



HAL
open science

Distributed Quantum Advantage for Local Problems

Alkida Balliu, Sebastian Brandt, Xavier Coiteux-Roy, Francesco d'Amore, Massimo Equi, François Le Gall, Henrik Lievonen, Augusto Modanese, Dennis Olivetti, Marc-Olivier Renou, et al.

► **To cite this version:**

Alkida Balliu, Sebastian Brandt, Xavier Coiteux-Roy, Francesco d'Amore, Massimo Equi, et al.. Distributed Quantum Advantage for Local Problems. 2024. hal-04768228

HAL Id: hal-04768228

<https://hal.science/hal-04768228v1>

Preprint submitted on 5 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed Quantum Advantage for Local Problems

Alkida Balliu · Gran Sasso Science Institute, Italy

Sebastian Brandt · CISA Helmholtz Center for Information Security, Germany

Xavier Coiteux-Roy · Technical University of Munich, Germany · Munich Center for Quantum Science and Technology, Germany

Francesco d'Amore · Bocconi University, Italy · BIDSa, Italy

Massimo Equi · Aalto University, Finland

François Le Gall · Nagoya University, Japan

Henrik Lievonen · Aalto University, Finland

Augusto Modanese · Aalto University, Finland

Dennis Olivetti · Gran Sasso Science Institute, Italy

Marc-Olivier Renou · Inria Paris-Saclay, France · CPHT, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris, France

Jukka Suomela · Aalto University, Finland

Lucas Tendick · Inria Paris-Saclay, France

Isadora Veeren · Inria Paris-Saclay, France · CPHT, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris, France

Abstract. We present the first *local* problem that shows a super-constant separation between the classical randomized LOCAL model of distributed computing and its quantum counterpart. By prior work, such a separation was known only for an artificial graph problem with an inherently *global* definition [Le Gall et al. 2019].

We present a problem that we call *iterated GHZ*, which is defined using only local constraints. Formally, it is a family of *locally checkable labeling* problems [Naor and Stockmeyer 1995]; in particular, solutions can be verified with a constant-round distributed algorithm.

We show that in graphs of maximum degree Δ , any classical (deterministic or randomized) LOCAL model algorithm will require $\Omega(\Delta)$ rounds to solve the iterated GHZ problem, while the problem can be solved in 1 round in quantum-LOCAL.

We use the *round elimination* technique to prove that the iterated GHZ problem requires $\Omega(\Delta)$ rounds for classical algorithms. This is the first work that shows that round elimination is indeed able to separate the two models, and this also demonstrates that round elimination *cannot* be used to prove lower bounds for quantum-LOCAL. To apply round elimination, we introduce a new technique that allows us to discover appropriate *problem relaxations* in a mechanical way; it turns out that this new technique extends beyond the scope of the iterated GHZ problem and can be used to e.g. reproduce prior results on maximal matchings [FOCS 2019, PODC 2020] in a systematic manner.

1 Introduction

The key question we study in this work is this: **which problems admit a distributed quantum advantage?** That is, if we have a computer network, and we replace classical computers with quantum computers and classical communication channels with quantum communication channels, which distributed tasks can be now solved faster?

1.1 Setting: LOCAL vs. quantum-LOCAL

There is a range of prior work that has studied such questions in *bandwidth-limited* networks, e.g., comparing the relative power of the classical CONGEST and CONGESTED CLIQUE models with their quantum counterparts [33, 42, 50–52, 55, 61, 62, 64]. In essence, the question is: Can a quantum network which can send b qubits (quantum bits) per time unit over each communication link outperform a classical network which can send the same number of bits per time unit? It turns out that in many cases we can indeed demonstrate quantum advantage in this sense, see, e.g., [33, 50–52].

However, quantum advantage in settings in which *large distances* are the key limitation (instead of bandwidth) are much less understood. We emphasize that large distances (i.e., network latency, or time for information to propagate from one node to another) are a fundamental physical limitation, and not merely an engineering challenge. One can at least in principle increase bandwidth by installing multiple parallel communication channels (say, a larger number of fiber-optic links between a pair of nodes), and it might be significantly harder for quantum communication channels, while there is no way to increase the speed of information propagation beyond the speed of light in a spatially distributed system. Can quantum computation and communication still help? This is our focus in this work.

We study the classical **LOCAL model** of distributed computing, and compare it with its quantum counterpart. In the LOCAL model computation proceeds in synchronous rounds, and in each round each computer can send an *arbitrarily large message* to each neighbor, receive a message from each neighbor, and do *unbounded local computation*. Initially all nodes are only aware of their own local input, and eventually all nodes must stop and announce their own local output (e.g. in the graph coloring problem: what is its own color). We are interested in **how many communication rounds** are needed (in the worst case) until all nodes stop.

There is extensive line of research since the early 1990s on understanding the classical LOCAL model and the round complexity of various graph problems in this model, see, e.g., [2, 6, 7, 9, 12, 13, 17, 26–28, 35, 36, 41, 44–46, 54, 57]. The main focus has been on problems defined with local constraints (e.g. graph coloring), and especially in the recent years one of the main research themes has been understanding the distributed complexity landscape of all such problems.

But do any of these problems strictly benefit from quantum computation and communication? In the quantum-LOCAL model, the local state of a node consists of an arbitrary number of qubits, each message consists of an arbitrary number of qubits, and after each communication round, each node can do arbitrary quantum operations and measurements with the qubits that it holds. Does this help in comparison with the classical randomized LOCAL model, where the nodes have only access to a classical source of randomness?

1.2 Prior work vs. our main contribution

By prior work [53], it is known that there exists a (very artificial) graph problem with the following properties: in n -node graphs, any classical (deterministic or randomized) LOCAL algorithm will

require $\Omega(n)$ rounds to solve the problem, while there is an $O(1)$ -round quantum-LOCAL algorithm for solving it.

However, the problem from prior work is not only artificial, but it is also very different from the problems usually studied in the context of the LOCAL model. In particular, it has got an inherently *global* definition: to decide if a given output is valid, one has to inspect the joint outputs of nodes in distant parts of the networks. Such a problem is very different from problems typically studied in the context of the LOCAL model. Usually we are interested in problems similar to graph coloring, which can be defined with *local* constraints (e.g., a vertex coloring is valid if and only if each node has an output label different from the labels of its neighbors). Put otherwise, these are problems where the solution may be hard to find for a distributed algorithm, but it is at least easy to verify (cf. class FNP in the centralized setting).

The idea of problems defined with local constraints was formalized in the seminal paper by Naor and Stockmeyer in 1995 [57]. They defined the family of *locally checkable labeling* problems, or in brief LCLs. These are problems that can be defined by specifying a finite set of valid labeled neighborhoods. For example, proper vertex coloring with 11 colors in graphs of degree at most 10 is a concrete example of a graph problem that is formally an LCL in the strict sense (very often we study *families* of LCLs parameterized by the maximum degree, and e.g. $(\Delta + 1)$ -coloring in graphs of degree at most Δ is a natural example of such a problem family). Most of the problems of interest in the context of the LOCAL model are (families of) LCL problems and there is a huge body of work that has established all kinds of structural results related to the complexity of LCL problems (see, e.g., [5–8, 15, 16, 18, 18–23, 28, 34–37, 48]). Now the key question is:

Does quantum-LOCAL help in comparison with classical LOCAL for some LCL problem?

Before this work, the answer was: we do not know, beyond constant advantage. There are several *negative* results that exclude any quantum advantage for a specific problem, or at least put limits on it (see, e.g., [3, 40, 43]). It is also easy to engineer a problem that exhibits a *constant* quantum advantage, e.g., a problem that is solvable in 1 round with quantum-LOCAL and requires 2 rounds with classical LOCAL. One can also amplify this if we cheat by defining a family of LCL in which the checking radius (i.e., how local the constraints are, or what is the maximum radius of a valid labeled neighborhood in the definition of the LCL) increases as a function of Δ . But beyond this, *nothing positive was known*; the global graph problem from [53] was the only known example of any asymptotic quantum advantage in the LOCAL model. Our main result is this:

We present a family of LCL problems \mathcal{P}_Δ called **iterated GHZ**, parameterized by the maximum degree Δ , with these three properties:

1. the checking radius of \mathcal{P}_Δ is $O(1)$, independently of Δ ,
2. in quantum-LOCAL, we can solve \mathcal{P}_Δ in $O(1)$ rounds, independently of Δ ,
3. in classical (deterministic or randomized) LOCAL, any algorithm that solves \mathcal{P}_Δ requires $\Omega(\Delta)$ rounds.

This gives the **first local problem** (in any reasonable sense) that demonstrates **distributed quantum advantage** in the LOCAL model, also answering the open questions from [1, 40] positively. We emphasize that the problem is engineered so that the feasibility of solutions is *easy to verify* also for classical distributed algorithms, but it is easy to compute only for distributed quantum algorithms.

1.3 Key techniques and new ideas

As a gentle introduction to our work, let us first introduce a toy version of our problem, called **iterated CHSH**. While this problem does not quite give us distributed quantum advantage, it already contains all the key ingredients; and the final step that we need to take from iterated CHSH to iterated GHZ is relatively straightforward (albeit technical).

CHSH game. The CHSH game [32, 39] is the following two-player game: Alice gets an input bit x and Bob gets an input bit y . Alice has to produce an output bit a and Bob has to produce an output bit b . The bits have to satisfy $xy = a \oplus b$, that is, if both inputs are 1, the players must produce different outputs, while otherwise the players must produce the same output. This game has two important properties:

1. It is “classically hard” in the following sense: Alice and Bob cannot win this game without any communication (for any classical strategy, Alice and Bob will fail on at least one combination (x, y) of inputs, and thus win with probability at most 75% on uniformly random inputs).
2. It is “non-signaling” in the following sense: if a friendly oracle picks a uniformly random valid output, and informs Alice about her output a and informs Bob about his output b , then Alice will learn absolutely nothing about Bob’s input and vice versa.

The non-signaling property implies that there *might* be a quantum strategy that enables Alice and Bob to win this game (with a nontrivial probability); at least the existence of such a strategy would not violate causality. And indeed this is the case: if Alice and Bob share a single entangled qubit pair, they can win the game with up to $(1 + \sqrt{2})/2 \approx 85\%$ on uniformly random inputs, which is better than the classical strategy [38].

Now 85% probability is not going to suffice for us, and this is exactly why this is just a toy version of the real problem. But let us put this aside for now. It may be useful to imagine for a while that we are living in a *super-quantum* world where the laws of physics actually allowed Alice and Bob to win the CHSH game with probability 1 if they hold some magic pre-shared entangled super-qubits. We will eventually get back from the super-quantum world to the usual quantum world. (Slightly more formally, an entangled pair super-qubit is going to be exactly what is commonly called in the literature a *PR-box* first introduced by Popescu and Rohrlich [59], but we do not assume familiarity with this concept.)

Iterated CHSH. Now we are ready to explain what the *iterated CHSH* problem looks like. Fix some maximum degree Δ . We can without loss of generality assume that our input graph is Δ -regular and it is Δ -edge-colored, that is, for each $c = 1, 2, \dots, \Delta$, each node is incident to exactly one edge that is labeled with color c . (We note that we do not really need a promise here: we can simply define the problem so that all nodes that violate this property are unconstrained, and hence hard instances are exactly those graphs in which all nodes satisfy this locally checkable property, and we do not have any unconstrained nodes.)

We define a graph problem in which each node v has to output a sequence of bits $v_0, v_1, \dots, v_\Delta$ that satisfies these constraints:

- For each node v we have $v_0 = 1$.
- If the color of the $e = \{u, v\}$ is c , then $u_{c-1}v_{c-1} = u_c \oplus v_c$.

This can be interpreted as follows:

- Each edge $\{u, v\}$ of color 1 represents an instance of the CHSH game where the inputs are hardcoded to $u_0 = 1$ and $v_0 = 1$, and the outputs are u_1 and v_1 .
- Each edge $\{u, v\}$ of color $c > 1$ represents an instance of the CHSH game where the inputs are u_{c-1} and v_{c-1} and the outputs are u_c and v_c .

That is, the outputs of the CHSH games that we play on the edges of color 1 form the inputs of the CHSH games that we play on the edges of color 2, and the outputs of those games form the inputs of the CHSH games that we play on the edges of color 3, etc., all the way up to the games that we play on the edges of color Δ . We will now study some properties of the iterated CHSH problem.

Iterated CHSH is a local problem. First, we observe that the iterated CHSH problem is a locally verifiable problem. Indeed, each edge represents the constraints of the CHSH game, and if we are given some solution, it is sufficient that we check that for each edge the relevant bits satisfy these constraints; hence the validity of a solution is locally verifiable.

Furthermore, for each fixed maximum degree Δ we satisfy the usual technical requirements of the definition of LCL problems in [57], namely, the set of input labels (Δ edge colors) and the set of output labels (2^Δ bit strings) are finite. In summary, we have a family of LCL problems parameterized by Δ , and the checking radius does not depend on Δ .

Iterated CHSH in super-quantum-LOCAL. Now if we were in the hypothetical super-quantum world in which pre-shared super-qubits enabled Alice and Bob to win the CHSH game with probability 1, we could also solve the iterated CHSH problem in only **one round of super-quantum communication**, as follows.

First, each node locally prepares Δ pairs of entangled super-qubits. Then we use one round of communication so that each node exchanges one entangled super-qubit with each of its neighbors (it is critical here that we can do this in parallel with each neighbor). Additionally, each node can share e.g. its own unique identifier that we can later use for symmetry-breaking purposes. This way for each edge we have prepared two pairs of entangled super-qubits, and with the help of unique identifiers we can consistently agree on which of these pairs to use. In summary, for each node v there are super-qubits $v^1, v^2, \dots, v^\Delta$ so that if $\{u, v\}$ is an edge of color v , super-qubits u^c and v^c form a pair of entangled super-qubits powerful enough to win the CHSH game.

Then all the rest is local computation. Consider a node v . It initializes $v_0 = 1$, and then iteratively repeats the following step for each $c = 1, 2, \dots, \Delta$: Let $e = \{u, v\}$ be the edge of color c incident to v ; in the first round v already learned the unique identifier of u . We check the unique identifiers to decide which of the nodes takes the role of Alice and which takes the role of Bob; let us say that v becomes Alice. Now v knows that it shares an entangled pair of super-qubits (v^c, u^c) with node u , who plays the role of Bob. Alice knows its own input v_{c-1} , and Bob knows its own input u_{c-1} . Node v , Alice, uses its own super-qubit v^c appropriately to figure out its output v_c , and without any communication node u will do the same using its own entangled super-qubit u^c to find u_c .

This way just 1 round of communication, followed by Δ steps of purely local super-quantum operations is enough for all nodes to produce valid solutions. Hence we have constructed a family of LCL problems that is easy in super-quantum-LOCAL. This is not particularly surprising, as the whole intuition here is that the problem is *by construction* easy in the super-quantum world. All the heavy lifting goes to showing that it is hard in the classical world.

Iterated CHSH in classical LOCAL. To gain more intuition on the iterated CHSH problem, let us first make the following observation: in graphs of maximum degree Δ , we can solve this problem in $\Theta(\Delta)$ rounds in the classical deterministic LOCAL model.

The algorithm is very simple: Each node v first initializes $v_0 = 1$. Then for each $c = 1, 2, \dots, \Delta$, we consider all edges $e = \{u, v\}$ of color c . For each such edge, v will send its own identifier and the value of v_{c-1} to u , and conversely u will send its own identifier and u_{c-1} to v . This way u and v both know the values of u_{c-1} and v_{c-1} , plus their own identifiers. Now assume that the identifier of u is smaller than v ; then the nodes will find the lexicographically smallest pair (u_c, v_c) such that this is a valid output for the CHSH game assuming that node u is Alice with input u_{c-1} and node v is Bob with input v_{c-1} . In particular, node v learns the value of v_c and node u learns the value of u_c so that these satisfy the constraints for edge $\{u, v\}$.

In Δ rounds this will result in a valid solution. Essentially, each node plays CHSH games in a sequential order, starting with the game on the edge of color 1, and eventually reaching the game on the edge of color Δ .

Incidentally, it is good to note that both this $\Theta(\Delta)$ -round classical algorithm and the $O(1)$ -round super-quantum algorithm send $\Theta(\Delta)$ messages—each node communicates with each neighbor exactly once. The key difference is that in the super-quantum algorithm, the messages that we send (entangled super-qubits) do not depend on each other, and we can send all of them in the first round, simultaneously in parallel, while in the classical algorithm, the messages that we send in round c depend on the messages sent in round $c - 1$, and hence they cannot be parallelized.

Now what we would like to do is to show that **the sequential strategy is the best possible classical algorithm**, i.e., there is no $o(\Delta)$ -round classical algorithm that in some clever way circumvents the need for coordinating with each neighbor in a sequential manner. But let us first see how we get from the imaginary super-quantum world to the familiar quantum world, and then discuss how to show the analogous statement for the real problem.

From super-quantum to quantum: replacing CHSH with GHZ. So far we have discussed the toy problem, iterated CHSH, that (at best) can prove a separation between the classical world and an imaginary super-quantum world in which the CHSH game can be solved with probability 1. Let us now see how to apply the same idea to separate the classical world and the real quantum world. To do that, we need a suitable non-signaling game with all these properties:

- It is impossible to solve without any communication using classical randomness.
- It can be solved with probability 1 with the help of pre-shared entangled qubits.
- The input and output values come from some finite set, so that we can turn this into an LCL problem.
- The set of possible input values has to be the set of possible output values, so that we can use the output of one game as the input of another game.
- The game does not require any promises on allowed input combinations, so that the game is always well-defined.

The first two conditions define games called quantum pseudo-telepathy [30]. One good candidate to fulfill also the last three conditions is the well-known GHZ game [47, 56]. This is a three-player game, where Alice, Bob, and Charlie each get one input bit, x , y , and z , and each of them has to produce one output bit, a , b , and c . The usual phrasing of the game is as follows: we promise that $x \oplus y \oplus z = 0$, and the requirement is that $a \oplus b \oplus c = 0$ if and only if $x = y = z = 0$. This phrasing of the game comes with a promise, but we can eliminate it by specifying that any output is valid if $x \oplus y \oplus z = 1$.

The good news is that this game is indeed solvable with probability 1 given a suitable entangled quantum state [47]. The bad news is that it is very relaxed: it is not clear if this game can be used to construct any problem that is classically hard. We now follow the blueprint of the iterated CHSH problem; in essence we replace each instance of the CHSH game with an instance of the GHZ game. We obtain a local problem that we call *iterated GHZ*. We prove that this problem is classically hard.

In the CHSH problem, each edge represented a 2-player game, and we need to embed 3-player games. In essence, we need to switch from graphs to 3-uniform hypergraphs (where each hyperedge represents a 3-player game), and then we will use a bipartite graph to represent the hypergraph: we will have *white nodes* of degree Δ that represent players and *black nodes* of degree 3 that represent games. Black nodes are *colored* so that for each $c = 1, 2, \dots, \Delta$ and for each white node there is exactly one black neighbor of color c .

Again the intuition is that we play GHZ games in a sequential order, so that for each white node v , the output of the GHZ game represented by the black neighbor of color $c - 1$ will be the input of the GHZ game represented by the black neighbor of color c .

This time extra care is needed with bootstrapping. In the CHSH setting, we simply assumed that the first game has all inputs hardcoded to 1. If we did the same here, the iterated GHZ problem would become trivial to solve (note that all-1 input implies that all-1 output is valid, which means that all nodes could blindly output a vector of 1 bits). Similarly, all-0 input leads to a trivial problem. To force a nontrivial solution, we make the first game special: black nodes of color 1 represent the task in which we have no inputs, and exactly one of the players Alice, Bob, and Charlie has to output 1 and the other two must output 0. (We note that this is in some sense analogous to a CHSH game in which the inputs are hardcoded to 1, as there exactly one of Alice and Bob has to output 1 and the other one has to output 0.)

Iterated GHZ. Let us now summarize the definition of the iterated GHZ problem. We have a bipartite graph in which white nodes have degree Δ , and black nodes have degree 3. For each $c = 1, 2, \dots, \Delta$ and for each white node there is exactly one black neighbor of color c .

Each white node v has to output a bit vector $v_1, v_2, \dots, v_\Delta$, so that the following holds. Assume that s, t , and u are the three white neighbors of a black node x , and let c be the color of x . We require the following:

- If $c = 1$, then exactly one of s_1, t_1 , and u_1 has to be 1.
- If $c > 1$, then s_c, t_c , and u_c have to be a valid output for the GHZ game with inputs s_{c-1}, t_{c-1} , and u_{c-1} :

$$\begin{aligned} s_c \oplus t_c \oplus u_c &= 0 && \text{if } s_{c-1} + t_{c-1} + u_{c-1} = 0, \\ s_c \oplus t_c \oplus u_c &= 1 && \text{if } s_{c-1} + t_{c-1} + u_{c-1} = 2, \\ s_c \oplus t_c \oplus u_c &\in \{0, 1\} && \text{if } s_{c-1} + t_{c-1} + u_{c-1} \in \{1, 3\}. \end{aligned}$$

This problem can be solved with a one-round quantum strategy, as follows. Black nodes of color $c = 1$ inform their first white neighbor v that it should set $v_1 = 1$, and inform other neighbors to set $v_1 = 0$. Black nodes of color $c > 1$ prepare a suitable three-qubit GHZ state that enables Alice, Bob, and Charlie to win the GHZ game with probability 1. Additionally, black nodes assign the roles (Alice, Bob, Charlie) to their white neighbors, and also send their own color c , so that white nodes are aware of which games they are supposed to play, in which roles and in which order. After this one communication step, everything is ready, each white node v already knows v_1 , it holds $\Delta - 1$ suitably entangled qubits, and it can locally play its own part in $\Delta - 1$ GHZ games and this way discover valid outputs $v_2, v_3, \dots, v_\Delta$.

This problem also admits a $\Theta(\Delta)$ -round classical solution. Black nodes of color 1 start and choose which of their neighbors will set $v_1 = 1$; all other white nodes merely indicate their colors. Then white nodes send v_1 to black neighbors of color 2, black nodes of color 2 can find a valid solution to the GHZ game that they represent, and they reply with v_2 , etc. After Δ iterations (and approximately 2Δ communication rounds) all white nodes know their final output.

