to the rules in Item 9, which can all be accomplished in constant time. Since each 0-marked node u has exactly one incident edge in the ruling set, Item 9a guarantees

that at most $\deg(u) - 1$ incident edges are in S . On the other hand, since each 1-marked node v has at least one incident edge towards a 0-marked neighbor, Item 9b guarantees that exactly $\deg(v) - 1$ incident edges are in S . Hence, nodes satisfy Claim 6.1. Also, edges incident to 1-marked nodes satisfy Claim 6.1, since each marked 1 node has exactly $\deg(v) - 1$ incident edges in S . The only remaining edges to be analyzed are the ones in the ruling set, and according to Item 9c, they enter the set S if they can. Hence, the set S satisfies Claim 6.1. Regarding the running time, except for the first item, all other items can be performed in $O(1)$ rounds, achieving a total runtime of $O(\log^*(\Delta r))$ rounds. This implies the following lemma.

LEMMA 6.1. *The MIS problem can be solved in $O(\log^*(\Delta r))$ deterministic rounds on hypergraphs with maximum degree 2 and rank r , assuming that an $O((\Delta r)^2)$ -coloring is given in input.*

6.2.2 Generalization for $\Delta > 2$ On a high level, we show that, if we have an algorithm that solves MIS on $O((\Delta r)^2)$ -colored hypergraphs of maximum degree $\Delta - 1$ and rank r , then we can use it as a black box for constructing another algorithm that solves MIS on $O((\Delta r)^2)$ -colored hypergraphs of maximum degree Δ and rank r . Hence, by starting from our base-case algorithm for hypergraphs of degree at most 2, and applying this reasoning in an iterative way, we obtain an algorithm for solving MIS on $O((\Delta r)^2)$ -colored hypergraphs of any maximum degree Δ and any rank r .

Therefore, let $\mathcal{A}_{\Delta-1}$ be an algorithm that is able to solve MIS on any $O((\Delta r)^2)$ -colored hypergraph $H_{\Delta-1} = (V_{\Delta-1}, E_{\Delta-1})$ of maximum degree $\Delta - 1$ and rank r . In the following, we present an algorithm that solves MIS on $O((\Delta r)^2)$ -colored hypergraphs $H_{\Delta} = (V_{\Delta}, E_{\Delta})$ with maximum degree Δ and rank r . At the beginning, all nodes are *active*.

1. Compute a $(2, \Delta + 4)$ -ruling set of H_{Δ} .
2. Mark hyperedges with their distance from the nearest node in the ruling set, that is, hyperedges get marked with a number in $\{0, 1, \dots, \Delta + 4\}$. Let E^i be the set of hyperedges marked i .
3. Initialize $\text{MIS} := \emptyset$.
4. For $i = \Delta + 4$ to 0 do:
 - (a) Let S^i be the set of active nodes of V_{Δ} incident to *at least one* i -marked hyperedge, and let $X \subseteq S^i$ be the set of active nodes that have *all* incident hyperedges marked with i .
 - (b) Consider the hypergraph $H^i = (S^i \setminus X, \hat{E}^i)$, where $\hat{E}^i = \{e \cap (S^i \setminus X) \mid e \in E^i\}$, that is, the hypergraph induced by nodes in $S^i \setminus X$ and hyperedges in E^i . This hypergraph has maximum degree $\Delta - 1$. On H^i , simulate $\mathcal{A}_{\Delta-1}$.
 - (c) Nodes that are part of the MIS of H^i join the MIS.
 - (d) All nodes of X join the MIS.
 - (e) All nodes of S^i become passive.
 - (f) All active nodes that are incident to at least one hyperedge e having $\text{rank}_{H_{\Delta}}(e) - 1$ incident nodes in the MIS become passive.

Correctness. In the algorithm, after performing the ruling set on the hypergraph (Item 1), and after marking the hyperedges with their distance to the nearest ruling-set node (Item 2), we start to construct our solution. We go through the marked hyperedges in order, from the largest to the smallest, and Items 4a and 4f guarantee that all nodes incident to at least one hyperedge marked at least $i + 1$ are now passive.