However, as we discussed earlier, the GHZ game itself is very relaxed, and the details of the definition of the iterated GHZ game are subtle; minor changes in the way we define the first game lead to trivial problems. It is not at all clear intuitively if this problem is hard for classical algorithms, or if there is some clever way to cheat. This is what we show in this work: no classical $o(\Delta)$ -round algorithm exist for this problem.

Proving the classical lower bound. Now we are given a specific problem with a fairly complicated specification, namely iterated GHZ. There is a simple $\Theta(\Delta)$ -round classical algorithm, and we would like to show that this is also (asymptotically) optimal.

Here it is also critical to note that we are heavily limited in the availability of possible proof techniques that one could apply here. In distributed graph algorithms, one commonly uses, e.g., propagation arguments (fixing the solution in one place implies something in a distant part of the graph) or arguments based on indistinguishability (two neighborhoods look locally identical, so any distributed algorithm has to produce the same outputs). However, many of these techniques are also (directly or indirectly) applicable in the quantum-LOCAL model [40, 43]. If the proof technique would also imply that the problem is hard not only in classical LOCAL but also in quantum-LOCAL, it cannot work here, as we already know that the problem is easy in quantum-LOCAL.

There are very few lower-bound proof techniques that even in principle *could* separate classical LOCAL and quantum-LOCAL. One rare example is the *round elimination* technique [26]. It is promising in the informal sense that at least known proofs that use round elimination do not generalize to quantum-LOCAL, as a direct generalization would violate the no-cloning theorem.

In this work we show that **round elimination is indeed able to separate classical LOCAL and quantum-LOCAL**: we are able to use round elimination to show that iterated GHZ cannot be solved in $o(\Delta)$ rounds in classical (deterministic or randomized) LOCAL, while the problem can be solved in $O(1)$ rounds in quantum-LOCAL. This is both good news and bad news: the good news is that we indeed have a proof technique that can demonstrate distributed quantum advantage, but the bad news is that none of previously-known lower bounds that were proved with round elimination can be directly generalized to quantum-LOCAL.

Round elimination: basic idea. Round elimination is a proof technique that is specifically applicable to the study of LCL problems in the classical LOCAL model. It is commonly used to prove lower bounds (see, e.g., [4, 7, 10–14, 17, 24, 27, 54]), but it can also be used to e.g. synthesize efficient algorithms (see, e.g., [17]). It is closely related to the fanout inflation technique [63] developed to study network nonlocality [60].

In essence, round elimination is a function \mathcal{R} that answers the following question:

Assume that an LCL problem Π admits a (black box) classical LOCAL algorithm A working in T rounds. What is the most general problem $\Pi' = \mathcal{R}(\Pi)$ which can be solved in $T - 1$ rounds using A ?

Round elimination precisely constructs that new problem Π' . Informally, we construct a $(T - 1)$ -round algorithm A' that works as follows: each node outputs the set of all values that A might output, given what we see in the radius- $(T - 1)$ neighborhood (by considering all possible extensions that

we might see at distance exactly T and simulating A for each of them). Now function \mathcal{R} is defined so that assuming that A indeed solved Π , algorithm A' will solve $\mathcal{R}(\Pi)$, and given any solution to $\mathcal{R}(\Pi)$, we can solve the original problem Π with only one extra round.

We note that the definition of A' crucially needs to *clone* the information that we have in the $(T - 1)$ -radius neighborhood. Hence it is natural to expect that round elimination does not extend to quantum algorithms. Yet, other properties of quantum information could (at least in principle) imply that if a quantum LOCAL algorithm can solve Π in T rounds, then there is a quantum algorithm that can solve $\Pi' = \mathcal{R}(\Pi)$ in $T - 1$ rounds. Our main result implies that this cannot be the case.

Round elimination: formalism. To make use of round elimination, we need to have an LCL problem that is written in a specific formalism:

- We have a bipartite graph, where the two sets of nodes are called *active* and *passive*.
- Each active node has to label its incident edges with labels from some finite alphabet Σ .
- We have an *active constraint*, which is a collection of multisets, where each multiset is one possible valid assignment of edge labels incident to an active node.
- Similarly, we have a *passive constraint*, which specifies the collection of multisets that are valid on the edges incident to passive nodes.

Now if we have any problem Π written in this formalism, we can apply in a mechanical way round elimination to obtain another problem $\Pi' = \mathcal{R}(\Pi)$. We emphasize that this part is just syntactic manipulation of the problem description, without any references to any specific model of computation or any distributed algorithm. Yet as explained above, if the round complexity of Π is T rounds in the classical LOCAL model (for any sufficiently small T), then the round complexity of Π' will be exactly $T - 1$ rounds.

Round elimination: basic strategy. To apply round elimination in our task, we could try to construct a sequence of problems $\Pi_0, \Pi_1, \dots, \Pi_T$ such that the following holds:

- Π_0 is the original problem of interest, iterated GHZ,
- $\Pi_{i+1} = \mathcal{R}(\Pi_i)$,
- we can show that Π_T is not trivial, i.e., it cannot be solved in 0 rounds.

Now if we could solve Π_0 in T rounds, it would imply that Π_T can be solved in 0 rounds, but as this is not the case, Π_0 has to require at least $T + 1$ rounds. If we can do this for, e.g., $T = \Delta$, we would have the result that we want: iterated GHZ cannot be solved in $o(\Delta)$ rounds.

However, what typically happens is that $\mathcal{R}(\Pi)$ is *much* more complicated than Π , in the sense that the size of the alphabet of labels used to describe $\mathcal{R}(\Pi)$ may increase *exponentially* compared to the size of the alphabet used to describe Π , and the number of valid multisets in the active and passive constraints may increase accordingly. Even if the original problem Π_0 has got a simple description, there is typically no simple way to describe Π_i , other than saying that it can be obtained from Π_0 with i applications of \mathcal{R} . In particular, if we do not have a “closed-form” description of Π_T , how do we argue that it is not trivial? Hence, what typically happens is that we augment this mechanical procedure with some well-chosen *relaxations*.

Round elimination: relaxations. We say that Π' is a relaxation of Π if given a feasible solution to Π we can construct (in a distributed way, without any communication) a feasible solution to Π' . A simple example of a relaxation is what can be obtained by *merging labels*. Let $a, b \in \Sigma$ be two distinct labels that we can use in Π . To define Π' , we relax all active and passive constraints so that

whenever a can be used, we can also use b , and vice versa; now trivially Π' is at least as easy as Π . Then we can simplify the description of Π' by completely removing label b , as it is redundant.

The revised process now looks as follows. We construct a sequence of problems $\Pi_0, \Pi_1, \dots, \Pi_T$ such that the following holds:

- Π_0 is some relaxation of the original problem of interest, iterated GHZ,
- Π_{i+1} is some relaxation of $\mathcal{R}(\Pi_i)$,
- we can show that Π_T is not trivial, i.e., it cannot be solved in 0 rounds.

Now if we could solve the original problem in T rounds, we could also solve its relaxation Π_0 in T rounds, and hence we could solve $\mathcal{R}(\Pi_0)$ in $T - 1$ rounds, and hence we could solve its relaxation Π_1 in $T - 1$ rounds, etc., and eventually we could solve Π_T in 0 rounds, but this was not possible. If we can do this for, e.g., $T = \Delta$, we would have the result that we want: iterated GHZ cannot be solved in $o(\Delta)$ rounds. In essence, our task is to **guess a good sequence of problems** such that all of the above holds. How to do that?

Comparison with bipartite maximal matching. At this point it is useful to compare the situation with the *bipartite maximal matching* problem [9, 10]. There is a very simple classical LOCAL algorithm that solves the problem in $O(\Delta)$ rounds, but for decades it was unknown whether this is optimal. The breakthrough result from 2019 [9] eventually established this, by constructing a suitable problem sequence $\Pi_0, \Pi_1, \dots, \Pi_T$, and showing that it indeed has all the right properties: Π_{i+1} is a relaxation of $\mathcal{R}(\Pi_i)$.

Bipartite maximal matching is conceptually a very simple problem; it can be specified in the appropriate round elimination formalism with only 3 labels, independently of the degree. Hence there is a very simple starting point Π_0 . It is possible to compute $\mathcal{R}(\Pi_0)$, $\mathcal{R}(\mathcal{R}(\Pi_0))$, etc., write down explicitly some of the first problems in this sequence for some small value of Δ , study their structure, make some educated guesses on possible relaxations, and eventually with a sufficient amount of trial and error, human intuition, and computational help one can make a reasonable guess of the problem sequence $\Pi_0, \Pi_1, \dots, \Pi_T$, generalize this to an arbitrary Δ , and verify that it indeed has the right properties.

Iterated GHZ seems to be fundamentally different. The problem description is relatively long and complicated, and when we increase Δ , the size of the problem description increases exponentially. Studying the sequence $\mathcal{R}(\Pi_0), \mathcal{R}(\mathcal{R}(\Pi_0)), \dots$ is not particularly enlightening to a human being, and even if there are educated guesses of suitable relaxations, they are very hard to verify, even with the help of computers.

In this work we made the following surprising discovery: even though the maximal matching problem and the iterated GHZ problem do not seem to have any relation between them, beyond the fact that both of them admit an $O(\Delta)$ -round classical algorithm, we can nevertheless **apply qualitatively similar relaxations in both maximal matching and iterated GHZ** to construct a suitable problem sequence! This technique is likely to generalize to a broad range of other problems.

Finding the right relaxations. Now we come to the mysterious part: when we look at the problems at the *right abstraction level*, we can prove a lower bound for both maximal matching and iterated GHZ using the “same” relaxations, in a sense, even though the problems, their alphabets, and their constraints do not seem to have anything in common.

We do not know why this technique works, and we do not know how far it generalizes. But we do know that it seems to be applicable to many previously-studied variants of the maximal matching problem, and it seems to be applicable to both iterated CHSH and iterated GHZ. The reader does

not need to believe in this technique or trust that it is correct—this was merely a heuristic rule that enabled us to guess the right problem sequence $\Pi_0, \Pi_1, \dots, \Pi_T$, and once we have the sequence, we can prove that it indeed satisfies all the right properties. But as the technique is apparently applicable to two very different problem families, it may find other surprising uses, and hence we outline it here.

The following description is primarily intended for a reader who is already experienced in the use of round elimination; we refer to [26, 49, 58] for additional background. But for all other readers this hopefully demonstrates that even though the problem sequence that we will present later in this paper may look at first arbitrary and appearing out of thin air, there is a relatively short procedure that can be used to mechanically generate it.

When we compute $\Pi' = \mathcal{R}(\Pi)$, each label in Π' represents a *nonempty subset* of labels in Π ; singleton sets represent *old labels* that already existed in Π , while all other sets are *new labels*. Now we can construct a directed graph G that represents the relative strength of labels in Π' from the perspective of passive constraints: we will have an arrow $a \rightarrow b$ in the graph if the passive constraints satisfy the property that any configuration remains valid if you replace any a with b (of course this might violate active constraints). The transitive reduction of graph G is called the *diagram* for Π' in this context. Equipped with these definitions, we then find any pair of labels a and b that satisfy all these properties:

- a is a new label,
- a has a nonzero indegree in the diagram,
- b is the only successor of a in the diagram,
- b has outdegree zero or b is a new label.

If such a pair exists, we merge the labels a and b to obtain a relaxation of Π' . We then update the diagram and repeat the procedure until no such pair of labels a and b exists.

It turns out that following exactly this procedure will be able to reproduce the results of [9, 29], but it will also let us find exactly the right relaxations that can be used to show that the iterated GHZ problem cannot be solved in $o(\Delta)$ rounds with classical LOCAL algorithms. As a final remark we note that odd and even steps of the round elimination sequence are fundamentally different and here we applied this rule in even steps.

1.4 Implications and discussion

Our work presents the first local problem (more precisely, a family of LCL problems) that admits a distributed quantum advantage in the LOCAL model of distributed computing. In this section we will discuss some further implications and connections with other areas.

Communication complexity. As we discussed earlier, there is much more prior work on the quantum advantage in the CONGEST model than in the LOCAL model. One reason for this is that the CONGEST model is similar in spirit to questions in communication complexity. In essence, if we have a result in communication complexity that shows that b qubits can be used to solve a task that $O(b)$ classical bits cannot solve, one can try to turn that into an analogous result that separates the classical CONGEST model from the quantum-CONGEST model. However, when we compare classical LOCAL and quantum-LOCAL, we cannot use any communication complexity arguments. In essence our goal is to show that *unbounded quantum communication is stronger than unbounded classical communication*. In essence, we show that being able to communicate b qubits is strictly stronger than being able to communicate 2^b or even 2^{2^b} classical bits!

Circuit complexity. The question of distributed quantum advantage for LCLs is also related to establishing an *efficiently-verifiable* quantum advantage in circuit complexity. The quantum advantage obtained in [53] was obtained by “lifting” to the distributed setting a quantum advantage for circuits by Bravyi, Gosset and König [31], which was inspired by an older work in quantum information [25]. Ref. [31] showed that a variant of the global graph problem from [53] can be solved by a constant-depth quantum circuit but cannot be solved by any small-depth classical circuit. This computational problem, however, could not be checked by a quantum constant-depth circuit, leading to the question whether a similar advantage can be proved for an efficiently-verifiable computational problem. Since this question is still wide open in quantum circuit complexity, such a simple lifting strategy cannot be used to tackle our question. Indeed, our separation uses an approach that does not seem to be applicable to circuit complexity.

Lower-bound proof techniques. Our work is the first proof that shows that round elimination is able to separate classical LOCAL and quantum-LOCAL. As a corollary, we now know that round elimination is not applicable in the quantum-LOCAL model: it is possible that there is a round elimination sequence $\Pi_0, \Pi_1, \dots, \Pi_T$, but the existence of this sequence has no implications on the round complexity in the quantum-LOCAL model. This also represents a new formal barrier for resolving major open questions related to distributed quantum advantage: it is not possible to prove quantum lower bounds for e.g. 3-coloring cycles or for finding sinkless orientations by constructing a round elimination sequence.

Networks of quantum games. Our problem, iterated GHZ, is based on the idea that we build an LCL problem that asks us to play a number of quantum games, so that the games are represented by black nodes and the players are represented by white nodes. Then inside each white node the output of one game is fed as input to another game. We call such a problem a *network of quantum games*. Such a problem can be always solved in $O(1)$ rounds in quantum-LOCAL, as we can create shared entangled qubits in the first round, and then all further steps can be done without any additional communication.

Our work shows that the classical complexity of such an LCL can be as high as $\Omega(\Delta)$. Now the obvious question is whether we can push this further: can we maybe engineer a network of quantum games such that its classical complexity is super-constant as a function of the number of nodes n , even if Δ is constant? For example, could we show a lower bound of $\Omega(n)$ or $\Omega(\log n)$ or maybe $\Omega(\log^* n)$ for such a problem?

It turns out that the answer is **no**. We generalize the idea so that instead of quantum games, we can have a network of any non-signaling games (regardless of whether they happen to have quantum strategies or not); we call such a problem a *network of non-signaling games*. We show that if Π is any LCL problem that is constructed from networks of non-signaling games (even if we connect the games with each other using arbitrarily complicated arithmetic circuits), it can be always solved in $O(1)$ rounds with classical LOCAL algorithms; here the constant depends on the specific LCL problem Π (and in particular, if we have a family of LCL problems parameterized by Δ , the constant may depend on Δ). However, the running time will be independent of n .

The key reason for this is that any non-signaling game is *completable*: for example, for a two-player game this means that for any Alice’s input x there is some Alice’s output a such that for any Bob’s input y there is still a valid Bob’s output b , and vice versa.

We can exploit completeness to solve any Π in a distributed manner as follows. Each white node starts to process its own arithmetic circuit in a topological order. As soon as it encounters a step that involves a game, it sends a message to the black neighbor responsible for that specific

game, together with its own input for that game. The black nodes keep track of the inputs they have seen so far, and they always pick *safe* outputs for those players. This way in one pair of rounds all white nodes can learn their own output for the game that appears first in their own circuit, and we can repeat this for each game in a sequential order—thanks to completeness, black nodes will be always able to assign valid outputs also for players that join the game late. This way the running time will be proportional to the size of the circuit held by a single white node, and hence for a fixed LCL (with a finite set of possible local circuits) it will be bounded by some constant.

The result may at first seem counterintuitive: A non-signaling game may require the players to *break symmetry* (e.g., in the CHSH game, if the inputs are both 1, the outputs must be different). Hence it would seem that we could put together non-signaling games in cycle so that the input does not break symmetry, but the output has to break symmetry. Then we could apply the classical result that shows that any symmetry-breaking LCL requires $\Omega(\log^* n)$ rounds [54]. However, the above algorithm demonstrates that *any network of non-signaling games admits a symmetric solution in those neighborhoods where the input (together with the port numbering) does not break symmetry*.

This represents a new barrier for proving distributed quantum advantage: the idea of networks of non-signaling games can only prove $f(\Delta)$ separations between classical LOCAL and quantum-LOCAL, not $f(n)$ separations.

1.5 Roadmap

We introduce some basic definitions in Section 2. In Section 3 we give the formal definition of the iterated GHZ problem in the right formalism that we need for round elimination, and we show that there is indeed a one-round quantum-LOCAL algorithm for solving this problem. In Section 4 we prove our main result: iterated GHZ requires $\Omega(\Delta)$ rounds in classical LOCAL. Finally, in Section 5 we formalize the idea of networks of non-signaling games, and show that any fixed LCL in this problem family can be solved in constant time in classical LOCAL.

2 Preliminaries

We write \mathbb{N}_0 for the set of non-negative integers, \mathbb{N}_+ for that of positive integers, and $[n] = \{1, \dots, n\}$ for the set of the first $n \in \mathbb{N}_0$ positive integers. For a relation $R \subseteq A \times B$ and $x \in A$, we write $R(x)$ for the set $\{y \in B \mid (x, y) \in R\}$.

For $m \in \mathbb{N}_0$ and a set Σ , we write Σ^m for the set of strings over elements of Σ that have length m . We will refer to such strings in two different ways, both as concatenations of m many elements of Σ (e.g., $x = x_1 \cdots x_m$ where $x_i \in \Sigma$) and as maps $[m] \rightarrow \Sigma$. These two formalisms may be used interchangeably without confusion. For a permutation $\sigma: [m] \rightarrow [m]$ and a string $x = x_1 \cdots x_m \in \Sigma^m$, we let $x^\sigma = x_{\sigma(1)} \cdots x_{\sigma(m)}$.

Let $G = (V, E)$ be a graph and $v \in V$ a node of G . Since E is a relation, recall we may use our notation for relations to write $E(v)$ as shorthand for the set $\{e \in E \mid e \text{ is incident to } v\}$. For any node v of a graph $G = (V, E)$, we denote by $\deg_G(v)$ its degree in G . If the underlying graph is clear from the context, we omit the subscript and write simply $\deg(v)$.

In this paper we are interested in locally checkable labeling (LCL) problems, first introduced by Naor and Stockmeyer in [57]. We now define this class of problems using a formalism that is slightly different from the original one.

LCLs in the black-white formalism. While there are different ways to define LCLs, in this paper we consider the so-called *LCLs in the black-white formalism*. Such a formalism, compared

to the definition given by Naor and Stockmeyer in [57], is more suitable for applying the round elimination technique.

In this formalism, problems are defined on bipartite graphs. We refer to nodes in the first set as *white* nodes and to nodes in the other as *black* nodes. The connection to a general graph G is that we see white nodes as nodes of G and black nodes as edges of G . In this sense, we are able to handle constraints for nodes and edges in a uniform manner.

Definition 2.1 (LCL problem in the black-white formalism). An *LCL problem in the black-white formalism* is a tuple $\Pi = (\Sigma_{\text{in}}, \Sigma_{\text{out}}, \mathcal{C}_w, \mathcal{C}_b)$ where:

- Σ_{in} and Σ_{out} are finite sets of *labels*.
- \mathcal{C}_w is the *white constraint*, formally a set of multisets where each element of a multiset is a pair of labels $(\ell_{\text{in}}, \ell_{\text{out}}) \in \Sigma_{\text{in}} \times \Sigma_{\text{out}}$.
- \mathcal{C}_b is the *black constraint*, formally a set of multisets where each element of a multiset is a pair of labels $(\ell_{\text{in}}, \ell_{\text{out}}) \in \Sigma_{\text{in}} \times \Sigma_{\text{out}}$.

Let $G = (V, E)$ be a bipartite graph where $V = W \cup B$ is a partition of V that implicitly “colors” the nodes white and black by placing them in W and B , respectively. In addition, let $i: E \rightarrow \Sigma_{\text{in}}$ be an *input*, that is, a labeling of edges such that $i(e) \in \Sigma_{\text{in}}$ for each $e \in E$. An assignment $o: E \rightarrow \Sigma_{\text{out}}$ is a *solution* to Π if it satisfies both the white and the black constraints, that is:

- For every white node $w \in W$, if e_1, \dots, e_{d_w} are the edges that are incident to w , then $\{(i(e_1), o(e_1)), \dots, (i(e_{d_w}), o(e_{d_w}))\} \in \mathcal{C}_w$.
- For every black node $b \in B$, if e_1, \dots, e_{d_b} are the edges that are incident to b , then we have $\{(i(e_1), o(e_1)), \dots, (i(e_{d_b}), o(e_{d_b}))\} \in \mathcal{C}_b$.

In general graphs, the black-white formalism captures a strict subset of LCLs as defined originally by [57]; however, in the LOCAL model of computation (whose definition follows later in this section) the two notions are equivalent in general trees and high-girth graphs, that are the classes of graphs needed to apply round-elimination [12].

We first introduce the port-numbering model, and on top of it we define the LOCAL model of computation [54] and the quantum-LOCAL model [43]. We stick to the black-white formalism to keep all definitions consistent.

The port-numbering model. A port-numbered network is a triple $N = (V, P, p)$ where V is the set of nodes, P is the set of *ports*, and $p: P \rightarrow P$ is a function specifying connections between ports. Each element $x \in P$ is a pair (v, i) where $v \in V$, $i \in \mathbb{N}_+$. The connection function p between ports is an involution, that is, $p(p(x)) = x$ for all $x \in P$. If $(v, i) \in P$, we say that (v, i) is port number i in node v . The degree of a node v in the network N is $\deg_N(v)$ is the number of ports in v , that is, $\deg_N(v) = |\{i \in \mathbb{N} : (v, i) \in P\}|$. We assume that port numbers are consecutive, i.e., the ports of any node $v \in V$ are $(v, 1), \dots, (v, \deg_N(v))$. Clearly, a port-numbered network identifies an *underlying graph* $G = (V, E)$ where, for any two nodes $u, v \in V$, $\{u, v\} \in E$ if and only if there exists ports $x_u, x_v \in P$ such that $p(x_u) = x_v$. Here, the degree of a node $\deg_N(v)$ corresponds to $\deg_G(v)$.