Let H^m be the hypergraph induced by E^i , that is, the hypergraph containing all hyperedges marked i and active nodes incident to at least one of them. This hypergraph has maximum degree Δ , but then, we construct H^i by removing from it all nodes of degree exactly Δ (these nodes are then part of X). Hence, H^i has maximum degree $\Delta - 1$. Therefore we can compute an MIS of H^i using algorithm $\mathcal{A}_{\Delta-1}$ (Item 4b), and obtain that, in H^i , each hyperedge e has at most $\text{rank}_{H^i}(e) - 1$ incident nodes in the MIS.

We claim that the operations performed in Items 4c and 4d produce an MIS in H^m . By construction, the set is an MIS in H^i , and hence an independent set in H^m . Since we add all the nodes of X , we only need to show that we do not ruin the independence property. Observe that, for each i -marked hyperedge, it holds that there

exists at least one incident node connected to an $(i - 1)$ -marked hyperedge. Hence, each hyperedge e is connected to at most $\text{rank}_{H^{m_i}}(e) - 1$ nodes in X , or in other words, it is never the case that a hyperedge e , in H^{m_i} , has only neighbors that are part of X . Also, for each i -marked hyperedge incident to $x > 0$ nodes in X , it holds that at most $\text{rank}_{H^{m_i}}(e) - x - 1$ incident nodes are selected. Hence, by adding the nodes in X , we still obtain that at most $\text{rank}_{H^{m_i}}(e) - 1$ incident nodes are selected.

Now we need to argue that the solution that we get is indeed an MIS of H_Δ , and we show this by induction. As a base case of our inductive argument, as just showed, our algorithm computes an MIS of the hypergraph induced by hyperedges marked with $\Delta + 4$. Hence, suppose that we have an MIS $M^{>i}$ on the hypergraph induced by hyperedges marked at least $i + 1$. We show that, after computing an MIS M^i of the hypergraph induced by i -marked hyperedges, we get an MIS on the hypergraph induced by hyperedges marked at least i . In order to see that this holds it is enough to notice that, at each phase of the for-loop, we consider all nodes that could potentially be part of the MIS, taking into consideration the already computed partial solution. More precisely, when we execute Item 4f in phase $i + 1$, we ensure that, in phase i , we construct the MIS set M^i by considering *all* nodes that are incident to at least one i -marked hyperedge, and that can potentially be part of the MIS. Therefore, the set $M^{>i} \cup M^i$ is an MIS of the hypergraph induced by hyperedges marked at least i . Hence, after the $\Delta + 5$ steps performed in the for-loop, we get an MIS of our hypergraph H_Δ .

Time Complexity. We now provide a recursive formula for the above algorithm and prove the following lemma.

LEMMA 6.2. *Let $\mathcal{A}_{\Delta-1}$ be a T -rounds algorithm that solves MIS on $O((\Delta r)^2)$ -colored hypergraphs of maximum degree $\Delta - 1$ and maximum rank r . Then, the algorithm described in Section 6.2.2 solves MIS on $O((\Delta r)^2)$ -colored hypergraphs of maximum degree Δ and maximum rank r in time $O(\Delta(1 + T) + \log^*(\Delta r))$.*

Proof. Item 1 of the algorithm in Section 6.2.2 can be done in $O(\Delta + \log^*(\Delta r))$ rounds by using the algorithm of [47, Corollary 1.6]¹³. Then, nodes spend $O(\Delta)$ rounds to perform Item 2. In the for-loop, we then execute $O(\Delta)$ times $\mathcal{A}_{\Delta-1}$, while the other items of the for-loop can be done in constant time. Hence, we get a total runtime of $O(\Delta(1 + T) + \log^*(\Delta r))$ rounds. \square

Putting Things Together. Finally, we show that the recursive formula that expresses the runtime of our algorithm obtained in Lemma 6.2 can be upper bounded by $2^{O(\Delta \log \Delta)} \log^* r$. We prove the following lemma.

LEMMA 6.3. *The MIS problem can be solved in $2^{O(\Delta \log \Delta)} \log^*(\Delta r) = 2^{O(\Delta \log \Delta)} \log^* r$ deterministic rounds on $O((\Delta r)^2)$ -colored hypergraphs of maximum degree Δ and maximum rank r .*

Proof. Let T_d be the runtime of an algorithm that solves MIS on hypergraphs of maximum degree d and maximum rank r , given an $O((\Delta r)^2)$ -coloring in input. By Lemma 6.1, we know that $T_2 \leq c_1 \log^*(\Delta r)$, and by Lemma 6.2, $T_{d+1} \leq c_1 \Delta(1 + T_d) + c_1 \log^*(\Delta r)$, for some constant c_1 . We show, by induction, that this results in the runtime stated in this lemma.

As a base case we consider hypergraphs with maximum degree 2, that clearly holds by Lemma 6.1. Hence, assuming that $T_\Delta \leq 2^{c\Delta \log \Delta} \log^*(\Delta r)$, we show that $T_{\Delta+1} \leq 2^{c(\Delta+1) \log(\Delta+1)} \log^*(\Delta r)$.

$$\begin{aligned}
 T_{\Delta+1} &\leq c_1 \Delta(1 + T_\Delta) + c_1 \log^*(\Delta r) \text{ (by Lemma 6.2)} \\
 &\leq c_1 \Delta(1 + 2^{c\Delta \log \Delta} \log^*(\Delta r)) + c_1 \log^*(\Delta r) \text{ (by inductive hypothesis)} \\
 &\leq c_1 (\Delta 2^{c\Delta \log \Delta+1} + 1) \log^*(\Delta r) \\
 &\leq 2^{c\Delta \log \Delta + \log c_1 + \log \Delta + 2} \log^*(\Delta r) \\
 &\leq 2^{c\Delta \log \Delta + c_1 \log \Delta} \log^*(\Delta r) \\
 &\leq 2^{c(\Delta+1) \log(\Delta+1)} \log^*(\Delta r), \text{ for a large enough } c.
 \end{aligned}$$

\square

¹³Their algorithm is phrased as a $(2, \ell + 4)$ -ruling hyperedge set on hypergraphs of rank at most ℓ . This algorithm can be trivially converted, in the LOCAL model, into an algorithm that finds a $(2, \ell + 4)$ -ruling set on hypergraphs of degree at most ℓ , by reversing the roles of nodes and hyperedges.

Up until now we have assumed an $O((\Delta r)^2)$ -coloring of the nodes in input. We can compute this coloring in $O(\log^* n)$ rounds using the algorithm of [55], and together with Lemma 6.3, we obtain Theorem 1.4, which is restated here for completeness.

THEOREM 6.3. *In the LOCAL model, the MIS problem on n -node hypergraphs of maximum degree Δ and rank r can be solved in $2^{O(\Delta \log \Delta)} \log^* r + O(\log^* n)$ deterministic rounds.*

7 Open Questions

In this work, we showed that, for hypergraph maximal matching, the $O(\Delta r + \log^* n)$ algorithm is optimal, in the sense that in order to improve the Δr dependency it is required to spend much more as a function of n . Also, we showed that the same does not hold for hypergraph MIS, by providing two algorithms, one with complexity $O(\Delta^2 \log r + \Delta \log r \log^* r + \log^* n)$, and the other with complexity $2^{O(\Delta \log \Delta)} \log^* r + O(\log^* n)$. Hence, if Δ is constant but r very large, we can still solve hypergraph MIS in just $O(\log^* n)$ rounds. Unfortunately, these algorithms do not match the lower bound, since the only known lower bound comes from standard MIS, and says that this problem requires $\Omega(\min\{\Delta, \log_\Delta n\})$ rounds [5, 7]. Observe that this lower bound is tight for MIS on graphs, that can in fact be solved in $O(\Delta + \log^* n)$ rounds, and there is the possibility that also hypergraph MIS could be solved in this time (i.e., there is no lower bound that prevents this).

We now show a possible direction that could lead to solving hypergraph MIS in just $O(\Delta + \log^* n)$ rounds. Consider the following variant of hypergraph coloring, for which our lower bounds do *not* apply.

DEFINITION 7.1. *A c -unique-maximum coloring is a labeling of the nodes of the hypergraph, such that each node has a label in $\{1, \dots, c\}$ and each hyperedge e satisfies that, if C is the set of colors used by the nodes incident to e , then the maximum element of C is used by only one node incident to e .*

We informally state a curious fact about hypergraph MIS. By applying round elimination on it for k times, it seems that we obtain a problem that can be decomposed into three parts:

- The *original* problem;
- A *natural* part, the k -unique-maximum coloring;
- An *unnatural* part, of size roughly equal to a power tower of height k , that cannot be easily understood.

We do not formally prove that this is indeed the result that we get by applying round elimination on hypergraph MIS, but we instead use this observation as a suggestion for a possible algorithm. In fact, we now prove that we can solve hypergraph MIS fast if we are given a suitable unique maximum coloring.