In the port-numbering model we are given a distributed system consisting of a port-numbered network of $|V| = n$ *processors* (or *nodes*) that operates in a sequence of synchronous rounds. In each round the processors may perform unbounded computations on their respective local state variables and subsequently exchange of messages of arbitrary size along the links given by the underlying input graph. Nodes identify their neighbors by using ports as defined before, where the port assignment

may be done adversarially. Barring their degree, all nodes are identical and operate according to the same local computation procedures. Initially all local state variables have the same value for all processors; the sole exception is a distinguished local variable $x(v)$ of each processor v that encodes input data.

Let Σ_{in} be a set of input labels. The input of a problem is defined in the form of a labeled, bipartite graph (G, i) where $G = (V = B \cup W, E)$, V is the set of processors (hence it is specified as part of the input), and $i: E \rightarrow \Sigma_{\text{in}}$ is an assignment of an input label $i(v) \in \Sigma_{\text{in}}$ to each edge $e \in E$. The nodes in V are properly 2-colored with either color *black* (elements of B) or *white* (elements of W). The output of the algorithm is given in the form of an assignment of local output labels $o: E \rightarrow \Sigma_{\text{out}}$, and the algorithm is assumed to terminate once all labels $o(v)$ are definitely fixed. These outputs are determined and assigned by the white nodes, which are also referred to as *active nodes*, whereas black nodes are called *passive nodes* (they do not output anything). We assume that nodes and their links are fault-free. The local computation procedures may be randomized by giving each processor access to its own set of random variables; in this case, we are in the *randomized* port-numbering model as opposed to the *deterministic* port-numbering model. If the algorithm is randomized, we also require that the failure probability while solving any problem is at most $1/n$, where n is the size of the input graph.

The running time of an algorithm is the number of synchronous rounds required by all nodes to produce output labels. If an algorithm running time is T , we also say that the algorithm has locality T . Notice that T can be a function of the size of the input graph.

We remark that the notion of an (LCL) problem is a graph problem, and does not depend on the specific model of computation we consider (hence, the problem definition cannot depend on, e.g., port numbers).

The LOCAL model. The LOCAL model was first introduced by [54], here adapted to the black-white formalism. It is just the port-numbering model augmented with an assignment of unique identifiers to nodes. Let $c \geq 1$ be a constant, and let Σ_{in} be a set of input labels. The nodes of the input graph $G = (V, E)$ are given as input also unique identifiers specified by an injective function $\text{id}: V \rightarrow [n^c]$. This assignment might be adversarial and is stored in the local state variable $x(v)$, and nodes can exploit these values during their local computation.

The local computation procedures may be randomized by giving each processor access to its own set of random variables; in this case, we are in the *randomized* LOCAL model as opposed to the *deterministic* LOCAL model. If the algorithm is randomized, we also require that the failure probability while solving any problem is at most $1/n$, where n is the size of the input graph.

The quantum-LOCAL model. The quantum-LOCAL of computing is similar to the deterministic LOCAL model above, but now with quantum computers and quantum communication links. More precisely, the quantum computers manipulate local states consisting of an unbounded number of qubits with arbitrary unitary transformations, the communication links are quantum communication channels (adjacent nodes can exchange any number of qubits), and the local outputs can be the result of any quantum measurement.

3 Iterated GHZ problem

We start by defining the *iterated GHZ* problem. While in Section 2 we provided the definition of LCLs in their full generality, in this section we consider a restricted setting, i.e., the case in which the constraints do not depend on the provided input, and the considered graph is biregular. Hence,

in this section (and in the following one), an LCL problem Π in the black-white formalism is simply a tuple $(\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi)$ where Σ_Π is a finite set of elements, called (output) *labels*, \mathcal{N}_Π is a collection of cardinality- Δ multisets of labels from Σ_Π , and \mathcal{E}_Π is a collection of cardinality-3 multisets of labels from Σ_Π . We call \mathcal{N}_Π the *white constraint* of Π and \mathcal{E}_Π the *black constraint* of Π . We will use the term *configuration* to refer to any multiset of labels of cardinality Δ , resp. 3; we will call a configuration of cardinality Δ , resp. 3, an *allowed configuration* if it is contained in the white constraint, resp. black constraint, of the respectively considered LCL problem. For simplicity, we will write a configuration $\{\ell_1, \dots, \ell_k\}$ (which, technically, is a multiset) in the form $\ell_1 \dots \ell_k$ (where $k \in \{3, \Delta\}$).

For degree Δ , the iterated GHZ problem is defined as an LCL $\mathcal{P}_\Delta = (\Sigma_{\mathcal{P}_\Delta}, \mathcal{N}_{\mathcal{P}_\Delta}, \mathcal{E}_{\mathcal{P}_\Delta})$ in the black-white formalism (in the aforementioned restricted setting). While the constraints of this problem do not depend on the provided input, we will consider lower and upper bounds for this problem in the setting in which black nodes B are provided with a labeling $c : B \rightarrow \{1, \dots, \Delta\}$ satisfying that, if two black nodes u and v are incident to the same white node, then $c(u) \neq c(v)$. In the following, this input will simply be called Δ -edge coloring (where the name comes from the fact that we could think of a $(\Delta, 3)$ -biregular graph as a hypergraph of degree Δ and rank 3, and we consider a coloring of the hyperedges).

Labels. The label set is defined as:

$$\begin{aligned} \Sigma_{\mathcal{P}_\Delta} = & \{(-y)_1 \mid y \in \{0, 1\}\} \\ & \cup \{(x y)_j \mid 2 \leq j \leq \Delta \text{ and } x, y \in \{0, 1\}\}. \end{aligned}$$

White constraint. For all vectors $(y_1, \dots, y_\Delta) \in \{0, 1\}^\Delta$, $\mathcal{N}_{\mathcal{P}_\Delta}$ contains the configuration $L_1 \dots L_\Delta$ satisfying the following properties.

- $L_1 = (-y_1)_1$.
- For all j such that $2 \leq j \leq \Delta$, $L_j = (y_{j-1} y_j)_j$.

Black constraint. The black constraint $\mathcal{E}_{\mathcal{P}_\Delta}$ contains the following configurations.

- $(-0)_1 \quad (-0)_1 \quad (-1)_1$.
- $(x_1 y_1)_j \quad (x_2 y_2)_j \quad (x_3 y_3)_j$, for all $2 \leq j \leq \Delta$ and $x_1, x_2, x_3, y_1, y_2, y_3$ satisfying:
 - $x_i \in \{0, 1\}$ and $y_i \in \{0, 1\}$ for all $1 \leq i \leq 3$;
 - if $x_1 + x_2 + x_3$ is even, then $y_1 \oplus y_2 \oplus y_3 = x_1 \vee x_2 \vee x_3$.

Some intuition. Informally, each label encodes an input and an output to a GHZ game. The white constraint enforces that the output of one game flows correctly to become the input of the next game. The black constraint enforces that, on a black node, the GHZ game is played correctly by the neighboring white nodes, except for the first black constraint, which is a special case and only requires to break symmetry between adjacent white nodes. This special case is required to keep the game from becoming trivial.

Considered setting. The problem \mathcal{P}_Δ is only defined on $(\Delta, 3)$ -biregular graphs. However, it is trivial to lift this definition to arbitrary graphs: white nodes of degree Δ and black nodes of degree 3 need to satisfy the constraints of \mathcal{P}_Δ ; all other nodes can produce an arbitrary output.

Intuitively, the lower bound that we will prove can be interpreted as follows: for any n and any small-enough Δ , there exists a balanced tree of n nodes where all white nodes that are “far enough” from the leaves have degree Δ , all black nodes that are “far enough” from the leaves have degree 3, and there exists a white node that is “far enough” from the leaves that needs to spend $\Omega(\Delta)$ rounds. Observe that restricting the possible input instances only makes our lower bound stronger (i.e., this problem is hard *already* on regular trees).

We will provide upper bounds for $(\Delta, 3)$ -biregular graphs. However, given that the definition of \mathcal{P}_Δ lifted to arbitrary graphs allows an arbitrary output for white nodes of degree different from Δ and black nodes of degree different from 3, it is trivial to adapt our algorithm to the case of graphs with nodes of arbitrary degree.

3.1 Quantum algorithm for iterated GHZ

We now prove that the iterated GHZ problem can be solved in $O(1)$ rounds in the quantum LOCAL model, no matter how large Δ is. More formally, we devote the rest of the section to proving the following theorem.

Theorem 3.1. *For any integer $\Delta > 0$, the iterated GHZ problem can be solved in $O(1)$ rounds in the quantum LOCAL model, if a Δ -edge coloring is provided.*

Quantum algorithm for iterated GHZ problem. Each black node receives a color $c \in [\Delta]$ as an input; recall that this is what we call a Δ -edge coloring. Each black node prepares three qubits q_1, q_2, q_3 in the GHZ state, that is

$$|0\rangle_{q_1} |0\rangle_{q_2} |0\rangle_{q_3} \rightarrow \frac{1}{\sqrt{2}} \left(|0\rangle_{q_1} |0\rangle_{q_2} |0\rangle_{q_3} + |1\rangle_{q_1} |1\rangle_{q_2} |1\rangle_{q_3} \right).$$

It then sends message (c, i, q_i) to port i . Note that this is a message containing a qubit belonging to a three-party entangled state. Then the black node stops.

The white nodes wait for the messages from the black nodes. Let (c'_j, i'_j, q'_j) be the message a white node receives from port j . Recall that c_j is just the color of the Δ -edge coloring, and hence all colors $c'_1, c'_2, \dots, c'_\Delta$ form a permutation. Let us reindex the messages according to this permutation, i.e. (k, i_k, q_k) is the message received from the black neighbor with color k .

The white node first breaks the symmetry based on i_1 . If $i_1 = 2$, it sets $g_1 = 1$, otherwise $g_1 = 0$. Now the white node plays sequentially the GHZ games locally with no further communication by exploiting the GHZ states distributed by the black nodes. For each game $k \in \{2, 3, \dots, \Delta\}$, in order starting with $k = 2$, the white node measures the qubit q_k , with input g_{k-1} , according to the well-known quantum strategy [30, 56] for winning the GHZ game perfectly. This gives the measurement outcome g_k , which is then used for the $(k + 1)^{\text{th}}$ GHZ game, until all Δ games have been played.

Finally, the white node produces an output for each of its incident edges. For edge neighboring color 1, the node outputs $(-g_1)_1$. For edge neighboring color $k \in \{2, 3, \dots, \Delta\}$, the node outputs $(g_{k-1} g_k)_k$. Then the node stops.

Now the only thing left is to prove that this algorithm produces a correct output.

Proof. It is easy to see that the white constraint is satisfied: there is exactly one edge with each subscript, and the labels g_i propagate correctly between the labels.

The black constraint is also easy to check: Nodes with input color 1 satisfy the first constraint, as exactly one neighbor received message $(1, 2, q_2)$, and only that node outputs $(-1)_1$; the other two neighbors output $(-0)_1$. Nodes with input color $k \in \{2, \dots, \Delta\}$ satisfy the latter constraint. In particular, all of them have subscript k and describe the input and output to a GHZ game. Moreover, as the algorithm uses the quantum GHZ strategy to produce the outputs from the inputs, they satisfy the GHZ constraint. \square

4 Classical lower bound for iterated GHZ

In this section, we prove that the iterated GHZ problem requires $\Omega(\Delta)$ rounds in the LOCAL model. More formally, we prove the following.

Theorem 4.1. *Let $\Delta \geq 3$ be an integer. The iterated GHZ problem requires $\Omega(\min\{\Delta, \log_\Delta n\})$ rounds in the deterministic LOCAL model and $\Omega(\min\{\Delta, \log_\Delta \log n\})$ rounds in the randomized LOCAL model, even if a Δ -edge coloring is provided.*

The lower bound is based on the round elimination technique [26]. In Section 4.1, we first provide definitions for a variety of helpful concepts and objects and introduce round elimination. For simplicity, we will introduce all the definitions in the context of the LCL problems relevant in this paper; in particular, we will assume that black nodes have degree 3. After giving an overview of the structure of the proof in Section 4.2, we present the whole proof in Sections 4.3 to 4.10.

4.1 Preliminaries

Definitions. We will use so-called *condensed configurations* to represent certain sets of configurations efficiently. Formally, a condensed configuration is a configuration of sets of labels from Σ_Π . For a condensed configuration $C = S_1 \dots S_k$, by *configurations represented by C* we denote the set of all configurations $\ell_1 \dots \ell_k$ such that $\ell_j \in S_j$ for each $1 \leq j \leq k$. To distinguish a condensed configuration (which represents a set of configurations of labels from Σ) from a configuration of sets, we will use the term *disjunction* instead of “set” and write the disjunction consisting of some labels $\ell^{(1)}, \dots, \ell^{(z)}$ as $[\ell^{(1)}, \dots, \ell^{(z)}]$. However, when convenient, we may consider a disjunction of labels as the set of the same labels.

Definition 4.2 (Picking a configuration). We say that we can *pick* a configuration ℓ_1, \dots, ℓ_k from a configuration S_1, \dots, S_k of sets of labels, if there exists a permutation $\sigma: \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ such that, for each $1 \leq i \leq k$, it holds that $\ell_i \in S_{\sigma(i)}$. Moreover, we say that we can pick a configuration C from a set \mathcal{C} of configurations of sets of labels if there exists a configuration $C' \in \mathcal{C}$ from which we can pick C .

By configurations represented by a set of condensed configurations, we denote the set of configurations that can be picked from the condensed configurations.

Round elimination. On a high level, the round elimination framework introduced in [26] provides a mechanism to prove a lower bound for a given LCL problem Π by generating a sequence of problems starting with Π such that every problem has a complexity of exactly one round less than the previous problem (under some mild assumptions). The underlying intuition for obtaining a lower bound with this mechanism is that if one can show that the problem obtained after k steps in the sequence cannot be solved in 0 rounds, then this yields a lower bound of $k - 1$. We remark that this is a highly simplified description of the overall approach with a lot of details omitted; in the following, we introduce round elimination more formally.

The round elimination operator $\mathcal{R}()$. For a given LCL problem $\Pi = (\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi)$, the problem $\mathcal{R}(\Pi) = (\Sigma_{\mathcal{R}(\Pi)}, \mathcal{N}_{\mathcal{R}(\Pi)}, \mathcal{E}_{\mathcal{R}(\Pi)})$ is defined as follows.

- $\mathcal{E}_{\mathcal{R}(\Pi)}$ is the set of all configurations $S_1 S_2 S_3$ such that
 - for all $1 \leq j \leq 3$, S_j is a nonempty subset of Σ_Π ,
 - for each tuple $(\ell_1, \ell_2, \ell_3) \in S_1 \times S_2 \times S_3$, we have that $\ell_1 \ell_2 \ell_3 \in \mathcal{E}_\Pi$, and
 - $S_1 S_2 S_3$ is *maximal*, i.e., there exists no configuration $S'_1 S'_2 S'_3$ of nonempty subsets of Σ_Π such that
 - * for all $1 \leq j \leq 3$, $S_j \subseteq S'_j$,
 - * there exists some $1 \leq j \leq 3$ such that $S_j \subsetneq S'_j$, and
 - * for each tuple $(\ell'_1, \ell'_2, \ell'_3) \in S'_1 \times S'_2 \times S'_3$, we have that $\ell'_1 \ell'_2 \ell'_3 \in \mathcal{E}_\Pi$.
- $\Sigma_{\mathcal{R}(\Pi)}$ is the set of all nonempty subsets of Σ that appear in at least one configuration contained in $\mathcal{E}_{\mathcal{R}(\Pi)}$.
- $\mathcal{N}_{\mathcal{R}(\Pi)}$ is the set of all configurations $S_1 \dots S_\Delta$ of sets contained in $\Sigma_{\mathcal{R}(\Pi)}$ such that there exists some tuple $(\ell_1, \dots, \ell_\Delta) \in S_1 \times \dots \times S_\Delta$ satisfying $\ell_1 \dots \ell_\Delta \in \mathcal{N}_\Pi$.

The round elimination operator $\overline{\mathcal{R}}()$. For a given LCL problem $\Pi = (\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi)$, the problem $\overline{\mathcal{R}}(\Pi) = (\Sigma_{\overline{\mathcal{R}}(\Pi)}, \mathcal{N}_{\overline{\mathcal{R}}(\Pi)}, \mathcal{E}_{\overline{\mathcal{R}}(\Pi)})$ is defined as follows.

- $\mathcal{N}_{\overline{\mathcal{R}}(\Pi)}$ is the set of all configurations $S_1 \dots S_\Delta$ such that
 - for all $1 \leq j \leq \Delta$, S_j is a nonempty subset of Σ_Π ,
 - for each tuple $(\ell_1, \dots, \ell_\Delta) \in S_1 \times \dots \times S_\Delta$, we have that $\ell_1 \dots \ell_\Delta \in \mathcal{N}_\Pi$, and
 - $S_1 \dots S_\Delta$ is *maximal*, i.e., there exists no configuration $S''_1 \dots S''_\Delta$ of nonempty subsets of Σ_Π such that
 - * for all $1 \leq j \leq \Delta$, $S_j \subseteq S''_j$,
 - * there exists some $1 \leq j \leq \Delta$ such that $S_j \subsetneq S''_j$, and
 - * for each tuple $(\ell''_1, \dots, \ell''_\Delta) \in S''_1 \times \dots \times S''_\Delta$, we have that $\ell''_1 \dots \ell''_\Delta \in \mathcal{N}_\Pi$.
- $\Sigma_{\overline{\mathcal{R}}(\Pi)}$ is the set of all nonempty subsets of Σ_Π that appear in at least one configuration contained in $\mathcal{N}_{\overline{\mathcal{R}}(\Pi)}$.
- $\mathcal{E}_{\overline{\mathcal{R}}(\Pi)}$ is the set of all configurations $S_1 S_2 S_3$ of sets contained in $\Sigma_{\overline{\mathcal{R}}(\Pi)}$ such that there exists some tuple $(\ell_1, \ell_2, \ell_3) \in S_1 \times S_2 \times S_3$ satisfying $\ell_1 \ell_2 \ell_3 \in \mathcal{E}_\Pi$.

A simple way to compute $\mathcal{N}_{\mathcal{R}(\Pi)}$ and $\mathcal{E}_{\overline{\mathcal{R}}(\Pi)}$. It has been observed in previous works that use the round elimination technique (see, e.g, [12]), that $\mathcal{N}_{\mathcal{R}(\Pi)}$ can be easily computed as follows. Start from all the configurations in \mathcal{N}_Π , and for each configuration $C \in \mathcal{N}_\Pi$, add to $\mathcal{N}_{\mathcal{R}(\Pi)}$ all the configurations that can be picked from the condensed configuration obtained by replacing each label ℓ in C with the disjunction of all label sets in $\Sigma_{\mathcal{R}(\Pi)}$ that contain ℓ .

Similarly, $\mathcal{E}_{\overline{\mathcal{R}}(\Pi)}$ can be easily computed as follows. Start from all the configurations in \mathcal{E}_Π , and for each configuration $C \in \mathcal{E}_\Pi$, add to $\mathcal{E}_{\overline{\mathcal{R}}(\Pi)}$ all the configurations that can be picked from the condensed configuration obtained by replacing each label ℓ in C with the disjunction of all label sets in $\Sigma_{\overline{\mathcal{R}}(\Pi)}$ that contain ℓ .

Throughout this section, in order to compute $\mathcal{N}_{\mathcal{R}(\Pi)}$ and $\mathcal{E}_{\overline{\mathcal{R}}(\Pi)}$, we will implicitly use the described procedures.

Further notation and terminology. For a set \mathcal{L} of configurations of labels and a configuration $S = S_1 \dots S_k$ of sets of labels, we say that S *satisfies the universal quantifier* if, for each choice $(\ell_1, \dots, \ell_k) \in S_1 \times \dots \times S_k$, we have $\ell_j \in S_j$ for each $1 \leq j \leq k$. A configuration $S_1 \dots S_k$ satisfying the universal quantifier is called *maximal* if there is no configuration $S'_1 \dots S'_k$ satisfying the universal quantifier for which $S_j \subseteq S'_j$ for all $1 \leq j \leq k$ and $S_j \subsetneq S'_j$ for at least one $1 \leq j \leq k$. In particular, $\mathcal{N}_{\overline{\mathcal{R}}(\Pi)}$, resp. $\mathcal{E}_{\mathcal{R}(\Pi)}$, is precisely the set of all maximal configurations satisfying the universal quantifier (for $\mathcal{L} = \mathcal{N}_{\Pi}$, resp. $\mathcal{L} = \mathcal{E}_{\Pi}$). Utilizing the fact that a disjunction is nothing else than a set, we will call a set \mathcal{S} of condensed configurations *maximal-complete* if the set of configurations of sets obtained from \mathcal{S} by interpreting each disjunction as a set contains all the set of all maximal configurations satisfying the universal quantifier (w.r.t. the set of label configurations represented by \mathcal{S}).

Let $(\Sigma_{\Pi}, \mathcal{N}_{\Pi}, \mathcal{E}_{\Pi})$ be an LCL problem, and ℓ, ℓ' two labels from Σ_{Π} . We say that ℓ' *is at least as strong as ℓ w.r.t. \mathcal{E}_{Π}* (resp. *w.r.t. \mathcal{N}_{Π}*) if the following holds: for any configuration $\ell \ell_2 \ell_3 \in \mathcal{E}_{\Pi}$ (resp. for any configuration $\ell \ell_2 \dots \ell_{\Delta} \in \mathcal{N}_{\Pi}$), we have $\ell' \ell_2 \ell_3 \in \mathcal{E}_{\Pi}$ (resp. $\ell' \ell_2 \dots \ell_{\Delta} \in \mathcal{N}_{\Pi}$). We may omit the dependency on the constraint when clear from the context; in particular, for the problems $\Pi_{i,\Delta}$ we define later and all problems obtained from applying $\overline{\mathcal{R}}$, the relative strength of labels will be considered w.r.t. the respective black constraint, while for all problems obtained from applying \mathcal{R} it will be considered w.r.t. the respective white constraint. If ℓ' is at least as strong as ℓ but ℓ is not at least as strong as ℓ' , we say that ℓ' *is stronger than ℓ* . We write $\ell \leq \ell'$ to represent that ℓ' is at least as strong as ℓ ; we write $\ell < \ell'$ to represent that ℓ' is stronger than ℓ .