THEOREM 7.1. *If the $O(\Delta)$ -unique-maximum coloring problem can be solved in $O(\Delta + \log^* n)$ rounds, then also hypergraph MIS can be solved in $O(\Delta + \log^* n)$ rounds.*

Proof. We show that, after spending $O(\Delta + \log^* n)$ rounds to compute an $O(\Delta)$ -unique-maximum coloring, then we can spend only $O(\Delta)$ rounds to solve hypergraph MIS.

The algorithm is exactly the same as the trivial algorithm: we process nodes by color classes, from the smallest to the largest, and we add the nodes to the MIS, if possible. This clearly requires $O(\Delta)$ rounds. The unique-maximum property of the coloring guarantees that, for each hyperedge e , when processing the node v of the largest color incident to it, all other nodes in $e \setminus v$ have been already processed, and hence there could be two scenarios:

- either there are some nodes incident to e that did not join the MIS, and in that case v joins the MIS if allowed by the other hyperedges incident to v , or
- all the other incident nodes to e have already been processed and joined the MIS, and hence v decides not to join.

The independence property is guaranteed by the fact that, for each hyperedge e , among all nodes in e , the last one to be processed is v , which does not enter the MIS if all nodes in $e \setminus v$ are in the MIS. The maximality property is obtained for the following reason. If a node does not have the largest color in any of its incident hyperedges, then

it always enters the MIS, since, for each incident hyperedge e , there is at least one node that is not (potentially yet) in the MIS, and that is the node with the largest color in e . This means that, if there is a node v that decides not to join the MIS, then there must exist a hyperedge e such that v is the node with the unique maximum color among all nodes in e , and all nodes in $e \setminus v$ already joined the MIS, hence guaranteeing maximality. \square

We find very fascinating that a problem that has been studied in very different contexts (see, e.g., [21, 22, 35]) appears as a natural subproblem when applying round elimination on hypergraph MIS. Unfortunately, we do not know what is the complexity of $O(\Delta)$ -unique-maximum coloring, and we leave it as an open question to determine its complexity. Note that a $(\Delta + 1)$ -unique-maximum coloring always exists and that it can be computed distributedly in $O(\Delta \cdot T_{\text{hmis}})$ rounds, where T_{hmis} is the time required to compute a hypergraph MIS. In fact, it can be computed by doing the following for $\Delta + 1$ times: at step i compute a hypergraph MIS, color the nodes in the obtained set with color i , remove the colored nodes and the hyperedges of rank 1, and continue in the residual graph. The reason why it takes at most $\Delta + 1$ iterations is that for every node v that remains it must hold that for at least one incident hyperedge all incident nodes except v joined the set, and hence the degree of v decreases by at least 1.

OPEN PROBLEM 7.1. *What is the distributed complexity of $O(\Delta)$ -unique-maximum coloring?*

It could be that our algorithms are actually optimal, and that the current lower bound is not tight. Another possibility is that there is a fast algorithm, that is not based on first computing a unique-maximum coloring. Hence, we leave as open question determining the exact complexity of hypergraph MIS.

OPEN PROBLEM 7.2. *Is it possible to solve hypergraph MIS in $O(\Delta + \log^* n)$? What are the possible tradeoffs between r and Δ in the complexity of hypergraph MIS, when restricting the n dependency to be $O(\log^* n)$?*

Finally, proofs based on round elimination are getting harder and harder. We believe that there is still a lot to be understood about the round elimination technique. Specifically, it would be nice to find a way to prove round elimination statements without writing a tedious case analysis. We leave as an open question finding simpler ways to apply round elimination.

OPEN PROBLEM 7.3. *Is there a way to apply round elimination that does not require to write complicated proofs?*

References

- [1] Geir Agnarsson and Magnús M. Halldórsson. Strong colorings of hypergraphs. In *Approximation and Online Algorithms*, 2005.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. Distributed Δ -coloring plays hide-and-seek. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, (STOC 2018)*, 2022.
- [8] 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.
- [9] Alkida Balliu, Mohsen Ghaffari, Fabian Kuhn, and Dennis Olivetti. Node and Edge Averaged Complexities of Local Graph Problems. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing (PODC)*, 2022.
- [10] 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.