We will make use of so-called *diagrams* to represent the strength relations between labels. Formally, the diagram of a problem $(\Sigma_{\Pi}, \mathcal{N}_{\Pi}, \mathcal{E}_{\Pi})$ (w.r.t. \mathcal{N}_{Π} , resp. \mathcal{E}_{Π}) is the directed graph where the set of nodes is Σ_{Π} and there is a directed edge (ℓ, ℓ') between two labels $\ell, \ell' \in \Sigma_{\Pi}$ if

- $\ell < \ell'$ and
- there is no label $\ell'' \in \Sigma_{\Pi}$ satisfying $\ell < \ell'' < \ell'$

(where the strength relations are w.r.t. \mathcal{N}_{Π} , resp. \mathcal{E}_{Π}). We remark that the definition of strength implies that any diagram is acyclic.

We call two labels ℓ, ℓ' *comparable* if $\ell \leq \ell'$ or $\ell' \leq \ell$, and *incomparable* otherwise. From the definition of a diagram, it follows that two labels ℓ, ℓ' are comparable if and only if there is a directed path between ℓ and ℓ' (in one of the two possible directions) in the diagram. Furthermore, for two configurations $L = \ell_1 \dots \ell_k$ and $L' = \ell'_1 \dots \ell'_k$ of labels, we say that L' *dominates* L if $\ell_j \leq \ell'_j$ for all $1 \leq j \leq k$. Similarly, for configurations $C = S_1 \dots S_k$ and $C' = S'_1 \dots S'_k$ of sets of labels, we say that C' *dominates* C if $S_j \subseteq S'_j$ for all $1 \leq j \leq k$. (Note that as configurations are multisets, a configuration already dominates another if there is a permutation $\rho: \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ such that $\ell_j \leq \ell'_{\rho(j)}$ for all $1 \leq j \leq k$, resp. $S_j \subseteq S'_{\rho(j)}$ for all $1 \leq j \leq k$.) Moreover, for two sets \mathcal{C} and \mathcal{C}' of configurations, we say that \mathcal{C}' *dominates* \mathcal{C} if, for all $C \in \mathcal{C}$, there exists a configuration $C' \in \mathcal{C}'$ that dominates C .

For labels ℓ_1, \dots, ℓ_k , the set of all labels that are at least as strong as at least one of the ℓ_j is denoted by $\langle \ell_1, \dots, \ell_k \rangle$. In other words, $\langle \ell_1, \dots, \ell_k \rangle$ is the set of all labels that are identical to one of the ℓ_j or a (not necessarily direct) successor of one of the ℓ_j in the respective diagram. We say that $\langle \ell_1, \dots, \ell_k \rangle$ is *generated* by ℓ_1, \dots, ℓ_k . As before, we will omit the dependency on the constraint when clear from the context, and in particular use the conventions mentioned above. Moreover, for a set $L = \{\ell_1, \dots, \ell_k\}$ by writing $\langle \ell \mid \ell \in L \rangle$ we denote $\langle \ell_1, \dots, \ell_k \rangle$. Observe that $\langle L \rangle$ is different from $\langle \ell \mid \ell \in L \rangle$.

An alternative method to compute $\mathcal{R}(\Pi)$ and $\overline{\mathcal{R}}(\Pi)$. In a recent work [14, Section 4], a new method was developed for computing, for a given problem $\Pi = (\Sigma_{\Pi}, \mathcal{N}_{\Pi}, \mathcal{E}_{\Pi})$, the problems

$\mathcal{R}(\Pi)$ and $\overline{\mathcal{R}}(\Pi)$ in a different way. On a high level, this new method is based on a very simple operation that if applied iteratively will provide the same result as the more complex application of the universal quantifier in the definitions of $\mathcal{N}_{\mathcal{R}(\Pi)}$ and $\mathcal{E}_{\overline{\mathcal{R}}(\Pi)}$ above. The simplicity of said operation makes it considerably easier to apply round elimination and prove the correctness of claims such as that a certain problem is indeed what is obtained by applying \mathcal{R} or $\overline{\mathcal{R}}$. In the following, we describe this new method formally. We start by defining the notion of a combination of two configurations.

Let $C = S_1 \dots S_k$ and $C' = S'_1 \dots S'_k$ be two configurations of sets of labels. Let $1 \leq u \leq k$, and let $\sigma: \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ be a bijective function, i.e., a permutation. The *combination of C and C' w.r.t. u and σ* is defined as the configuration $C'' = S''_1 \dots S''_k$ satisfying

- $S''_j = S_j \cap S'_{\sigma(j)}$ for each $j \in \{1, \dots, k\} \setminus \{u\}$, and
- $S''_u = S_j \cup S'_{\sigma(j)}$.

Now consider the following process. Start with \mathcal{E}_{Π} , given as a collection of condensed configurations, where, as usual, disjunctions are considered as sets. (If \mathcal{E}_{Π} is given as a collection of uncondensed configurations, it is straightforward to transform it into a list of condensed configurations by interpreting each configuration of labels as the configuration of singleton disjunctions/sets containing those labels.) Now arbitrarily choose two configurations C and C' from the collection, combine them w.r.t. some arbitrarily chosen u and σ , and add the obtained configuration to the maintained collection of configurations of sets (unless it is already present). Repeat this, for all possible choices of C , C' , u , and σ (including all newly obtained combined configurations as choices for C and C') until no new configuration (i.e., no configuration not already obtained in the collection) is produced anymore. Then, remove from the obtained collection all *non-maximal configurations*, i.e., all configurations $S_1 \dots S_k$ for which there exists a configuration $S'_1 \dots S'_k$ satisfying $S_j \subseteq S'_j$ for all $1 \leq j \leq k$ and $S_j \subsetneq S'_j$ for at least one $1 \leq j \leq k$.

Now [14, Theorem 4.1] states that the obtained collection of configurations is identical to $\mathcal{N}_{\mathcal{R}(\Pi)}$, and that the same procedure, when started from \mathcal{E}_{Π} instead of \mathcal{N}_{Π} , yields $\mathcal{E}_{\overline{\mathcal{R}}(\Pi)}$. In other words, the described simple iterative process can be used to replace the traditional and more complex computation of $\mathcal{N}_{\mathcal{R}(\Pi)}$, resp. $\mathcal{E}_{\overline{\mathcal{R}}(\Pi)}$, when computing $\mathcal{R}(\Pi)$, resp. $\overline{\mathcal{R}}(\Pi)$ (while Σ_{Π} and $\mathcal{E}_{\mathcal{R}(\Pi)}$, resp. $\mathcal{N}_{\overline{\mathcal{R}}(\Pi)}$, are computed in the same way in both the traditional and the new method). We obtain the following observation (which follows directly from [14]).

Observation 4.3. *Let \mathcal{C} be a collection of configurations of sets. Assume that \mathcal{C} satisfies the following: for all pairs of configurations $C_1, C_2 \in \mathcal{C}$, all permutations σ and all indices u , the combination of C_1 and C_2 w.r.t. u and σ is dominated by some configuration in \mathcal{C} . Then, \mathcal{C} is maximal-complete.*

We call two sets S, S' of labels incomparable if $S \not\subseteq S'$ and $S' \not\subseteq S$, and comparable otherwise. Based on the notion of (in)comparable sets, the following observation (which is analogous to [14, Observation 5.4]) essentially shows that many choices for C , C' , u , and σ do not need to be considered in the described process as they cannot produce configurations that will be contained in the final collection.

Observation 4.4. *Let $C = S_1 \dots S_k$ and $C' = S'_1 \dots S'_k$ be two configurations of sets of labels, $1 \leq u \leq k$ an index, and $\sigma: \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ a permutation. Let $C'' = S''_1 \dots S''_k$ be the configuration obtained by combining C and C' w.r.t. u and σ . Then, if S_u and S'_u are comparable, C'' is dominated by C or C' .*

Proof. If S_u and S'_u are comparable, then $S_u \subseteq S'_u$ or $S'_u \subseteq S_u$. Now the definition of C'' ensures that, in the former case, C'' is dominated by C' and, in the latter case, C'' is dominated by C . \square

From the definition of the above process and Observation 4.4, it follows that the final collection obtained when the process terminates will not change if we modify the process so that combinations where S_u and S'_u are comparable are not considered. Note that it is possible that some configurations that are produced in the unmodified process along the way are not produced in the modified process; however, all such configurations are configurations that will not be maximal, i.e., that are removed in the last step of the unmodified process. In particular, all configurations that are maximal in the collection produced by the unmodified process will also be produced (and will be maximal) in the modified process.

We now define what it means for a problem to be the relaxation of another problem. Informally, for a problem Π' to be a relaxation of $\mathcal{R}(\Pi)$ it must hold that there exists a 0-round algorithm operating on black nodes that, given a solution for $\mathcal{R}(\Pi)$, it is able to solve Π' . Similarly, for a problem Π' to be a relaxation of $\overline{\mathcal{R}}(\Pi)$ it must hold that there exists a 0-round algorithm operating on white nodes that, given a solution for $\overline{\mathcal{R}}(\Pi)$, it is able to solve Π' . In this paper, we consider a stricter (i.e., less permissive) form of relaxation, defined as follows.

Definition 4.5 (Relaxation of a problem). Let $\Pi = (\Sigma, \mathcal{N}, \mathcal{E})$. A problem $\Pi' = (\Sigma', \mathcal{N}', \mathcal{E}')$ is a relaxation of $\mathcal{R}(\Pi)$ if the following holds:

- For all configurations C in the black constraint of $\mathcal{R}(\Pi)$, there exists a configuration $C' \in \mathcal{E}'$ that dominates C ;
- Σ' is the collection of sets that appear in at least one configuration contained in \mathcal{E}' ;
- \mathcal{N}' is the set of all configurations $S_1 \dots S_\Delta$ of sets contained in Σ' such that there exists some tuple $(\ell_1, \dots, \ell_\Delta) \in S_1 \times \dots \times S_\Delta$ satisfying $\ell_1 \dots \ell_\Delta \in \mathcal{N}$.

Similarly, a problem $\Pi' = (\Sigma', \mathcal{N}', \mathcal{E}')$ is a relaxation of $\overline{\mathcal{R}}(\Pi)$ if the following holds:

- For all configurations C in the white constraint of $\overline{\mathcal{R}}(\Pi)$, there exists a configuration $C' \in \mathcal{N}'$ that dominates C ;
- Σ' is the collection of sets that appear in at least one configuration contained in \mathcal{N}' ;
- \mathcal{E}' is the set of all configurations $S_1 S_2 S_3$ of sets contained in Σ' such that there exists some tuple $(\ell_1, \ell_2, \ell_3) \in S_1 \times S_2 \times S_3$ satisfying $\ell_1 \ell_2 \ell_3 \in \mathcal{E}$.

Definition 4.6 (Right-closed). Let $\Pi = (\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi)$ be an LCL problem. We say that a set $S = \{\ell_1, \dots, \ell_k\} \subseteq \Sigma_\Pi$ is *right-closed* if and only if, for each label $\ell_i \in S$, it holds that all successors of ℓ_i in the diagram (taken w.r.t. the constraint on which the universal quantifier is applied to) are also contained in S . More precisely, a set $S = \{\ell_1, \dots, \ell_k\} \subseteq \Sigma_\Pi$ is called right-closed if $S = \langle \ell_1, \dots, \ell_k \rangle$.

We observe the following statement, which was shown in [17].

Observation 4.7 ([17]). *Let $\Pi = (\Sigma_\Pi, \mathcal{N}_\Pi, \mathcal{E}_\Pi)$ be an LCL problem. Consider an arbitrary collection of labels $L_1, \dots, L_k \in \Sigma_\Pi$. If $\{L_1, \dots, L_k\} \in \Sigma_{\mathcal{R}(\Pi)}$, then the set $\{L_1, \dots, L_k\}$ is right-closed w.r.t. \mathcal{E}_Π . If $\{L_1, \dots, L_k\} \in \Sigma_{\overline{\mathcal{R}}(\Pi)}$, then the set $\{L_1, \dots, L_k\}$ is right-closed w.r.t. \mathcal{N}_Π .*

Definition 4.8 (A configuration generated by a condensed one). Let $C = S_1, \dots, S_k$ be a configuration of sets. We say that the configuration $C' = S'_1, \dots, S'_k$ is *generated by the condensed configuration* C if there exists a permutation $\sigma: \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ such that, for every $1 \leq i \leq k$, it holds that $S'_i = \langle c \mid c \in S_{\sigma(i)} \rangle$.

We are now ready to state the theorem that relates round elimination with the LOCAL model of distributed computing.

Theorem 4.9 ([12], 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 $f(\Delta)$.*

Also, assume that Π_t has at most $f(\Delta)$ labels and is not 0-round solvable in the deterministic port numbering model, even when given a Δ -edge coloring.¹ Then, Π_0 requires $\Omega(\min\{t, \log_\Delta n - \log_\Delta \log f(\Delta)\})$ rounds in the deterministic LOCAL model and $\Omega(\min\{t, \log_\Delta \log n - \log_\Delta \log f(\Delta)\})$ rounds in the randomized LOCAL model, even when given a Δ -edge coloring.

4.2 Roadmap

Informally, in order to prove a lower bound for \mathcal{P}_Δ , we introduce a family of problems $\Pi_{i,\Delta}$ (where $0 \leq i \leq \Delta - 2$) with the following properties.

- \mathcal{P}_Δ is at least as hard as $\Pi_{0,\Delta}$.
- By applying round elimination to $\Pi_{i,\Delta}$, we obtain a problem that is at least as hard as $\Pi_{i+1,\Delta}$, for all $0 \leq i \leq \Delta - 3$.
- $\Pi_{\Delta-2,\Delta}$ cannot be solved in 0 rounds given a Δ -edge coloring.

Our result will follow by applying Theorem 4.9 to such a sequence of problems.

More in detail, in Section 4.3 we define the problems $\Pi_{i,\Delta}$, and in Section 4.4 we prove the strength relation between their labels. Then, in Section 4.5 we define a problem $\Pi'_{i,\Delta}$ and show that it is a relaxation of $\mathcal{R}(\Pi_{i,\Delta})$. In Section 4.6 we prove the strength relation between the labels of $\Pi'_{i,\Delta}$. In Section 4.7 we prove that $\overline{\mathcal{R}}(\Pi'_{i,\Delta})$ can be relaxed to $\Pi_{i+1,\Delta}$. Then, in Section 4.8 we show that $\Pi_{\Delta-2,\Delta}$ cannot be solved in 0 rounds given a Δ -edge coloring. In Section 4.9, we show that \mathcal{P}_Δ is at least as hard as $\Pi_{0,\Delta}$. Finally, in Section 4.10, we combine the results of all previous sections to prove Theorem 4.1.

4.3 Problem definition

For each $0 \leq i \leq \Delta - 2$, we define the problems $\Pi_{i,\Delta}$ in the black-white formalism. We remark that while the problem $\Pi_{i,\Delta}$ is defined also for $i = \Delta - 2$, from Section 4.4 onwards we will assume $i \leq \Delta - 3$ for technical reasons.

Labels. The label set of $\Pi_{i,\Delta}$ is defined as

$$\begin{aligned} \Sigma_{i,\Delta} = & \{ (!!)_j \mid 1 \leq j \leq \Delta - i - 1 \} \\ & \cup \{ (!)_j \mid \Delta - i + 1 \leq j \leq \Delta \} \\ & \cup \{ (--)_j \mid 1 \leq j \leq \Delta - i \} \\ & \cup \{ (-y)_j \mid 1 \leq j \leq \Delta - i \text{ and } y \in \{0, 1\} \} \\ & \cup \{ (x-)_j \mid 2 \leq j \leq \Delta - i - 1 \text{ and } x \in \{0, 1\} \} \\ & \cup \{ (xy)_j \mid 2 \leq j \leq \Delta - i \text{ and } x, y \in \{0, 1\} \} \\ & \cup \{ (?)_j \mid \Delta - i + 1 \leq j \leq \Delta \}. \end{aligned}$$

¹In [12], this theorem is stated in its full generality, and it talks about restricted assignments of port numbers. It is easy to see that a Δ -edge coloring is indeed a restricted assignment of port numbers.

We call the subscript of a label its *color*. For example, label $(--)_j$ has color j . When considering some fixed i , we say that a color j is *gone* if $j > \Delta - i$, otherwise it is *present*. Moreover, color $\Delta - i$ is called *special*. Throughout the section, when writing a configuration $L_1 \dots L_\Delta$ we may implicitly assume that label L_i is of color i .

White constraint. We now define the white constraint $\mathcal{N}_{i,\Delta}$ of problem $\Pi_{i,\Delta}$. For an example, see the paragraph after the definition. $\mathcal{N}_{i,\Delta}$ contains node configurations of three different kinds.

First, for all integers a, b satisfying $1 \leq a < b \leq \Delta - i + 1$ and $b \neq \Delta - i$, and for all vectors $(y_a, \dots, y_{b-1}) \in \{0, 1\}^{b-a}$, $\mathcal{N}_{i,\Delta}$ contains the configuration $L_1 \dots L_\Delta$ satisfying the following properties.

- For all j such that $1 \leq j < a$, $L_j = (--)_j$,
- $L_a = (-y_a)_a$,
- for all j such that $a < j < b$, $L_j = (y_{j-1} y_j)_j$,
- if $b \leq \Delta - i$, $L_b = (y_{b-1} -)_b$,
- for all j such that $b < j \leq \Delta - i$, $L_j = (--)_j$, and
- for all j such that $\Delta - i + 1 \leq j \leq \Delta$, $L_j = (?)_j$.

Second, for all integers a satisfying $1 \leq a \leq \Delta - i - 1$, $\mathcal{N}_{i,\Delta}$ contains the configuration $L_1 \dots L_\Delta$ satisfying the following properties.

- $L_a = (!!)_a$,
- for all j such that $1 \leq j \leq \Delta - i$ and $j \neq a$, $L_j = (--)_j$, and
- for all j such that $\Delta - i + 1 \leq j \leq \Delta$, $L_j = (?)_j$.

Third, for all integers a satisfying $\Delta - i + 1 \leq a \leq \Delta$, $\mathcal{N}_{i,\Delta}$ contains the configuration $L_1 \dots L_\Delta$ satisfying the following properties.

- $L_a = (!)_a$,
- for all j such that $1 \leq j \leq \Delta - i$, $L_j = (--)_j$, and
- for all j such that $\Delta - i + 1 \leq j \leq \Delta$ and $j \neq a$, $L_j = (?)_j$.

Example for the white constraint. To make the above definition more accessible, we now informally show how to obtain the white constraint in a concrete way, illustrated with an example for specific values of Δ and i . Let $\Delta = 7$ and $i = 2$. We obtain $\mathcal{N}_{2,7}$ by setting $\mathcal{N}_{2,7} = \emptyset$ and then performing the following procedure for each bit string B of $\Delta - i$ bits (which we illustrate for the case $B = 01100$).

Add to $\mathcal{N}_{2,7}$ the configuration C of the first kind described above where we set $a := 1$, $b := \Delta - i + 1$, and $(y_a, \dots, y_{b-1}) := B$, i.e., we add

$$C = (-0)_1 (01)_2 (11)_3 (10)_4 (00)_5 (?)_6 (?)_7.$$

Starting from C , we now create additional configurations by *canceling* some bits, which we will then all add to $\mathcal{N}_{2,7}$ (if not already present). We can cancel a prefix of the bit string B , a suffix, or both at the same time. However, at least one bit must remain. Moreover, if we cancel a suffix, the canceled suffix cannot be only the last bit. Each time we obtain a new bit string, we transform it into a configuration analogously to how we obtained C from B , and add it to $\mathcal{N}_{2,7}$. In this way, we add the following configurations.

- Prefix canceling:

$$\begin{aligned}
& (--)_1 (-1)_2 (11)_3 (10)_4 (00)_5 (?)_6 (?)_7 \\
& (--)_1 (--)_2 (-1)_3 (10)_4 (00)_5 (?)_6 (?)_7 \\
& (--)_1 (--)_2 (--)_3 (-0)_4 (00)_5 (?)_6 (?)_7 \\
& (--)_1 (--)_2 (--)_3 (--)_4 (-0)_5 (?)_6 (?)_7
\end{aligned}$$

- Suffix canceling (note that the rules do not allow a canceling of only the fifth bit in B):

$$\begin{aligned}
& (-0)_1 (01)_2 (11)_3 (1-)_4 (--)_5 (?)_6 (?)_7 \\
& (-0)_1 (01)_2 (1-)_3 (--)_4 (--)_5 (?)_6 (?)_7 \\
& (-0)_1 (0-)_2 (--)_3 (--)_4 (--)_5 (?)_6 (?)_7
\end{aligned}$$

- Prefix and suffix canceling (again, the rules do not allow a suffix canceling of only the fifth bit in B):

$$\begin{aligned}
& (--)_1 (-1)_2 (11)_3 (1-)_4 (--)_5 (?)_6 (?)_7 \\
& (--)_1 (-1)_2 (1-)_3 (--)_4 (--)_5 (?)_6 (?)_7 \\
& (--)_1 (--)_2 (-1)_3 (1-)_4 (--)_5 (?)_6 (?)_7
\end{aligned}$$

Finally, we add a small number of additional configurations to $\mathcal{N}_{2,7}$ that we obtain from canceling the entire bit string B . However, when we cancel B entirely, we are required to select in the respectively obtained configuration a label at some position $j \neq \Delta - i$ and replace this label by $(!)_j$ if $j < \Delta - i$, and by $(!)_j$ otherwise. Hence, we obtain the following additional configurations that we add to $\mathcal{N}_{2,7}$.

$$\begin{aligned}
& (!!)_1 (--)_2 (--)_3 (--)_4 (--)_5 (?)_6 (?)_7 \\
& (--)_1 (!!)_2 (--)_3 (--)_4 (--)_5 (?)_6 (?)_7 \\
& (--)_1 (--)_2 (!!)_3 (--)_4 (--)_5 (?)_6 (?)_7 \\
& (--)_1 (--)_2 (--)_3 (!!)_4 (--)_5 (?)_6 (?)_7 \\
& (--)_1 (--)_2 (--)_3 (--)_4 (--)_5 (!)_6 (?)_7 \\
& (--)_1 (--)_2 (--)_3 (--)_4 (--)_5 (?)_6 (!)_7
\end{aligned}$$

Black constraint. We now define the black constraint $\mathcal{E}_{i,\Delta}$ of problem $\Pi_{i,\Delta}$. Let j be an integer satisfying $\Delta - i + 1 \leq j \leq \Delta$, i.e., a color that is gone. The allowed configurations are exactly those that are described by the set of condensed configurations containing the following:

- $[(!)_j, (?)_j]^2 \quad [(?)_j]$

Then, the allowed configurations are all the ones described by the set of condensed configurations containing the following:

- $[(!)_1, (--)_1, (-0)_1, (-1)_1] \quad [(-)_1] \quad [(-)_1, (-0)_1]$
- $[(-)_1, (-0)_1]^2 \quad [(-)_1, (-1)_1]$

Let j be an integer satisfying $2 \leq j \leq \Delta - i - 1$. The allowed configurations are all the ones described by the set of condensed configurations containing the following:

- $[(-)j] \ [(!)j, (-)j, (-0)j, (-1)j, (0-)j, (00)j, (01)j, (1-)j, (10)j, (11)j]^2$
- $[(-)j, (-0)j, (00)j, (1-)j, (10)j, (11)j] \ [(-)j, (00)j]^2$
- $[(-)j, (-1)j, (01)j, (1-)j, (10)j, (11)j] \ [(-)j, (00)j] \ [(-)j, (01)j]$
- $[(-)j, (-0)j, (00)j, (1-)j, (10)j, (11)j] \ [(-)j, (01)j]^2$
- $[(-)j, (-1)j, (0-)j, (00)j, (01)j, (11)j] \ [(-)j, (00)j] \ [(-)j, (10)j]$
- $[(-)j, (-0)j, (0-)j, (00)j, (01)j, (10)j] \ [(-)j, (00)j] \ [(-)j, (11)j]$
- $[(-)j, (-0)j, (0-)j, (00)j, (01)j, (10)j] \ [(-)j, (01)j] \ [(-)j, (10)j]$
- $[(-)j, (-1)j, (0-)j, (00)j, (01)j, (11)j] \ [(-)j, (01)j] \ [(-)j, (11)j]$
- $[(-)j, (-1)j, (01)j, (1-)j, (10)j, (11)j] \ [(-)j, (10)j]^2$
- $[(-)j, (-0)j, (00)j, (1-)j, (10)j, (11)j] \ [(-)j, (10)j] \ [(-)j, (11)j]$
- $[(-)j, (-1)j, (01)j, (1-)j, (10)j, (11)j] \ [(-)j, (11)j]^2$
- $[(-)j, (-0)j, (00)j, (10)j] \ [(-)j, (00)j] \ [(-)j, (00)j, (11)j]$
- $[(-)j, (-1)j, (01)j, (11)j] \ [(-)j, (00)j] \ [(-)j, (01)j, (10)j]$
- $[(-)j, (-0)j, (00)j, (10)j] \ [(-)j, (01)j] \ [(-)j, (01)j, (10)j]$
- $[(-)j, (-1)j, (01)j, (11)j] \ [(-)j, (00)j, (11)j] \ [(-)j, (01)j]$
- $[(-)j, (-0)j, (00)j, (10)j] \ [(-)j, (-1)j, (01)j, (11)j] \ [(-)j, (10)j]$
- $[(-)j, (11)j] \ [(-)j, (-0)j, (00)j, (10)j]^2$
- $[(-)j, (10)j] \ [(-)j, (01)j, (10)j]^2$
- $[(-)j, (10)j] \ [(-)j, (00)j, (11)j]^2$
- $[(-)j, (00)j, (11)j] \ [(-)j, (01)j, (10)j] \ [(-)j, (11)j]$
- $[(-)j, (11)j] \ [(-)j, (-1)j, (01)j, (11)j]^2$
- $[(-)j, (1-)j, (10)j, (11)j] \ [(-)j, (0-)j, (00)j, (01)j]^2$
- $[(-)j, (1-)j, (10)j, (11)j]^3$

For the case $j = \Delta - i$, i.e., for the special color, the list of allowed configurations are all the ones described by the set of condensed configurations given by the following process:

- Start from the condensed configurations given above for the case $2 \leq j \leq \Delta - i - 1$.
- From each disjunction, remove labels $(0-)j$, $(1-)j$, and $(!)j$ (i.e., labels that for color $\Delta - i$ do not exist).

Alternative definition of the black constraint. We provided the allowed configurations for the case $2 \leq j \leq \Delta - i - 1$ as a long list. We now provide an equivalent but compact description of such configurations, that will be useful later. Such configurations are based on the allowed configurations of the iterated GHZ problem. Each configuration is composed of 3 labels with subscript j , and for the moment we only consider labels containing two bits, i.e., $(00)j$, $(01)j$, $(10)j$, and $(11)j$. Later we will provide additional configurations as a function of the ones that we provide now. The first bit of each label is called input bit, and the second is called output bit. If the number of input bits equal to 1 is even, is required that the XOR of the output bits is equal to the OR of the input bits. Instead, if the number of input bits equal to 1 is odd, there is no constraint on the output bits. The configurations satisfying these requirements, which we call *bit configurations*, are the following:

- $(00)j \ (00)j \ (00)j$
- $(00)j \ (01)j \ (01)j$
- $(00)j \ (00)j \ (10)j$
- $(00)j \ (00)j \ (11)j$
- $(00)j \ (01)j \ (10)j$
- $(00)j \ (01)j \ (11)j$
- $(01)j \ (01)j \ (10)j$

- $(01)_j$ $(01)_j$ $(11)_j$
- $(00)_j$ $(10)_j$ $(11)_j$
- $(01)_j$ $(10)_j$ $(10)_j$
- $(01)_j$ $(11)_j$ $(11)_j$
- $(10)_j$ $(10)_j$ $(10)_j$
- $(10)_j$ $(10)_j$ $(11)_j$
- $(10)_j$ $(11)_j$ $(11)_j$
- $(11)_j$ $(11)_j$ $(11)_j$

We now compute the set of maximal configurations satisfying the universal quantifier w.r.t. the above set of configurations.

Lemma 4.10. *Let S_j denote the set of maximal configurations satisfying the universal quantifier w.r.t. the set of bit configurations of color j . Then S_j contains precisely the following configurations.*

- $\{(00)_j, (10)_j, (11)_j\}$ $\{(00)_j\}^2$
- $\{(01)_j, (10)_j, (11)_j\}$ $\{(00)_j\}$ $\{(01)_j\}$
- $\{(00)_j, (10)_j, (11)_j\}$ $\{(01)_j\}^2$
- $\{(00)_j, (01)_j, (11)_j\}$ $\{(00)_j\}$ $\{(10)_j\}$
- $\{(00)_j, (01)_j, (10)_j\}$ $\{(00)_j\}$ $\{(11)_j\}$
- $\{(00)_j, (01)_j, (10)_j\}$ $\{(01)_j\}$ $\{(10)_j\}$
- $\{(00)_j, (01)_j, (11)_j\}$ $\{(01)_j\}$ $\{(11)_j\}$
- $\{(01)_j, (10)_j, (11)_j\}$ $\{(10)_j\}^2$
- $\{(00)_j, (10)_j, (11)_j\}$ $\{(10)_j\}$ $\{(11)_j\}$
- $\{(01)_j, (10)_j, (11)_j\}$ $\{(11)_j\}^2$
- $\{(00)_j, (10)_j\}$ $\{(00)_j\}$ $\{(00)_j, (11)_j\}$
- $\{(01)_j, (11)_j\}$ $\{(00)_j\}$ $\{(01)_j, (10)_j\}$
- $\{(00)_j, (10)_j\}$ $\{(01)_j\}$ $\{(01)_j, (10)_j\}$
- $\{(01)_j, (11)_j\}$ $\{(00)_j, (11)_j\}$ $\{(01)_j\}$
- $\{(00)_j, (10)_j\}$ $\{(01)_j, (11)_j\}$ $\{(10)_j\}$
- $\{(11)_j\}$ $\{(00)_j, (10)_j\}^2$
- $\{(10)_j\}$ $\{(01)_j, (10)_j\}^2$
- $\{(10)_j\}$ $\{(00)_j, (11)_j\}^2$
- $\{(00)_j, (11)_j\}$ $\{(01)_j, (10)_j\}$ $\{(11)_j\}$
- $\{(11)_j\}$ $\{(01)_j, (11)_j\}^2$
- $\{(10)_j, (11)_j\}$ $\{(00)_j, (01)_j\}^2$
- $\{(10)_j, (11)_j\}^3$

Proof. By inspecting each configuration of S_j , we can observe that the universal quantifier is satisfied on all of them. It is straightforward to verify that none of the displayed configurations dominates any of the other displayed configurations. Hence, it suffices to show that S_j contains all the maximal configurations satisfying the universal quantifier. Assume for a contradiction that there exists a maximal configuration C satisfying the universal quantifier that is not dominated by any configuration present in S_j .

We start by proving that C cannot contain two sets of size 1, let them be C_1 and C_2 . As it is straightforward to check, the set S_j contains all configurations satisfying the universal quantifier composed of two sets of size 1 and one set of size 3, and that all configurations satisfying the universal quantifier composed of two sets of size 1 and one set of size at most 2 are dominated by configurations in S_j . Hence, for the configuration C to not be in S_j , it must contain a set of size 4,

let it be C_3 . Consider the input bits i_1 and i_2 of the labels in the sets C_1 and C_2 . If the number of input bits in (i_1, i_2) equal to 1 is even, then C_3 cannot contain both $(00)_j$ and $(01)_j$. If the number of input bits in (i_1, i_2) equal to 1 is odd, then C_3 cannot contain both $(10)_j$ and $(11)_j$. Hence, C_3 has size at most 3.

Since a set cannot be empty, and hence it must have size at least 1, the above also implies that C cannot contain sets of size 4 (even if the other sets have size strictly larger than 1).

Moreover, consider a configuration C containing two sets of size 1 (C_1 and C_2) and one set of size 3 (C_3). By adding any element L to some set of size 1, w.l.o.g. to C_1 , the obtained configuration does not satisfy the universal quantifier anymore, for the following reasons.

- Consider the case in which L has an input bit that is the same as i_1 , and hence the opposite output bit. Since C_3 has size 3, from C we can always pick a configuration with an even number of input bits set to 1. If we now replace the element picked from C_1 with L , we obtain an invalid configuration.
- If L has an input bit that is different from i_1 , then we can pick two configurations that only differ in the element picked from C_3 and that have an even number of input bits set to 1, but that have different parity in the amount of output bits set to 1, and hence one of the two is invalid.

Hence, if the number of sets of size 1 is strictly less than two, it is not possible to have sets of size 3.

We now consider the case in which C contains exactly one set of size 1, let it be C_1 . The other sets, C_2 and C_3 , must be of size 2. Since C is maximal, it is not possible to add an element to C_1 , and in particular the label with the same input and opposite output bit of the label in C_1 , without losing the property that the configuration satisfies the universal quantifier. This implies that it must be possible to pick, from C , a configuration with an even number of input bits set to 1. Let such a configuration be $P_1 P_2 P_3$, and let the non-picked labels be $N_2 \in C_2$ and $N_3 \in C_3$. It cannot be that P_2 and N_2 (resp. P_3 and N_3) have the same input and opposite output bit, as replacing P_2 with N_2 (resp. P_3 with N_3) in C would result in a configuration that is not allowed (since the amount of input bits equal to 1 stays even, but the parity of the output changes). Moreover, the configuration $P_1 N_2 N_3$ still needs to be allowed. It is straightforward to verify that all configurations satisfying these requirements are listed.

Finally, let us consider the case in which C contains only sets of size exactly 2. We prove that each set must satisfy that all its elements have the same input bit, and that the number of sets whose elements have input bit 1 must be odd. Given such a proof, we get a contradiction to our initial assumption by observing that there are only two configurations satisfying these requirements, and they are both listed.

Suppose that all configurations that can be picked from C have an odd number of input bits set to 1. This implies that all sets have elements with the same input bit, since otherwise we could replace a picked element with a non-picked one and change the parity. Now, suppose for a contradiction that we can pick a configuration with an even number of input bits set to 1. If any set contains a non-picked element with the same input bit as the picked one but the opposite output bit, we would get a contradiction. Hence, each non-picked element must have opposite input bit from the picked one of the same set. Hence, it must be possible to pick $(00)_j^3$, or $(00)_j (01)_j^2$. In the former case, we must be able to pick $(00)_j (10)_j^2$ or $(00)_j (11)_j^2$, which are both not allowed. In the latter case, we can pick $(00)_j (10)_j^2$, $(00)_j (11)_j^2$, or $(10)_j (01)_j (11)_j$, which are all not allowed. This completes our proof by contradiction. \square

We now modify the set S_j defined in Lemma 4.10 by applying the following rules.

- If a set contains both $(00)_j$ and $(01)_j$, add $(0-)_j$.
- If a set contains both $(10)_j$ and $(11)_j$, add $(1-)_j$.
- If a set contains both $(00)_j$ and $(10)_j$, add $(-0)_j$.
- If a set contains both $(01)_j$ and $(11)_j$, add $(-1)_j$.
- Always add $(--)_j$.

Finally, we add the following configuration C^* :

$$\{(--)_j\} \quad \{(!)_j, (--)_j, (-0)_j, (-1)_j, (0-)_j, (00)_j, (01)_j, (1-)_j, (10)_j, (11)_j\}^2.$$

Let the result be S'_j . Observe that, for all $2 \leq j \leq \Delta - i - 1$, S'_j is exactly equal to the set of condensed configurations listed when defining $\mathcal{E}_{i,\Delta}$. We observe the following (which will be useful later).

Observation 4.11. *For all $2 \leq j \leq \Delta - i - 1$, the set S'_j is maximal-complete.*

Proof. We prove that, if we take two arbitrary configurations C^1 and C^2 in S'_j , and we combine them w.r.t. an arbitrary u and an arbitrary σ , we obtain a configuration $C = C_1 C_2 C_3$ that is dominated by some other configuration in S'_j . By Observation 4.3 this implies that S'_j is maximal-complete.

We start by observing that, by Observation 4.4, we do not need to consider the case in which $C^1 = C^*$ or $C^2 = C^*$.

If an intersection results in $(--)_j$, then C is dominated by the configuration C^* . Otherwise, let C'^1 and C'^2 be the configurations obtained by starting from C^1 and C^2 and removing the labels $(0-)_j$, $(1-)_j$, $(-0)_j$, $(-1)_j$, and $(--)_j$ from all the sets of C^1 and C^2 . We obtain that C'^1 and C'^2 are configurations listed in Lemma 4.10. Let C' be the configuration obtained by combining C'^1 and C'^2 w.r.t. u and σ , and let D' be the configuration dominating C' among the configurations listed in Lemma 4.10 (which is guaranteed to exist by Lemma 4.10). In each set of D' add the labels $(0-)_j$, $(1-)_j$, $(-0)_j$, $(-1)_j$, $(--)_j$ by using the same rules applied when defining S'_j as a function of S_j , obtaining the configuration $D = D_1 D_2 D_3$. Clearly, the configuration D is present in S'_j . Moreover, the following observations show that D dominates C .

- The label $(--)_j$ is present in all sets of both C and D , and hence does not affect domination.
- If some label $L \in \{(00)_j, (01)_j, (10)_j, (11)_j\}$ is present in C_i , for some $i \in \{1, 2, 3\}$, then L is present also in D_i .
- If C_u contains a label L in $\{(0-)_j, (1-)_j, (-0)_j, (-1)_j\}$, then $L \in D_u$. To see this, consider w.l.o.g. $L = (0-)_j$. We get that $(00)_j \in C_u^1$ or $(00)_j \in C_{\sigma(u)}^2$, and $(01)_j \in C_u^1$ or $(01)_j \in C_{\sigma(u)}^2$. Hence, $\{(00)_j, (01)_j\} \subseteq D_u$, and hence $(0-)_j \in D_u$.
- Let $i \in \{1, 2, 3\}$ and $i \neq u$. If C_i contains a label L in $\{(0-)_j, (1-)_j, (-0)_j, (-1)_j\}$, then $L \in D_i$. To see this, consider w.l.o.g. $L = (0-)_j$. We get that $(00)_j \in C_i^1$, $(00)_j \in C_{\sigma(i)}^2$, $(01)_j \in C_i^1$, and $(01)_j \in C_{\sigma(i)}^2$. Hence, $\{(00)_j, (01)_j\} \subseteq D_i$, and hence $(0-)_j \in D_i$.

Observe that this covers all cases, since $(!)_j$ only occurs in C^* . □

4.4 The diagram of $\Pi_{i,\Delta}$

Throughout the rest of Section 4, we assume that $i \leq \Delta - 3$. We now provide the diagram of $\Pi_{i,\Delta}$, that is, we prove the strength relation between the labels of $\Pi_{i,\Delta}$ w.r.t. its black constraint. In Appendix A we will first observe that, a necessary condition for a label to be at least as strong as another label, is for the two labels to be of the same color. Then, we will prove Lemmas 4.12 to 4.15 (which are just a long case analysis).




Figure 1: The strength relation of the first color.




Figure 2: The strength relation of each present color j different from the first and the special color.




Figure 3: The strength relation of the special color $j = \Delta - i - 1$.




Figure 4: The strength relation of gone colors.

Lemma 4.12. *The strength relation of the labels of color 1 is given by the diagram of Figure 1.*

Lemma 4.13. *Let j be a color satisfying $2 \leq j \leq \Delta - i - 1$. The strength relation of the labels of color j is given by the diagram of Figure 2.*

Lemma 4.14. *Let j be the special color, that is, $j = \Delta - i$. The strength relation of the labels of color j is given by the diagram of Figure 3.*

Lemma 4.15. *Let j be a gone color, i.e., an integer satisfying $\Delta - i + 1 \leq j \leq \Delta$. The strength relation of the labels of color j is given by the diagram of Figure 4.*

4.5 The first round elimination step

In this section, we define a family of problems $\Pi'_{i,\Delta}$, and we prove that $\mathcal{R}(\Pi_{i,\Delta})$ can be relaxed to $\Pi'_{i,\Delta}$. Each label of $\Pi'_{i,\Delta}$ is going to be a set of labels of $\Pi_{i,\Delta}$. Recall that with $\langle L_1, \dots, L_k \rangle$ we denote the set of labels containing L_1, \dots, L_k and all of their (not necessarily direct) successors in the diagram of $\Pi_{i,\Delta}$.

Black constraint. We now define the black constraint $\mathcal{E}'_{i,\Delta}$ of $\Pi'_{i,\Delta}$. For each j where j is a gone color, the following configuration is allowed:

- $\langle (!)_j \rangle^2 \quad \langle (?)_j \rangle$

For color 1, the following configurations are allowed:

- $\langle (!)_1 \rangle \quad \langle (--)_1 \rangle \quad \langle (-0)_1 \rangle$
- $\langle (-0)_1 \rangle^2 \quad \langle (-1)_1 \rangle$

For each color j satisfying $2 \leq j \leq \Delta - i - 1$, the following configurations are allowed:

- $\langle (--)_j \rangle \quad \langle (!)_j \rangle^2$
- $\langle (-0)_j, (1-)_j \rangle \quad \langle (00)_j \rangle^2$
- $\langle (-1)_j, (1-)_j \rangle \quad \langle (00)_j \rangle \quad \langle (01)_j \rangle$
- $\langle (-0)_j, (1-)_j \rangle \quad \langle (01)_j \rangle^2$
- $\langle (-1)_j, (0-)_j \rangle \quad \langle (00)_j \rangle \quad \langle (10)_j \rangle$
- $\langle (-0)_j, (0-)_j \rangle \quad \langle (00)_j \rangle \quad \langle (11)_j \rangle$
- $\langle (-0)_j, (0-)_j \rangle \quad \langle (01)_j \rangle \quad \langle (10)_j \rangle$
- $\langle (-1)_j, (0-)_j \rangle \quad \langle (01)_j \rangle \quad \langle (11)_j \rangle$
- $\langle (-1)_j, (1-)_j \rangle \quad \langle (10)_j \rangle^2$
- $\langle (-0)_j, (1-)_j \rangle \quad \langle (10)_j \rangle \quad \langle (11)_j \rangle$
- $\langle (-1)_j, (1-)_j \rangle \quad \langle (11)_j \rangle^2$
- $\langle (-0)_j \rangle \quad \langle (00)_j \rangle \quad \langle (00)_j, (11)_j \rangle$
- $\langle (-1)_j \rangle \quad \langle (00)_j \rangle \quad \langle (01)_j, (10)_j \rangle$
- $\langle (-0)_j \rangle \quad \langle (01)_j \rangle \quad \langle (01)_j, (10)_j \rangle$
- $\langle (-1)_j \rangle \quad \langle (00)_j, (11)_j \rangle \quad \langle (01)_j \rangle$
- $\langle (-0)_j \rangle \quad \langle (-1)_j \rangle \quad \langle (10)_j \rangle$
- $\langle (11)_j \rangle \quad \langle (-0)_j \rangle^2$
- $\langle (10)_j \rangle \quad \langle (01)_j, (10)_j \rangle^2$
- $\langle (10)_j \rangle \quad \langle (00)_j, (11)_j \rangle^2$
- $\langle (00)_j, (11)_j \rangle \quad \langle (01)_j, (10)_j \rangle \quad \langle (11)_j \rangle$
- $\langle (11)_j \rangle \quad \langle (-1)_j \rangle^2$
- $\langle (1-)_j \rangle \quad \langle (0-)_j \rangle^2$
- $\langle (1-)_j \rangle^3$

For the special color $j = \Delta - i$, the following configurations are allowed:

- $\langle (00)_j, (01)_j \rangle \quad \langle (-0)_j, (-1)_j \rangle^2$
- $\langle (10)_j, (11)_j \rangle \quad \langle (-0)_j, (-1)_j \rangle^2$

This concludes the definition of the black constraint.