- [11] Leonid Barenboim, Michael Elkin, and Fabian Kuhn. Distributed $(\Delta+1)$ -Coloring in Linear (in Δ) Time. *SIAM Journal on Computing*, 43(1):72–95, 2014.
- [12] Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. The locality of distributed symmetry breaking. In *Proc. 53rd IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 321–330, 2012.
- [13] Paul Beame and Michael Luby. Parallel search for maximal independence given minimal dependence. In *Proc. 1st ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 212–218, 1990.
- [14] Ioana O. Bercea, Navin Goyal, David G. Harris, and Aravind Srinivasan. On computing maximal independent sets of hypergraphs in parallel. *ACM Trans. Parallel Comput.*, 3(1), 2017.
- [15] 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.
- [16] Sebastian Brandt, Orr Fischer, Juho Hirvonen, Barbara Keller, Tuomo Lempiäinen, 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.
- [17] Sebastian Brandt, Barbara Keller, Joel Rybicki, Jukka Suomela, and Jara Uitto. Efficient load-balancing through distributed token dropping. In *Proc. 33rd ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, pages 129–139, 2021.
- [18] Sebastian Brandt and Dennis Olivetti. Truly tight-in- Δ bounds for bipartite maximal matching and variants. In *Proc. 39th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 69–78, 2020.
- [19] Armando Castañeda, Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum, Matthieu Roy, and Corentin Travers. A topological perspective on distributed network algorithms. *Theor. Comput. Sci.*, 849:121–137, 2021.
- [20] Armando Castañeda and Sergio Rajsbaum. New combinatorial topology bounds for renaming: The upper bound. *J. ACM*, 59(1):3:1–3:49, 2012.
- [21] Panagiotis Cheilaris, Balázs Keszegh, and Dömötör Pálvölgyi. Unique-maximum and conflict-free coloring for hypergraphs and tree graphs. In *SOFSEM 2012: Theory and Practice of Computer Science*, 2012.
- [22] Ke Chen, Amos Fiat, Haim Kaplan, Meital Levy, Jiří Matoušek, Elchanan Mossel, János Pach, Micha Sharir, Shakh Smorodinsky, Uli Wagner, and Emo Welzl. Online conflict-free coloring for intervals. *SIAM Journal on Computing*, 36(5):1342–1359, 2007.
- [23] A. Czygrinow and M. Hańćkowiak. Distributed algorithm for better approximation of the maximum matching. In *Proc. 9th Annual Int. Computing and Combinatorics Conf. (COCOON)*, pages 242–251, 2003.
- [24] Andrzej Czygrinow, Michal Hanćkowiak, Edyta Szymańska, and Wojciech Wawrzyniak. Distributed 2-approximation algorithm for the semi-matching problem. In *Proc. 26th Symp. on Distributed Computing (DISC)*, pages 210–222, 2012.
- [25] Manuela Fischer. Improved deterministic distributed matching via rounding. In *Proc. 31st Symp. on Distributed Computing (DISC)*, pages 17:1–17:15, 2017.
- [26] Manuela Fischer, Mohsen Ghaffari, and Fabian Kuhn. Deterministic distributed edge-coloring via hypergraph maximal matching. In *Proc. 58th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 180–191, 2017.
- [27] 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.
- [28] 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.
- [29] Mohsen Ghaffari, Christoph Grunau, and Václav Rozhon. Improved deterministic network decomposition. In *Proc. 32nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2904–2923, 2021.
- [30] 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.
- [31] Mohsen Ghaffari, Fabian Kuhn, and Yannic Maus. On the complexity of local distributed graph problems. In *Proc. 49th ACM Symp. on Theory of Computing (STOC)*, pages 784–797, 2017.
- [32] Mohsen Ghaffari, Fabian Kuhn, Yannic Maus, and Jara Uitto. Deterministic distributed edge-coloring with fewer colors. In *Proc. 50th ACM Symposium on Theory of Computing (STOC)*, pages 418–430, 2018.
- [33] Michal Hanćkowiak, Michal Karonski, and Alessandro Panconesi. On the distributed complexity of computing maximal matchings. In *Proc. 9th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 219–225, 1998.
- [34] Michal Hanćkowiak, Michal Karonski, and Alessandro Panconesi. On the distributed complexity of computing maximal matchings. *SIAM Journal on Discrete Mathematics*, 15(1):41–57, 2001.
- [35] Sarel Har-Peled and Shakh Smorodinsky. Conflict-free coloring of points and simple regions in the plane. *Discret. Comput. Geom.*, 34(1):47–70, 2005.
- [36] David G. Harris. Derandomized concentration bounds for polynomials, and hypergraph maximal independent set. *ACM Trans. Algorithms*, 15(3), 2019.
- [37] David G. Harris. Distributed local approximation algorithms for maximum matching in graphs and hypergraphs.