Before proving the relation between the black constraint of $\mathcal{R}(\Pi_{i,\Delta})$ and $\mathcal{E}'_{i,\Delta}$, we will state a useful observation, based on the following definitions. For any set S that contains only labels of the same color, say j , we define the color of S as j . For any configuration C that contains only sets of color j , we define the color of C as j .

In order for a configuration of the black constraint of $\mathcal{R}(\Pi_{i,\Delta})$ to satisfy the universal quantifier, the following must clearly hold.

Observation 4.16. *All labels of $\mathcal{R}(\Pi_{i,\Delta})$ are sets of labels of $\Pi_{i,\Delta}$, where each set S satisfies that all labels in S have the same color. Moreover, each configuration of the black constraint of $\mathcal{R}(\Pi_{i,\Delta})$ consists of sets of the same color.*

In the following, we will implicitly use the fact that, for any set \mathcal{L} of labels, if two labels satisfy $L_i \leq L_j$, and $L_i, L_j \in \mathcal{L}$, then $\langle L \mid L \in \mathcal{L} \rangle = \langle L \mid L \in \mathcal{L} \setminus \{L_j\} \rangle$.

Lemma 4.17. *The black constraint of $\mathcal{R}(\Pi_{i,\Delta})$ is dominated by $\mathcal{E}'_{i,\Delta}$.*

Proof. By Observation 4.16, each configuration of $\mathcal{R}(\Pi_{i,\Delta})$ is composed of labels with the same color. In order to prove the statement, we consider each color j separately.

- For all $j \neq \Delta - i$, the configurations of $\mathcal{E}'_{i,\Delta}$ are exactly the condensed configurations provided when describing $\mathcal{E}_{i,\Delta}$. We prove that, by combining configurations of color j , we do not get any additional configuration. By Observation 4.3, this implies that the configurations of color j are maximal-complete.
- For $j = \Delta - i$, we prove that each configuration of color j of $\mathcal{R}(\Pi_{i,\Delta})$ is dominated by (at least) one of the two configurations of color j that we provided.

Let j be a gone color, i.e., $j > \Delta - i$. By Observation 4.4, when combining the unique configuration of color j with itself (for any index u and permutation σ), we do not obtain any new configuration, since $\langle (!)_j \rangle$ and $\langle (?)_j \rangle$ are comparable.

Let us now consider color 1. By Observation 4.4, by combining $\langle (!!)_1 \rangle \langle (--)_1 \rangle \langle (-0)_1 \rangle$ with itself, we cannot get additional configurations (not dominated by ones already present), since all sets in this configuration are comparable with each other. By Observation 4.4, in order to get new configurations when combining $\langle (-0)_1 \rangle^2 \langle (-1)_1 \rangle$ with itself, we need to take the union of $\langle (-0)_1 \rangle$ with $\langle (-1)_1 \rangle$, thus obtaining the configuration $\langle (-0)_1, (-1)_1 \rangle \langle (-0)_1 \rangle \langle (--)_1 \rangle$, which is dominated by $\langle (!!)_1 \rangle \langle (--)_1 \rangle \langle (-0)_1 \rangle$. Finally, if we combine $\langle (!!)_1 \rangle \langle (--)_1 \rangle \langle (-0)_1 \rangle$ with $\langle (-0)_1 \rangle^2 \langle (-1)_1 \rangle$, by Observation 4.4 we need to take the union of $\langle (-0)_1 \rangle$ with $\langle (-1)_1 \rangle$, thus obtaining the configuration $\langle (-0)_1, (-1)_1 \rangle \langle (-0)_1 \rangle \langle (--)_1 \rangle$, which again is dominated by $\langle (!!)_1 \rangle \langle (--)_1 \rangle \langle (-0)_1 \rangle$.

In the case $2 \leq j \leq \Delta - i - 1$, the claim follows by applying Observation 4.11 (and the definition of \mathcal{R}).

Finally, we consider the case $j = \Delta - i$, i.e., the special color. In this case, we prove that all the configurations satisfying the universal quantifier are dominated by (at least) one of the two configurations present in $\mathcal{E}'_{i,\Delta}$. We first observe that, since $\langle (-0)_j, (-1)_j \rangle$ contains all labels of color j (see Figure 3), any configuration $C = C_1 C_2 C_3$ of color j in the black constraint of $\mathcal{R}(\Pi_{i,\Delta})$ that is not dominated by the two configurations present in $\mathcal{E}'_{i,\Delta}$ must necessarily satisfy that all its sets C_i satisfy $C_i \not\subseteq \langle (00)_j, (01)_j \rangle$ and $C_i \not\subseteq \langle (10)_j, (11)_j \rangle$. By Observation 4.7, each set in such a configuration must be right-closed, and hence it holds that, for all i , C_i is a superset of at least one of the following four sets: $X_1 = \{(00)_j, (10)_j\}$, $X_2 = \{(00)_j, (11)_j\}$, $X_3 = \{(01)_j, (10)_j\}$, $X_4 = \{(01)_j, (11)_j\}$. We consider all possible choices in $\{X_1, X_2, X_3, X_4\}^3$ (excluding permutations):

- For any choice in $\{X_3, X_4\}^3$, we can pick the configuration $(01)_j^3$, which is not present in $\mathcal{E}_{i,\Delta}$.
- For any choice in $\{X_1, X_2\}^3$, we can pick the configuration $(00)_j (10)_j^2$ or the configuration $(00)_j (11)_j^2$, which are not present in $\mathcal{E}_{i,\Delta}$.
- For any choice in $\{X_1, X_2\}^2 \{X_3, X_4\}$, we can pick the configuration $(00)_j^2 (01)_j$, which is not present in $\mathcal{E}_{i,\Delta}$.
- For any choice in $\{X_1, X_2\} \{X_3\}^2$, we can pick the configuration $(00)_j (10)_j^2$, which is not present in $\mathcal{E}_{i,\Delta}$.
- For any choice in $\{X_1, X_2\} \{X_4\}^2$, we can pick the configuration $(00)_j (11)_j^2$, which is not present in $\mathcal{E}_{i,\Delta}$.

- For any choice in $\{X_1, X_2\} \{X_3\} \{X_4\}$, we can pick the configuration $(10)_j (01)_j (11)_j$, which is not present in $\mathcal{E}_{i,\Delta}$. \square

White constraint. We first define an intermediate constraint $\mathcal{N}'_{i,\Delta}$. Then, we define $\mathcal{N}''_{i,\Delta}$ as a function of $\mathcal{N}'_{i,\Delta}$. The constraint $\mathcal{N}''_{i,\Delta}$ is defined as follows. Let $x = \Delta - i$ be the special color. For each configuration $C \in \mathcal{N}'_{i,\Delta}$, the constraint $\mathcal{N}''_{i,\Delta}$ contains the condensed configuration C^* obtained by replacing each label of C with the disjunction of sets according to the following rules (where j denotes the color of L):

- $\langle \{L\} \rangle$, if $j \neq x$;
- $\langle \{(00)_j, (01)_j\} \rangle$, if $j = x$ and $L \in \{(00)_j, (01)_j\}$;
- $\langle \{(10)_j, (11)_j\} \rangle$, if $j = x$ and $L \in \{(10)_j, (11)_j\}$;
- $\langle \{(-0)_j, (-1)_j\} \rangle$, if $j = x$ and $L \in \{(-0)_j, (-1)_j\}$;
- $\langle \{(00)_j, (01)_j\}, \langle \{(10)_j, (11)_j\} \rangle$, if $j = x$ and $L = (--)_j$.

Let $\Sigma'_{i,\Delta}$ be the set of labels that appear in $\mathcal{E}'_{i,\Delta}$. The white constraint $\mathcal{N}'_{i,\Delta}$ of $\Pi'_{i,\Delta}$ contains all the configurations described by the set of condensed configurations given by the following process. Take a condensed configuration C^* in $\mathcal{N}''_{i,\Delta}$, and replace each disjunction of labels $[L_1 \dots L_k]$ with the disjunction containing L_1, \dots, L_k and all the labels $L \in \Sigma'_{i,\Delta}$ that are supersets of at least one label in $\{L_1, \dots, L_k\}$.

Lemma 4.18. $\mathcal{N}'_{i,\Delta}$ is the set of all configurations $S_1 \dots S_\Delta$ of sets contained in $\Sigma'_{i,\Delta}$ such that there exists some tuple $(\ell_1, \dots, \ell_\Delta) \in S_1 \times \dots \times S_\Delta$ satisfying that $\ell_1 \dots \ell_\Delta$ is contained in $\mathcal{N}_{i,\Delta}$.

Proof. Consider the collection \mathcal{C} of all configurations obtained from $\mathcal{N}_{i,\Delta}$, by replacing, in each configuration, each label L with the disjunction containing all sets $L' \in \Sigma'_{i,\Delta}$ satisfying $L \in L'$.

It is sufficient to prove that the collection of configurations that can be picked from \mathcal{C} is equal to $\mathcal{N}'_{i,\Delta}$. This follows directly by the definition of $\mathcal{N}'_{i,\Delta}$ and Lemmas 4.12 to 4.15. \square

By combining Lemma 4.17 and Lemma 4.18, by the definition of relaxation, we obtain the following.

Lemma 4.19. The problem $\Pi'_{i,\Delta}$ is a relaxation of $\mathcal{R}(\Pi_{i,\Delta})$.

White constraint example. We now provide an example of $\mathcal{N}'_{i,\Delta}$, for the case $\Delta = 7$ and $i = 2$, by showing some examples of replacements.

In $\mathcal{N}_{2,7}$, the following configuration is present:

$$(-0)_1 (01)_2 (11)_3 (10)_4 (00)_5 (?)_6 (?)_7.$$

We get that, in $\mathcal{N}'_{2,7}$, the following configuration is present:

$$\langle (-0)_1 \rangle \langle (01)_2 \rangle \langle (11)_3 \rangle \langle (10)_4 \rangle \langle (00)_5, (01)_5 \rangle \langle (?)_6 \rangle \langle (?)_7 \rangle.$$

Moreover, $\mathcal{N}'_{2,7}$ contains all configurations that can be obtained by starting from such a configuration and replacing some labels with arbitrary supersets. For example, the following configuration is also present:

$$\langle (-0)_1 \rangle \langle (01)_2 \rangle \langle (11)_3 \rangle \langle (10)_4 \rangle \langle (00)_5, (01)_5 \rangle \langle (!)_6 \rangle \langle (!)_7 \rangle.$$

For another example, consider the following configuration present in $\mathcal{N}_{2,7}$:

$$(-0)_1 (01)_2 (1-)_3 (--)_4 (--)_5 (?)_6 (?)_7.$$

We get that, in $\mathcal{N}'_{2,7}$, all the configurations given by the following condensed configuration (plus all the ones obtained by replacing labels with supersets) are present:

$$[\langle(-0)_1\rangle] [\langle(01)_2\rangle] [\langle(1-)_3\rangle] [\langle(--)_4\rangle] [\langle(00)_5, (01)_5, \langle(10)_5, (11)_5\rangle] [\langle(?)_6\rangle] [\langle(?)_7\rangle].$$

For another example, consider the following configuration present in $\mathcal{N}'_{2,7}$:

$$(--)_1 \quad (--)_2 \quad (--)_3 \quad (--)_4 \quad (-0)_5 \quad (?)_6 \quad (?)_7$$

We get that, in $\mathcal{N}'_{2,7}$, the following configuration is present, plus all the ones obtained by replacing labels with supersets:

$$\langle(--)_1\rangle \langle(--)_2\rangle \langle(--)_3\rangle \langle(--)_4\rangle \langle(-0)_5, (-1)_5\rangle \langle(?)_6\rangle \langle(?)_7\rangle$$

We observe that the constraint $\mathcal{N}'_{i,\Delta}$ can be described in the following alternative (and more explicit) way, that will be useful later. Let $\mathcal{N}'_{i,\Delta}$ be the set containing the following condensed configurations.

- **Strikethrough configurations.** For all integers a, b satisfying $1 \leq a < b \leq \Delta - i$, and for all vectors $(y_j \mid y_j \in \{0, 1\}$ and $a \leq j < b)$, the condensed configuration $L_1 \dots L_\Delta$ is allowed, where:

- For all j such that $1 \leq j < a$, $L_j = [\langle(--)_j\rangle]$;
- $L_a = [\langle(-y_a)_a\rangle]$;
- For all j such that $a < j < b$, $L_j = [\langle(y_{j-1} y_j)_j\rangle]$;
- If $b \leq \Delta - i - 1$, $L_b = [\langle(y_{b-1}-)_b\rangle]$;
- For all j such that $b < j \leq \Delta - i - 1$, $L_j = [\langle(--)_j\rangle]$;
- If $b \neq \Delta - i$, $L_{\Delta-i} = [\langle(00)_{\Delta-i}, (01)_{\Delta-i}, \langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle]$;
- If $b = \Delta - i$, $L_{\Delta-i} = [\langle(y_{b-1}0)_{\Delta-i}, (y_{b-1}1)_{\Delta-i}\rangle]$;
- For all j such that $\Delta - i + 1 \leq j \leq \Delta$, $L_j = [\langle(?)_j\rangle]$.

- **Grabbing-present configurations.** For all integers a satisfying $1 \leq a \leq \Delta - i - 1$, the condensed configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_a = [\langle(!!)_a\rangle]$;
- For all j such that $1 \leq j \leq \Delta - i - 1$ and $j \neq a$, $L_j = [\langle(--)_j\rangle]$;
- $L_{\Delta-i} = [\langle(00)_{\Delta-i}, (01)_{\Delta-i}, \langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle]$;
- For all j such that $\Delta - i + 1 \leq j \leq \Delta$, $L_j = [\langle(?)_j\rangle]$.

- **Grabbing-gone configurations.** For all integers a satisfying $\Delta - i + 1 \leq a \leq \Delta$, the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_a = [\langle(!)_a\rangle]$;
- For all j such that $1 \leq j \leq \Delta - i - 1$, $L_j = [\langle(--)_j\rangle]$;
- $L_{\Delta-i} = [\langle(00)_{\Delta-i}, (01)_{\Delta-i}, \langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle]$;
- For all j such that $\Delta - i + 1 \leq j \leq \Delta$ and $j \neq a$, $L_j = [\langle(?)_j\rangle]$.

- **Grabbing-special configuration.** The configuration $L_1 \dots L_\Delta$ is allowed where:

- For all j such that $1 \leq j \leq \Delta - i - 1$, $L_j = [\langle(--)_j\rangle]$;

- $L_{\Delta-i} = [((-0)_{\Delta-i}, (-1)_{\Delta-i})]$;
- For all j such that $\Delta - i + 1 \leq j \leq \Delta$, $L_j = [((?)_j)]$.

Then, let $\mathcal{N}_{i,\Delta}^*$ be the set of condensed configurations obtained as follows. Take a configuration C^+ from $\mathcal{N}_{i,\Delta}^+$. Replace, in C^+ , each disjunction of labels $[L_1 \dots L_k]$ with the disjunction containing L_1, \dots, L_k and all the labels $L \in \Sigma'_{i,\Delta}$ that are supersets of at least one label in $\{L_1, \dots, L_k\}$.

By the definitions of $\mathcal{N}_{i,\Delta}^*$ and $\mathcal{N}'_{i,\Delta}$, we obtain the following observation.

Observation 4.20. *The white constraint $\mathcal{N}'_{i,\Delta}$ is the set containing all the configurations C that can be picked from $\mathcal{N}_{i,\Delta}^*$.*

4.6 The diagram of $\Pi'_{i,\Delta}$

We now provide the diagram of $\Pi'_{i,\Delta}$, that is, we prove the strength relation between the labels of $\Pi'_{i,\Delta}$ w.r.t. its white constraint. Again, a necessary condition for a label to be at least as strong as another label is for the two labels to be of the same color.




Figure 5: The strength relation of the first color.




Figure 6: The strength relation of each present color j different from the first and the special color.




Figure 7: The strength relation of the special color $j = \Delta - i - 1$.

Lemma 4.21. *The strength relation of the labels of color 1 is given by the diagram of Figure 5.*

$$\langle(?)\rangle \longrightarrow \langle(!)\rangle$$

Figure 8: The strength relation of gone colors.

Lemma 4.22. *Let j be a color satisfying $2 \leq j \leq \Delta - i - 1$. The strength relation of the labels of color j is given by the diagram of Figure 6.*

Lemma 4.23. *Let j be the special color, that is, $j = \Delta - i$. The strength relation of the labels of color j is given by the diagram of Figure 7.*

Lemma 4.24. *Let j be a gone color, i.e., an integer satisfying $\Delta - i + 1 \leq j \leq \Delta$. The strength relation of the labels of color j is given by the diagram of Figure 8.*

The proofs of Lemmas 4.21 to 4.24 are just a long case analysis, and hence they are deferred to Appendix B. By the above lemmas, we get the following.

Observation 4.25. $L_1 \leq L_2 \iff L_1 \subseteq L_2$.

4.7 The second round elimination step

In this section, we define a family of problems $\Pi''_{i,\Delta}$, and we prove that $\overline{\mathcal{R}}(\Pi'_{i,\Delta})$ can be relaxed to $\Pi''_{i,\Delta}$. Then, we show how to relax $\overline{\mathcal{R}}(\Pi'_{i,\Delta})$ even more in order to obtain $\Pi_{i+1,\Delta}$. Each label of $\Pi''_{i,\Delta}$ is going to be a set of labels of $\Pi'_{i,\Delta}$. Recall that with $\langle L_1, \dots, L_k \rangle$ (where L_1, \dots, L_k are labels of $\Pi'_{i,\Delta}$) we denote the set of labels obtained by taking L_1, \dots, L_k and all their successors in the diagram of $\Pi'_{i,\Delta}$.

White constraint. The white constraint $\mathcal{N}''_{i,\Delta}$ of $\Pi''_{i,\Delta}$ is defined as the set $\mathcal{N}^*_{i,\Delta}$ described immediately before Observation 4.20. By Observation 4.25, we obtain that an alternative but equivalent definition of $\mathcal{N}''_{i,\Delta}$ is the following.

Observation 4.26. $\mathcal{N}''_{i,\Delta}$ is equal to the set containing, for each condensed configuration C^+ in $\mathcal{N}^+_{i,\Delta}$, the configuration generated by C^+ .

We will prove that the white constraint of $\overline{\mathcal{R}}(\Pi'_{i,\Delta})$ is dominated by $\mathcal{N}''_{i,\Delta}$. For this purpose, we first define a function f that takes as input two (condensed) strikethrough configurations $C_1 = L_{1,1} \dots L_{1,\Delta}$ and $C_2 = L_{2,1} \dots L_{2,\Delta}$ from $\mathcal{N}^+_{i,\Delta}$, an index h satisfying $1 \leq h \leq \Delta - i - 1$, an index $1 \leq u \leq \Delta$, and a value $k \in \{1, 2\}$ as follows. First, for $j \in \{1, 2\}$, let $t_{j,h} = (t_{j,h,1}, t_{j,h,2})$ be defined as:

- $t_{j,h} = (-, -)$ if $L_{j,h} = [\langle(- -)_h\rangle]$;
- $t_{j,h} = (b, -)$ if $L_{j,h} = [\langle(b -)_h\rangle]$;
- $t_{j,h} = (-, b)$ if $L_{j,h} = [\langle(- b)_h\rangle]$;
- $t_{j,h} = (b_1, b_2)$ if $L_{j,h} = [\langle(b_1 b_2)_h\rangle]$.

Then, let $f(C_1, C_2, h, u, k) =$

- - if $h = u$;
- - if $h \neq u$ and $\{t_{1,h,k}, t_{2,h,k}\} \in \{\{-\}, \{0, 1\}\}$;
- 0 if $h \neq u$ and $\{t_{1,h,k}, t_{2,h,k}\} \in \{\{0\}, \{0, -\}\}$;
- 1 if $h \neq u$ and $\{t_{1,h,k}, t_{2,h,k}\} \in \{\{1\}, \{1, -\}\}$.

Lemma 4.27. Let $T_{h,k} = f(C_1, C_2, h, u, k)$. Let $L = [\langle (T_{h,1} T_{h,2})_h \rangle]$. Then,

- if $h = u$, $\langle \ell \mid \ell \in L \rangle \supseteq \langle \ell \mid \ell \in L_{1,h} \rangle \cup \langle \ell \mid \ell \in L_{2,h} \rangle$;
- if $h \neq u$, $\langle \ell \mid \ell \in L \rangle \supseteq \langle \ell \mid \ell \in L_{1,h} \rangle \cap \langle \ell \mid \ell \in L_{2,h} \rangle$.

Proof. By the definition of f , if $h = u$, then $L = [\langle (--)_h \rangle]$, and the set $\langle \ell \mid \ell \in L \rangle$ contains all labels of color h (see Figures 5 and 6). Hence in this case the claim holds.

If $h \neq u$, by inspecting the diagrams in Figures 5 and 6, it is straightforward to verify that it always holds that $\langle \ell \mid \ell \in L_{1,h} \rangle \cap \langle \ell \mid \ell \in L_{2,h} \rangle$ (i.e., the intersection between the set containing the unique label contained in the disjunction $L_{1,h}$ and all successors of that label, and the set containing the unique label contained in the disjunction $L_{2,h}$ and all successors of that label) is a subset of the set containing the unique label contained in the disjunction L and all successors of that label. \square

Lemma 4.28. Assume that $1 \leq u \leq \Delta - i - 1$, and let ℓ_1 denote the unique label in $L_{1,u}$ and ℓ_2 the unique label in $L_{2,u}$. Assume further that ℓ_1 and ℓ_2 are incomparable. Then,

- for all $1 \leq u < \Delta - i - 1$, it holds that, if $T_{u+1,1} \neq -$, then $\{t_{1,u,1}, t_{2,u,1}\} = \{0, 1\}$ or $\{t_{1,u}, t_{2,u}\} = \{(x_1, -), (-, x_2)\}$ for some $x_1, x_2 \in \{0, 1\}$, and
- for all $1 < u \leq \Delta - i - 1$, it holds that, if $T_{u-1,2} \neq -$, then $\{t_{1,u,2}, t_{2,u,2}\} = \{0, 1\}$ or $\{t_{1,u}, t_{2,u}\} = \{(x_1, -), (-, x_2)\}$ for some $x_1, x_2 \in \{0, 1\}$.

Proof. For all $1 \leq u < \Delta - i - 1$, if $T_{u+1,1} \neq -$, it must hold that $\{t_{1,u,2}, t_{2,u,2}\}$ (i.e., the set of output bits in position u of C_1 and C_2) is in $\{\{0, -\}, \{1, -\}, \{0\}, \{1\}\}$. By Lemma 4.22 (and Lemma 4.21) it must hold that either $\{t_{1,u,1}, t_{2,u,1}\} = \{0, 1\}$ or $\{t_{1,u}, t_{2,u}\} = \{(x_1, -), (-, x_2)\}$ for some $x_1, x_2 \in \{0, 1\}$. The other case is symmetric. \square

Observation 4.29. For all $1 \leq j < \Delta - i - 1$ satisfying $j \neq u$ and $j + 1 \neq u$, $T_{j,2} = T_{j+1,1}$.

Lemma 4.30. The white constraint of $\overline{\mathcal{R}}(\Pi'_{i,\Delta})$ is dominated by $\mathcal{N}''_{i,\Delta}$.

Proof. By definition, $\mathcal{N}''_{i,\Delta}$ contains exactly the condensed configurations used to describe $\mathcal{N}'_{i,\Delta}$. Hence, we only need to prove that $\mathcal{N}''_{i,\Delta}$ is maximal-complete.

In the following, we use the description of $\mathcal{N}''_{i,\Delta}$ given by Observation 4.26. There are 4 types of configurations in $\mathcal{N}^+_{i,\Delta}$, and by Lemmas 4.21 to 4.24, for each color $1 \leq j \leq \Delta$, all sets of labels of color j appearing in at least one of the configurations generated by the condensed configurations of types grabbing-present, grabbing-gone, and grabbing-special are comparable with all sets of labels of color j appearing in at least one of the configurations generated by at least one of the condensed configurations in $\mathcal{N}^+_{i,\Delta}$. Moreover, in order to combine two configurations C_1 and C_2 and obtain a configuration C that does not contain any empty set, it must clearly hold that the union and all the intersections must be taken on labels of the same color (i.e., the permutation σ used to combine C_1 with C_2 must always pair labels of the same color). Hence, by Observation 4.4, the only way we could get additional configurations is by combining two configurations generated by strikethrough configurations.

Hence, let $C_1 = L_{1,1} \dots L_{1,\Delta}$ and $C_2 = L_{2,1} \dots L_{2,\Delta}$ be two strikethrough configurations. Let a_1, a_2, b_1, b_2 be their parameters and $y_{1,j}, y_{2,j}$ their bits. Let u be the position of the union, and note that the permutation σ used to combine the configurations generated by C_1 and C_2 must be the identity function, since otherwise, when taking intersections, we would obtain empty sets. Moreover, for the sake of contradiction, assume that the configuration obtained by combining the configuration generated by C_1 and the configuration generated by C_2 w.r.t. u and σ is not contained in $\mathcal{N}''_{i,\Delta}$.

This implies that, by Observation 4.4, $u \leq \Delta - i$. Let C be the combination of the configurations generated by C_1 and C_2 w.r.t. u and the identity function σ .

We start by handling a special case, that is, when $a_1 = b_2 = u$ or $a_2 = b_1 = u$. W.l.o.g., let $a_1 = b_2 = u$. Consider the configuration C^* generated by the condensed configuration $L_1 \dots L_\Delta$, where:

- $L_j = L_{2,j}$ for all $j < u$;
- $L_j = L_{1,j}$ for all $j > u$;
- $L_u = [\langle (y_{2,u-1} \ y_{1,u})_u \rangle]$.

We first prove that C is dominated by C^* , and then that C^* is contained in $\mathcal{N}''_{i,\Delta}$. For each $j \neq u$, it holds that the set $\langle \ell \mid \ell \in L_{1,j} \rangle \cap \langle \ell \mid \ell \in L_{2,j} \rangle$ is a subset of both $\langle \ell \mid \ell \in L_{1,j} \rangle$ and $\langle \ell \mid \ell \in L_{2,j} \rangle$. In position u of C_1 there is the disjunction $[\langle (-y_{1,u})_u \rangle]$, and in position u of C_2 there is the disjunction $[\langle (y_{2,u-1})_u \rangle]$. Observe that L_u is a superset of both (see Figure 6). Also, by definition, the obtained configuration C^* is generated by a strikethrough configuration with parameters a_2 and b_1 , and hence it is part of $\mathcal{N}''_{i,\Delta}$.

In the other cases, we proceed as follows. For $1 \leq h \leq \Delta - i - 1$, let

$$T_h = (T_{h,1}, T_{h,2}) = (f(C_1, C_2, h, u, 1), f(C_1, C_2, h, u, 2)).$$

We now consider all possible cases for the values of $T_1, \dots, T_{\Delta-i-1}$.

- **There exists an index $1 \leq h \leq \Delta - i - 1$ such that $T_{h,2} \neq -$ and $(h = \Delta - i - 1 \vee T_{h+1,1} \neq -)$.** Let a be the largest integer satisfying $T_a \in \{(-, 0), (-, 1)\}$, $a \leq h$, ($u \geq h \vee u < a$), if it exists, and \perp otherwise. Let b be the smallest integer satisfying $T_b \in \{(0, -), (1, -)\}$, $b \geq h + 1$, ($u \leq h + 1 \vee u > b$), if it exists, and \perp otherwise. In the following, we assume that, if there exists at least one index h that satisfies the requirements and that makes a and b (which are defined as a function of h) both different from \perp , then h is one of such values.

If a and b are both different from \perp , consider the configuration C^* generated by $L_1 \dots L_\Delta$, where:

- $L_j = [\langle (--)_j \rangle]$ if $j < a$;
- $L_j = [\langle (T_{j,1} \ T_{j,2})_j \rangle]$ if $a \leq j \leq b$;
- $L_j = [\langle (--)_j \rangle]$ if $b < j \leq \Delta - i - 1$;
- $L_{\Delta-i} = [\langle (00)_{\Delta-i}, (01)_{\Delta-i}, (10)_{\Delta-i}, (11)_{\Delta-i} \rangle]$;
- $L_j = [\langle (?)_j \rangle]$ if $j > \Delta - i$.

We start by arguing that u must satisfy $u < a$ or $u > b$. For a contradiction, assume otherwise. By the definition of a and b , this implies $u \in \{h, h + 1\}$. If $u = h$, then the assumption $T_{h,2} \neq -$ implies $T_u \neq (-, -)$. If $u = h + 1$, then the assumption $T_b \in \{(0, -), (1, -)\}$ implies $b \leq \Delta - i - 1$ (as otherwise T_b would not be defined), which, combined with the assumption $b \geq h + 1$ implies $h < \Delta - i - 1$, which implies $T_{h+1,1} \neq -$, and hence $T_u \neq (-, -)$. In either case, we obtain $T_u \neq (-, -)$, contradicting the definition of T_u , and we conclude that $u < a$ or $u > b$.

This in particular implies that for all $a < h < b$ it holds that $T_{h,1}$ and $T_{h,2}$ are both different from $-$, and we get that C^* is generated by a strikethrough configuration with parameters a, b (where we use Observation 4.29 in both arguments). Hence, C^* is contained in $\mathcal{N}''_{i,\Delta}$. Moreover, C^* dominates C , in fact:

- In all positions j satisfying $j < a$ or $j > b$, C^* contains all sets of color j . Hence, in such positions the set of C^* is a superset of the set of C . Moreover, observe that u satisfies $u < a$ or $u > b$ (as observed above).
- In all positions j satisfying $a \leq j \leq b$, we apply Lemma 4.27, obtaining that

$$\langle\langle (T_{j,1} T_{j,2})_j \rangle\rangle \supseteq \langle \ell \mid \ell \in L_{1,j} \rangle \cap \langle \ell \mid \ell \in L_{2,j} \rangle.$$

For similar reasons, if $a \neq \perp$, $u < a$, and $b = \perp$, C is dominated by the configuration generated by $L_1 \dots L_\Delta$, where:

- $L_j = [\langle\langle (-) \rangle\rangle_j]$ if $j < a$;
- $L_j = [\langle\langle (T_{j,1} T_{j,2})_j \rangle\rangle]$ if $a \leq j \leq \Delta - i - 1$;
- $L_j = [\langle\langle (T_{j-1,2} 0)_j, (T_{j-1,2} 1)_j \rangle\rangle]$ if $j = \Delta - i$;
- $L_j = [\langle\langle (?) \rangle\rangle_j]$ if $j > \Delta - i$.

Note that this configuration is contained in $\mathcal{N}''_{i,\Delta}$.

Now consider the case that $a = \perp$. Recall that, by assumption, it holds that there exists an index h such that $T_{h,2} \neq -$ and $(h = \Delta - i - 1 \vee T_{h+1,1} \neq -)$. By recursively applying Observation 4.29 starting with h and decreasing it step by step, since $a = \perp$ and since $T_{1,1} = -$, we get that $u < h$, and that for all $u < j \leq h$, $T_{j,1}$ and $T_{j,2}$ are both different from $-$ (since otherwise we would have found a value for a).

Similarly, if $b = \perp$ and $u > h$, it must also hold that for all $h + 1 \leq j < \min\{u, \Delta - i - 1\}$, $T_{j,1}$ and $T_{j,2}$ are both different from $-$. Moreover, the case $b = \perp$, $a \neq \perp$ and $u < a$ has already been handled, and, by the definition of a , it is not possible that $b = \perp$, $a \neq \perp$, and $a \leq u \leq h$ (where we use that $u \neq h$ since $T_{h,2} \neq -$ and $T_u = (-, -)$).

Hence, we need to handle the following two cases: the case $a = \perp$ and $u < h$, and the case $b = \perp$, $a \neq \perp$, and $u > h$.

Let us consider the case $a = \perp$ and $u < h$. Since $T_{u+1,1} \neq -$, by Lemma 4.28, it must hold that either $\{t_{1,u,1}, t_{2,u,1}\} = \{0, 1\}$ or $\{t_{1,u}, t_{2,u}\} = \{(x_1, -), (-, x_2)\}$ for some $x_1, x_2 \in \{0, 1\}$. The latter case is covered by the special case handled before, i.e., when $a_1 = b_2 = u$ or $a_2 = b_1 = u$.

Hence, consider the former case, which implies that $u \neq 1$ and that $T_{u-1,2} = -$. If there exists an index $h' < u$ such that $T_{h',2} \neq -$, which by Observation 4.29 implies $T_{h'+1,1} \neq -$, then, since $T_{u-1,2} = -$, we can find two indices $a \leq h'$ and $b \geq h' + 1$ satisfying the requirements stated at the beginning of this case, contradicting the assumption that $a = \perp$. Hence, $T_j = (-, -)$ for all $j \leq u$, which implies that $a_1 = a_2 < u$, because:

- if $a_1 > u$ or $a_2 > u$, then we get that the union is on comparable sets, which contradicts Observation 4.4;
- if $a_1 = u$ or $a_2 = u$, w.l.o.g. $a_1 = u$, then $t_{1,u,1} = t_{1,a_1,1} = -$, which contradicts $t_{1,u,1} \in \{0, 1\}$;
- if $a_1 < u$, $a_2 < u$, and $a_1 \neq a_2$, w.l.o.g. $a_1 < a_2$, then it must hold that $T_{a_2} \neq (-, -)$, which is a contradiction.

Therefore, in position $a_1 = a_2$, since we must have $\{L_{1,a_1}, L_{2,a_2}\} = \{\langle\langle (-0)_{a_1} \rangle\rangle, \langle\langle (-1)_{a_1} \rangle\rangle\}$ (otherwise we would not have $T_{a_1} = (-, -)$), we get $\langle \ell \mid \ell \in L_{1,a_1} \rangle \cap \langle \ell \mid \ell \in L_{2,a_2} \rangle = \langle\langle (!!)_{a_1} \rangle\rangle$, by Lemmas 4.21 and 4.22. Hence the result is dominated by a configuration generated by a grabbing-present configuration.

Let us now consider the case $b = \perp$, $a \neq \perp$, and $u > h$. If $u \leq \Delta - i - 1$ and $T_{\Delta-i-1,2} = -$, we consider two cases:

- $t_{1,\Delta-i-1,2} = t_{2,\Delta-i-1,2} = -$, i.e., the output bit in position $\Delta - i - 1$ is $-$ in both configurations. In this case, the claim holds for symmetric reasons as in the previous case (i.e., the case $a = \perp$ and $u < h$).
- $\{t_{1,\Delta-i-1,2} = t_{2,\Delta-i-1,2}\} = \{0, 1\}$. In this case, the output is dominated by the configuration generated by the grabbing-special configuration, because $\{L_{1,\Delta-i}, L_{2,\Delta-i}\} = \{[(\langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle), [(\langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle)]\}$.

We now consider the case that $u \leq \Delta - i - 1$ and $T_{\Delta-i-1,2} \neq -$. Since $b = \perp$ and $u > h$, we obtain that $T_{u-1,2} \neq -$. By Lemma 4.28, this implies that $\{t_{1,u,2}, t_{2,u,2}\} = \{0, 1\}$, or $\{t_{1,u}, t_{2,u}\} = \{(x_1, -), (-, x_2)\}$ for some $x_1, x_2 \in \{0, 1\}$. The latter case is covered by the special case handled before, i.e., when $a_1 = b_2 = u$ or $a_2 = b_1 = u$. In the former case, we obtain that $u < \Delta - i - 1$ (because $T_{\Delta-i-1,2} \neq -$) and that $T_{u+1,1} = -$. We get that we could have picked a different value of h that would have made $a \neq \perp$ and $u < a$, which is a case previously covered. For the case $u = \Delta - i$, by applying Observation 4.4, we get that $\{L_{1,\Delta-i}, L_{2,\Delta-i}\} = \{[(\langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle), [(\langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle)]\}$ and hence that $\{t_{1,u-1,2}, t_{2,u-1,2}\} = \{0, 1\}$. This implies that $T_{u-1,2} = -$, and hence b cannot be \perp , a contradiction. Finally, we can exclude the case $u > \Delta - i$ by Observation 4.4.

- **For all $1 \leq h \leq \Delta - i - 1$, $T_h = (-, -)$.** We consider positions a_1 and b_1 . Recall that $a_1 < b_1$. Either $a_1 \neq u$ or $b_1 \neq u$. In the former case, in order to obtain $T_h = (-, -)$ it must hold that $\{t_{1,a_1}, t_{2,a_1}\} \in \{(-, 0), (-, 1)\}$ and that $t_{1,a_1} \neq t_{2,a_1}$. In this case, the result is dominated by the configuration generated by a grabbing-present configuration (i.e., the one with label $[(\langle(!)_{a_1}\rangle)]$). In the latter case, we distinguish two cases. If $b_1 < \Delta - i$, it must hold that $\{t_{1,b_1}, t_{2,b_1}\} \in \{(0, -), (1, -)\}$ and that $t_{1,b_1} \neq t_{2,b_1}$. As before, the result is dominated by the configuration generated by a grabbing-present configuration (i.e., the one with label $[(\langle(!)_{b_1}\rangle)]$). If $b_1 = \Delta - i$, since it holds that $T_{\Delta-i-1,2} = -$, we get that $\{L_{1,\Delta-i}, L_{2,\Delta-i}\} = \{[(\langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle), [(\langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle)]\}$, and hence the result is dominated by the configuration generated by the grabbing-special configuration.
- **There exists an index $1 \leq h \leq \Delta - i - 1$ such that $T_h = (-, b)$ and $T_{h+1} = (-, -)$, or such that $T_h = (b, -)$ and $T_{h-1} = (-, -)$, for some $b \in \{0, 1\}$, and the previous cases do not apply.** We consider positions a_1 and b_1 . Recall that $a_1 < b_1$. Either $a_1 \neq u$ or $b_1 \neq u$. In the former case, if $\{t_{1,a_1}, t_{2,a_1}\} \in \{(-, 0), (-, 1)\}$ and $t_{1,a_1} \neq t_{2,a_1}$, the result is dominated by the configuration generated by a grabbing-present configuration. If $t_{1,a_1} = t_{2,a_1}$, w.l.o.g. assume $t_{1,a_1} = t_{2,a_1} = (-, 0)$, by Observation 4.29, in order to have $T_{a_1+1} = (-, -)$, it must hold that $u = a_1 + 1$. By Observation 4.4, $L_{1,u}$ and $L_{2,u}$ must be incomparable, which implies that $\{t_{1,u}, t_{2,u}\} \in \{(0, 0), (0, 1)\}$. If $b_1 = u + 1$, we get that $\{t_{1,b_1}, t_{2,b_1}\} \in \{(0, -), (1, -)\}$, and hence that the result is dominated by the configuration generated by the grabbing-special configuration. If $b_1 > u + 1$ and $\{t_{1,b_1}, t_{2,b_1}\} \in \{(0, -), (1, -)\}$, we again get that the result is dominated by the configuration generated by the grabbing-special configuration. The remaining case to handle is the one where $b_1 > u + 1$ and $t_{1,b_1} = t_{2,b_1}$. We get that $T_{b_1-1,2} \neq -$ and that $T_{b_1,1} \neq -$, which contradicts the assumption that previous cases do not apply. Finally, the case $b_1 \neq u$ is symmetric to the case $a_1 \neq u$, which we have already handled.
- **There exists an index $1 \leq h \leq \Delta - i - 1$ such that $T_h = (b_1, b_2)$ such that $b_h, b_j \in \{0, 1\}$ and for all $j \neq h$, it holds that $T_j = (-, -)$.** This case cannot happen by Observation 4.29.

We obtained that, in all cases, C is dominated by some configuration in $\mathcal{N}_{i,\Delta}''$, which is a contradiction. \square

Black constraint. We define the black constraint $\mathcal{E}_{i,\Delta}''$ as follows. Let $\Sigma_{i,\Delta}''$ be the set of labels that appear in $\mathcal{N}_{i,\Delta}''$. The black constraint $\mathcal{E}_{i,\Delta}''$ contains all the configurations described by the set of condensed configurations given by the following process. Take a configuration C^* in $\mathcal{E}'_{i,\Delta}$, and replace each label L with the disjunction containing the labels of $\Sigma_{i,\Delta}''$ that are supersets of (or equal to) $\langle L \rangle$. By the definition of $\mathcal{E}_{i,\Delta}''$, we obtain the following.

Lemma 4.31. $\mathcal{E}_{i,\Delta}''$ is the set of all configurations $S_1 S_2 S_3$ of sets contained in $\Sigma_{i,\Delta}''$ such that there exists some tuple $(\ell_1, \ell_2, \ell_3) \in S_1 \times S_2 \times S_3$ satisfying that $\ell_1 \ell_2 \ell_3$ is contained in $\mathcal{E}'_{i,\Delta}$.

By combining Lemma 4.30 and Lemma 4.31, by the definition of relaxation, we obtain the following.

Lemma 4.32. The problem $\Pi_{i,\Delta}''$ is a relaxation of $\overline{\mathcal{R}}(\Pi'_{i,\Delta})$.

More explicit black constraint. While for the white constraint we have an explicit definition (given by Observation 4.26), the definition of $\mathcal{E}_{i,\Delta}''$ is given implicitly. We now provide an explicit definition of $\mathcal{E}_{i,\Delta}''$. First, we explicitly list the labels appearing in $\mathcal{N}_{i,\Delta}''$, and hence the labels in $\Sigma_{i,\Delta}''$.

- For color 1, there are the following labels: $\langle\langle(!)_1\rangle\rangle, \langle\langle(-)_1\rangle\rangle, \langle\langle(-0)_1\rangle\rangle, \langle\langle(-1)_1\rangle\rangle$.
- For each present color j satisfying $2 \leq j \leq \Delta - i - 1$, there are the following labels: $\langle\langle(!)_j\rangle\rangle, \langle\langle(-)_j\rangle\rangle, \langle\langle(-0)_j\rangle\rangle, \langle\langle(-1)_j\rangle\rangle, \langle\langle(0)_j\rangle\rangle, \langle\langle(1)_j\rangle\rangle, \langle\langle(00)_j\rangle\rangle, \langle\langle(01)_j\rangle\rangle, \langle\langle(10)_j\rangle\rangle, \langle\langle(11)_j\rangle\rangle$.
- For each gone color j , there are the following labels: $\langle\langle(?)_j\rangle\rangle, \langle\langle(!)_j\rangle\rangle$.
- For the special color $j = \Delta - i$, there are the following labels: $\langle\langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle\rangle, \langle\langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle\rangle, \langle\langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle\rangle, \langle\langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle\rangle, \langle\langle(-0)_{\Delta-i}, (-1)_{\Delta-i}\rangle\rangle$.