- SIAM J. Comput.*, 49(4):711–746, 2020.
- [38] David G. Harris, Ehab Morsy, Gopal Pandurangan, Peter Robinson, and Aravind Srinivasan. Efficient computation of sparse structures. *Random Struct. Algorithms*, 49(2):322–344, 2016.
- [39] Maurice Herlihy, Dmitry N. Kozlov, and Sergio Rajsbaum. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann, 2013.
- [40] Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *J. ACM*, 46(6):858–923, 1999.
- [41] J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. Comput.*, 2(4):225–231, 1973.
- [42] S. Hougardy and D. E. Drake Vinkemeier. Approximating weighted matchings in parallel. *Inf. Process. Lett.*, 99(3):119–123, 2006.
- [43] Amos Israeli and A. Itai. A fast and simple randomized parallel algorithm for maximal matching. *Information Processing Letters*, 22(2):77–80, 1986.
- [44] Richard M. Karp, Eli Upfal, and Avi Wigderson. The complexity of parallel search. *J. Comput. Syst. Sci.*, 36(2):225–253, 1988.
- [45] Pierre Kelsen. On the parallel complexity of computing a maximal independent set in a hypergraph. In *Proc. 24th ACM Symp. on Theory of Computing (STOC)*, pages 339–350, 1992.
- [46] 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.
- [47] 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.
- [48] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. What cannot be computed locally! In *Proc. 23rd ACM Symp. on Principles of Distributed Computing (PODC)*, pages 300–309, 2004.
- [49] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local computation: Lower and upper bounds. *Journal of ACM*, 63:17:1–17:44, 2016.
- [50] Fabian Kuhn and Chaodong Zheng. Efficient distributed computation of MIS and generalized MIS in linear hypergraphs. *CoRR*, abs/1805.03357, 2018.
- [51] Shay Kutten, Danupon Nanongkai, Gopal Pandurangan, and Peter Robinson. Distributed symmetry breaking in hypergraphs. In *Proc. 28th Symp. on Distributed Computing (DISC)*, pages 469–483, 2014.
- [52] 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.
- [53] N. Linial and M. Saks. Low diameter graph decompositions. *Combinatorica*, 13(4):441–454, 1993.
- [54] 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.
- [55] Nathan Linial. Locality in Distributed Graph Algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- [56] Z. Lotker, B. Patt-Shamir, and S. Pettie. Improved distributed approximate matching. *J. ACM*, 62(5):38:1–38:17, 2015.
- [57] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.
- [58] Tomasz Luczak and Edyta Szymanska. A parallel randomized algorithm for finding a maximal independent set in a linear hypergraph. *J. Algorithms*, 25(2):311–320, 1997.
- [59] Yannic Maus and Tigran Tonoyan. Local conflict coloring revisited: Linial for lists. In *34th International Symposium on Distributed Computing, DISC*, pages 16:1–16:18, 2020.
- [60] Moni Naor. A lower bound on probabilistic algorithms for distributive ring coloring. *SIAM Journal on Discrete Mathematics*, 4(3):409–412, 1991.
- [61] Dennis Olivetti. Round Eliminator: a tool for automatic speedup simulation, 2019.
- [62] Alessandro Panconesi and Romeo Rizzi. Some simple distributed algorithms for sparse networks. *Distributed Computing*, 14(2):97–100, 2001.
- [63] Alessandro Panconesi and Aravind Srinivasan. On the Complexity of Distributed Network Decomposition. *Journal of Algorithms*, 20(2):356–374, 1996.
- [64] Václav Rozhoň and Mohsen Ghaffari. Polylogarithmic-time deterministic network decomposition and distributed derandomization. In *Proc. 52nd ACM Symp. on Theory of Computing (STOC)*, pages 350–363, 2020.
- [65] Michael E. Saks and Fotios Zaharoglou. Wait-free k-set agreement is impossible: The topology of public knowledge. *SIAM J. Comput.*, 29(5):1449–1483, 2000.