Now, by Observation 4.25 (see Figures 5 to 8), $\mathcal{E}_{i,\Delta}''$ can be obtained as follows. Take each configuration C^* in $\mathcal{E}'_{i,\Delta}$, and replace labels as follows. For gone colors j :

- $\langle\langle(!)_j\rangle\rangle \mapsto [\langle\langle(?)_j\rangle\rangle, \langle\langle(!)_j\rangle\rangle];$
- $\langle\langle(?)_j\rangle\rangle \mapsto [\langle\langle(?)_j\rangle\rangle].$

For color 1:

- $\langle\langle(!)_1\rangle\rangle \mapsto [\langle\langle(!)_1\rangle\rangle, \langle\langle(-)_1\rangle\rangle, \langle\langle(-0)_1\rangle\rangle, \langle\langle(-1)_1\rangle\rangle];$
- $\langle\langle(-)_1\rangle\rangle \mapsto [\langle\langle(-)_1\rangle\rangle];$
- For $b \in \{0, 1\}$, $\langle\langle(-b)_1\rangle\rangle \mapsto [\langle\langle(-b)_1\rangle\rangle, \langle\langle(-)_1\rangle\rangle].$

For each present color j satisfying $2 \leq j \leq \Delta - i - 1$:

- For $b_1 \in \{0, 1\}, b_2 \in \{0, 1\}$, $\langle\langle(b_1 b_2)_j\rangle\rangle \mapsto [\langle\langle(-)_j\rangle\rangle, \langle\langle(b_1 b_2)_j\rangle\rangle];$
- For $b \in \{0, 1\}$, $\langle\langle(b)_j\rangle\rangle \mapsto [\langle\langle(-)_j\rangle\rangle, \langle\langle(b 0)_j\rangle\rangle, \langle\langle(b 1)_j\rangle\rangle, \langle\langle(b)_j\rangle\rangle];$
- For $b \in \{0, 1\}$, $\langle\langle(-b)_j\rangle\rangle \mapsto [\langle\langle(-)_j\rangle\rangle, \langle\langle(0 b)_j\rangle\rangle, \langle\langle(1 b)_j\rangle\rangle, \langle\langle(-b)_j\rangle\rangle];$

- $\langle\langle 01 \rangle\rangle_j, \langle\langle 10 \rangle\rangle_j \mapsto [\langle\langle\langle 01 \rangle\rangle_j\rangle, \langle\langle\langle 10 \rangle\rangle_j\rangle, \langle\langle\langle -- \rangle\rangle_j\rangle];$
- $\langle\langle 00 \rangle\rangle_j, \langle\langle 11 \rangle\rangle_j \mapsto [\langle\langle\langle 00 \rangle\rangle_j\rangle, \langle\langle\langle 11 \rangle\rangle_j\rangle, \langle\langle\langle -- \rangle\rangle_j\rangle];$
- For $b_1 \in \{0, 1\}, b_2 \in \{0, 1\}, \langle\langle b_1 - \rangle\rangle_j, \langle\langle -b_2 \rangle\rangle_j \mapsto [\langle\langle\langle -- \rangle\rangle_j\rangle, \langle\langle\langle b_1 - \rangle\rangle_j\rangle, \langle\langle\langle -b_2 \rangle\rangle_j\rangle, \langle\langle\langle b_1 0 \rangle\rangle_j\rangle, \langle\langle\langle b_1 1 \rangle\rangle_j\rangle, \langle\langle\langle 0 b_2 \rangle\rangle_j\rangle, \langle\langle\langle 1 b_2 \rangle\rangle_j\rangle];$
- $\langle\langle !! \rangle\rangle_j \mapsto [\langle\langle\langle !! \rangle\rangle_j\rangle, \langle\langle\langle -- \rangle\rangle_j\rangle, \langle\langle\langle -0 \rangle\rangle_j\rangle, \langle\langle\langle -1 \rangle\rangle_j\rangle, \langle\langle\langle 0- \rangle\rangle_j\rangle, \langle\langle\langle 1- \rangle\rangle_j\rangle, \langle\langle\langle 00 \rangle\rangle_j\rangle, \langle\langle\langle 01 \rangle\rangle_j\rangle, \langle\langle\langle 10 \rangle\rangle_j\rangle, \langle\langle\langle 11 \rangle\rangle_j\rangle].$

For the special color $j = \Delta - i$:

- $\langle\langle 00 \rangle\rangle_j, \langle\langle 01 \rangle\rangle_j \mapsto [\langle\langle\langle 00 \rangle\rangle_j, \langle\langle 01 \rangle\rangle_j\rangle, \langle\langle\langle 00 \rangle\rangle_{\Delta-i}, \langle\langle 01 \rangle\rangle_{\Delta-i}, \langle\langle 10 \rangle\rangle_{\Delta-i}, \langle\langle 11 \rangle\rangle_{\Delta-i}\rangle];$
- $\langle\langle 10 \rangle\rangle_j, \langle\langle 11 \rangle\rangle_j \mapsto [\langle\langle\langle 10 \rangle\rangle_j, \langle\langle 11 \rangle\rangle_j\rangle, \langle\langle\langle 00 \rangle\rangle_{\Delta-i}, \langle\langle 01 \rangle\rangle_{\Delta-i}, \langle\langle 10 \rangle\rangle_{\Delta-i}, \langle\langle 11 \rangle\rangle_{\Delta-i}\rangle];$
- $\langle\langle -0 \rangle\rangle_j, \langle\langle -1 \rangle\rangle_j \mapsto [\langle\langle\langle 00 \rangle\rangle_{\Delta-i}, \langle\langle 01 \rangle\rangle_{\Delta-i}\rangle, \langle\langle\langle 10 \rangle\rangle_{\Delta-i}, \langle\langle 11 \rangle\rangle_{\Delta-i}\rangle, \langle\langle\langle -0 \rangle\rangle_{\Delta-i}, \langle\langle -1 \rangle\rangle_{\Delta-i}\rangle],$

Let B be the defined mapping. We summarize the above observation in the following statement.

Observation 4.33. *The black constraint $\mathcal{E}''_{i,\Delta}$ can be obtained as follows. Take each configuration C^* in $\mathcal{E}'_{i,\Delta}$, and replace labels using the mapping B .*

Relation between $\Pi''_{i,\Delta}$ and $\Pi_{i+1,\Delta}$. We define additional problems $\Pi''_{i,\Delta}$ and $\Pi'''_{i,\Delta}$, which are, like $\Pi'_{i,\Delta}$, also going to be relaxations of $\overline{\mathcal{R}}(\Pi'_{i,\Delta})$, and then we prove that $\Pi'''_{i,\Delta}$ is equivalent to $\Pi_{i+1,\Delta}$. Let j be the special color. The problem $\Pi'''_{i,\Delta}$ is obtained by replacing each instance of $\langle\langle\langle 00 \rangle\rangle_j, \langle\langle 01 \rangle\rangle_j\rangle$, and each instance of $\langle\langle\langle 10 \rangle\rangle_j, \langle\langle 11 \rangle\rangle_j\rangle$, with the label $\langle\langle\langle 00 \rangle\rangle_j, \langle\langle 01 \rangle\rangle_j, \langle\langle 10 \rangle\rangle_j, \langle\langle 11 \rangle\rangle_j\rangle$, in both the white and the black constraint of $\Pi''_{i,\Delta}$.

Observation 4.34. *In the white constraint of $\Pi'''_{i,\Delta}$, all configurations that, for $j = \Delta - i - 1$, contain the label $\langle\langle\langle 0- \rangle\rangle_j\rangle$ or $\langle\langle\langle 1- \rangle\rangle_j\rangle$ are non-maximal.*

Proof. Consider the definition of $\mathcal{N}''_{i,\Delta}$ given by Observation 4.26, and the replacement rules used to define $\Pi'''_{i,\Delta}$ as a function of $\Pi''_{i,\Delta}$. The proof follows by observing the following:

- Label $\langle\langle\langle 0- \rangle\rangle_j\rangle$ is a subset of both $\langle\langle\langle 00 \rangle\rangle_j\rangle$ and $\langle\langle\langle 01 \rangle\rangle_j\rangle$;
- Label $\langle\langle\langle 1- \rangle\rangle_j\rangle$ is a subset of both $\langle\langle\langle 10 \rangle\rangle_j\rangle$ and $\langle\langle\langle 11 \rangle\rangle_j\rangle$;
- Given a configuration C that in position j has a label in $\{\langle\langle\langle 0- \rangle\rangle_j\rangle, \langle\langle\langle 1- \rangle\rangle_j\rangle\}$, there is also a configuration that is equal to C in all positions except j , and in position j contains a label from $\{\langle\langle\langle 00 \rangle\rangle_j\rangle, \langle\langle\langle 01 \rangle\rangle_j\rangle, \langle\langle\langle 10 \rangle\rangle_j\rangle, \langle\langle\langle 11 \rangle\rangle_j\rangle\}$. \square

Let $\Pi'''_{i,\Delta}$ be the problem obtained by removing from the white constraint of $\Pi'''_{i,\Delta}$ all configurations that, for $j = \Delta - i - 1$, contain the label $\langle\langle\langle 0- \rangle\rangle_j\rangle$ or $\langle\langle\langle 1- \rangle\rangle_j\rangle$. By the definitions of $\Pi'''_{i,\Delta}$ and $\Pi'''_{i,\Delta}$, by Observation 4.34, and by Lemma 4.32, we obtain the following.

Lemma 4.35. *The problem $\Pi'''_{i,\Delta}$ is a relaxation of $\overline{\mathcal{R}}(\Pi'_{i,\Delta})$.*

We now define a renaming N of the labels of $\mathcal{N}'''_{i,\Delta}$. Let N be the following mapping.

- If j is not special, let $N(\langle\langle L \rangle\rangle) = L$.
- If j is special, let:

- $N(\langle\langle\langle(00)_j, (01)_j\rangle\rangle, \langle\langle(10)_j, (11)_j\rangle\rangle) = (?)_j$;
- $N(\langle\langle\langle(-0)_j, (-1)_j\rangle\rangle) = (!)_j$.

By combining the definition of $\mathcal{N}''_{i,\Delta}$ given by Observation 4.26, Observation 4.33, the replacement rules used to define $\Pi''_{i,\Delta}$ as a function of $\Pi''_{i,\Delta}$, and the definition of $\mathcal{N}''''_{i,\Delta}$ as a function of $\Pi''_{i,\Delta}$, we obtain the following.

Observation 4.36. *Under the renaming N , the white constraint $\mathcal{N}''''_{i,\Delta}$ of $\Pi''''_{i,\Delta}$ is equal to $\mathcal{N}_{i+1,\Delta}$.*

Observation 4.37. *Under the renaming N , the black constraint $\mathcal{E}''''_{i,\Delta}$ of $\Pi''''_{i,\Delta}$ is equal to $\mathcal{E}_{i+1,\Delta}$.*

By combining Lemma 4.19, Lemma 4.35, Observation 4.36, and Observation 4.37, we obtain the following.

Lemma 4.38. *For all $0 \leq i \leq \Delta - 3$, there exists a problem $\Pi_{i,\Delta}^{\text{rel}}$ satisfying the following:*

- $\mathcal{R}(\Pi_{i,\Delta})$ can be relaxed to $\Pi_{i,\Delta}^{\text{rel}}$;
- $\overline{\mathcal{R}}(\Pi_{i,\Delta}^{\text{rel}})$ can be relaxed to $\Pi_{i+1,\Delta}$;
- The number of labels of $\Pi_{i,\Delta}^{\text{rel}}$ are upper bounded by $O(\Delta)$.

4.8 The problem $\Pi_{\Delta-2,\Delta}$ is not trivial

We now prove that $\Pi_{\Delta-2,\Delta}$ is not trivial to solve, even when given a Δ -edge coloring.

Lemma 4.39. *The problem $\Pi_{\Delta-2,\Delta}$ is not 0-round solvable in the deterministic port numbering model, even if a Δ -edge coloring is given.*

Proof. Assume for a contradiction that there exists a deterministic 0-round algorithm solving $\Pi_{\Delta-2,\Delta}$ when given a Δ -edge coloring. Any 0-round algorithm (operating on white nodes) in the deterministic port numbering model satisfies that all white nodes output the same configuration $C = \{L_1, \dots, L_\Delta\} \in \mathcal{N}_{\Delta-2,\Delta}$. Moreover, there must exist a permutation φ such that all nodes output label L_i on port $\varphi(i)$.

By the definition of $\mathcal{N}_{\Delta-2,\Delta}$, the configuration C must contain at least one label among the following:

- $(!!)_1$;
- $(!)_j$, for some $3 \leq j \leq \Delta$;
- $(-0)_j$, for some $1 \leq j \leq 2$;
- $(-1)_j$, for some $1 \leq j \leq 2$.

Let L_j be such a label, and let $i = \varphi(j)$ be the port on which the algorithm outputs L_j .

Observe that, for any possible L_j listed above, the configuration $L_j L_j L_j$ is not contained in $\mathcal{E}_{\Delta-2,\Delta}$. However, by considering a graph in which there are three white nodes u, v, w connected via port i to the same black node b (satisfying that the Δ -edge coloring assigns color i to node b), we obtain a setting in which the algorithm fails to produce a valid solution for $\Pi_{\Delta-2,\Delta}$. \square

4.9 Relation between \mathcal{P}_Δ and $\Pi_{0,\Delta}$

Lemma 4.40. *Given a solution for \mathcal{P}_Δ , it is possible to solve $\Pi_{0,\Delta}$ in 0 rounds.*

Proof. It is easy to see that $\mathcal{N}_{\mathcal{P}_\Delta} \subsetneq \mathcal{N}_{0,\Delta}$ and that $\mathcal{E}_{\mathcal{P}_\Delta} \subsetneq \mathcal{E}_{0,\Delta}$. Hence, any solution for \mathcal{P}_Δ is also a solution for $\Pi_{0,\Delta}$. \square

4.10 Putting things together

By applying Lemma 4.38 multiple times, and then applying Lemma 4.39, we obtain the following.

Lemma 4.41. *Let $\Delta \geq 3$ be an integer. There exists a problem sequence $\Pi_{0,\Delta} \rightarrow \Pi_{1,\Delta} \rightarrow \dots \rightarrow \Pi_{\Delta-2,\Delta}$ such that, for all $0 \leq i < \Delta - 2$, the following holds:*

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

Also, $\Pi_{\Delta-2,\Delta}$ is not 0-round solvable in the deterministic port numbering model, when given a Δ -edge coloring.

By combining Lemma 4.41 with Theorem 4.9, we obtain the following.

Lemma 4.42. *Let $\Delta \geq 3$ be an integer. The problem $\Pi_{0,\Delta}$ requires $\Omega(\min\{\Delta, \log_\Delta n\})$ rounds in the deterministic LOCAL model and $\Omega(\min\{\Delta, \log_\Delta \log n\})$ rounds in the randomized LOCAL model, even if a Δ -edge coloring is given.*

By combining Lemma 4.42 with Lemma 4.40, we obtain Theorem 4.1.

5 Networks of non-signaling games

In this section, we show a classical upper bound for all networks of non-signaling games. We start by defining the notion of a game. Each player of a game receives a local input and must produce a local output such that the outputs of all players combined are consistent with their inputs.

Definition 5.1 (Game). Let Σ be a finite set, and let $m \in \mathbb{N}_+$ be the number of players. We call $\mathfrak{g} \subseteq \Sigma^m \times \Sigma^m$ a *game*. Each player i receives one input $x_i \in \Sigma$ and produces one output $y_i \in \Sigma$. A *move* $\mu = (x, y) \in \Sigma^m \times \Sigma^m$ is valid if $\mu \in \mathfrak{g}$. We overload the notation so that $\mathfrak{g}(x) = \{y \in \Sigma^m \mid (x, y) \in \mathfrak{g}\}$. We say \mathfrak{g} is *solvable* if, for every x , $\mathfrak{g}(x)$ is non-empty.

Note that our games last for a single move. (Later on we will connect a player to multiple games and hence each player will be part of moves in multiple games.) Since games that are not solvable cannot be solved by any strategy at all, in what follows we assume that we are dealing exclusively with solvable games.

Let us now move to the setting where a player takes part in multiple games. A game's inputs may depend arbitrarily (but in a computable way) on the outputs of games that were played before. Formally we model this using circuits.

Definition 5.2 (Circuit of half-games). Let $d \in \mathbb{N}_+$, and let D be a circuit that satisfies the following properties:

- D has a single input gate ξ and no output gates.
- Every wire in D carries a value in Σ .
- D may contain arbitrary gates over Σ (of unbounded fan-in and fan-out).
- D contains d unary gates H_1, \dots, H_d , which we call *half-games*.

A circuit C is a (*degree- d*) *circuit of half-games* if it can be obtained from such a D by detaching all the H_i , that is, by replacing every H_i with a pair of input and output gates x_i and y_i (that are connected to the rest of the circuit in the same way the input and output of H_i are). In a sense,




Figure 9: Transforming D (on the left) into C (in the middle) by detaching all half-games. Note how C may have multiple components and not use every one of its inputs. On the right, the prefix C^2 is presented.

C is just D where one can externally examine the inputs $x_1 \cdots x_d$ to half-games H_1, \dots, H_d , and the outputs of half-games are fixed at $y_1 \cdots y_d$. As a result, C has $d + 1$ inputs $\xi, y_1, \dots, y_d \in \Sigma$ and d outputs $x_1, \dots, x_d \in \Sigma$. Equivalently, we may also view C as having 2 inputs, $\xi \in \Sigma$ and $y = y_1 \cdots y_d \in \Sigma^d$, and a single output $x = x_1 \cdots x_d \in \Sigma^d$. See Figure 9 for an example.

The intuition is that, in addition to an external input ξ , D is connecting together d pairs (x_i, y_i) that each correspond to a half-game H_i . Note that the roles of x_i and y_i as input and output from the perspective of D , in contrast to that of the half-game H_i , are reversed: How y_i is obtained from x_i is extrinsic to D —indeed, it is determined by H_i —and hence what D does is describe how x_i (which is the input of H_i) is obtained from the external input ξ and the y_j (which are the outputs of the other half-games).

With this in place, we now turn to the definition of the network of games problem. Each player is modeled as a circuit of half-games of degree d . The input on white nodes describe what the circuit looks like for each white node and what is their local input. The input on black nodes specify which game is played with its neighbors. The output on each edge is the respective player's input and output pair to the game across that edge.

Definition 5.3 (Network of games problem). Fix the number of players $m \in \mathbb{N}_+$ and the number of games $d \in \mathbb{N}_+$ each player plays. Let $\mathcal{G} = \{\mathbf{g}_i \subseteq \Sigma^m \times \Sigma^m\}_i$ be a family of games on Σ .

The *network of games problem* $\text{GAMES}_d[\mathcal{G}]$ is an LCL problem defined on (d, m) -regular graphs as follows:

- Each edge must be labeled with a pair from $\Sigma \times \Sigma$.
- Each white node w is given an input element $\xi \in \Sigma$, a circuit of half-games C , and a permutation σ_w of its neighboring edges. The node satisfies the white constraint if the output labels $(x_1, y_1), \dots, (x_d, y_d) \in \Sigma \times \Sigma$ around it satisfy $C(\xi, y^{\sigma_w}) = x^{\sigma_w}$. This ensures that the outputs of games flow correctly through the circuit to following games.

- Each black node is given as input a game $\mathbf{g} \in \mathcal{G}$ and a permutation σ_b of its neighboring edges. The node satisfies the black constraint if the output labels $(x_1, y_1), \dots, (x_m, y_m) \in \Sigma \times \Sigma$ around it satisfy $y^{\sigma_b} \in \mathbf{g}(x^{\sigma_b})$. This ensures that each game is played correctly.

Note that, for simplicity, the problem above is defined with inputs on nodes. To make this an LCL problem in the sense of Definition 2.1, we can move the node input labels to the adjacent edges and require that all the inputs on the adjacent edges agree (and allow a node's output to be arbitrary if this is not the case). In particular, the permutation of neighboring edges is most naturally encoded by having the neighboring edges carry their indices in the permutation. Moreover, note that both the circuits \mathcal{C} and the games \mathbf{g} can be seen as tables over a finite alphabet, and hence there are only finitely many distinct circuits and games. This ensures that the input and output label sets are finite, as required by Definition 2.1.

We note here that the iterated GHZ problem, as defined in Section 3, is a relaxation of a network of games problem, namely $\text{GAMES}_\Delta[\{\text{SYMM}, \text{GHZ}\}]$. Here SYMM is the symmetry-breaking game, where exactly one of the players must output 1 and the other two output 0, regardless of the input. The easiest way to see that is indeed the case is to notice that the iterated GHZ problem is a $\text{GAMES}_\Delta[\{\text{SYMM}, \text{GHZ}\}]$ instance with the following promises on the input:

- The circuit in each white node is the same, namely one where the first game is the SYMM game and the rest are GHZ games. The circuit is linear and the output of one game flows directly to become the input of the next game.
- The Δ -edge coloring (recall that this is a coloring of the black nodes) gives the order of the games: Black node with color 1 plays the SYMM game while the rest of the black nodes play the GHZ game. Note that the setting is actually even more relaxed, as the algorithm is free to permute the order of the colors.
- The Δ -edge coloring also promises that each white agrees on the order of the games with its adjacent white nodes. In the network of games problems in general, the first half-game of one white node may be played with the last half-game of another white node.

Using similar reasoning, one can also get convinced that the iterated CHSH game, defined in Section 1.3, is a relaxation of a network of games problem, this time $\text{GAMES}_\Delta[\{\text{CHSH}\}]$. Finally, we note that CHSH, SYMM and GHZ are non-signaling games, as defined in the next section (Definition 5.5), and hence our upper bound algorithm works for solving both $\text{GAMES}_\Delta[\{\text{CHSH}\}]$ and $\text{GAMES}_\Delta[\{\text{SYMM}, \text{GHZ}\}]$, and therefore also for solving the iterated CHSH and GHZ games.

5.1 Solving networks of non-signaling games deterministically

An important special class of games are so-called non-signaling games. For a survey on the topic, see [32]. We now give a definition of non-signaling games adapted to our purposes. Before doing that, we introduce some notation and minor definitions. For any function $f : A \rightarrow B$ and any subset $S \subseteq A$, we denote by $f \upharpoonright_S$ the function $g : S \rightarrow B$ such that $f(x) = g(x)$ for all $x \in S$. We now define the notion of a *strategy* for a game.

Definition 5.4 (Strategy). For any fixed game $\mathbf{g} \subseteq \Sigma^m \times \Sigma^m$ of m players, consider a mapping from inputs to probability distributions over outputs $\rho^{(\mathbf{g})} : (x_1, \dots, x_m) \mapsto \{(\lambda_j : [m] \rightarrow \Sigma^m, p_j) : j \in J\}$. Here, $(x_1, \dots, x_m) \in \Sigma^m$ is an input for \mathbf{g} , $\lambda_j : [m] \rightarrow \Sigma^m$ is an assignment of outputs to the $[m]$ players, p_j is the probability that λ_j is the actual outputs of the players, and J is any set used for indexing. We call $\rho^{(\mathbf{g})}$ a *strategy* for \mathbf{g} . Since $\{(\lambda_j : [m] \rightarrow \Sigma^m, p_j) : j \in J\}$ is a probability distribution, $\sum_{j \in J} p_j = 1$ and, for all $j \in J$, $p_j \geq 0$.

For any game \mathbf{g} and any strategy $\rho^{(\mathbf{g})}$ for \mathbf{g} , we say that $\rho^{(\mathbf{g})}$ solves \mathbf{g} with probability at least p if, for any input $(x_1, \dots, x_m) \in \Sigma^m$ to \mathbf{g} , the distribution $\{(\lambda_j : [m] \rightarrow \Sigma^m, p_j) : j \in J\} = \rho^{(\mathbf{g})}((x_1, \dots, x_m))$ satisfies the following property:

$$\sum_{\substack{j \in J: \\ (\lambda_j(1), \dots, \lambda_j(m)) \in \mathbf{g}(x_1, \dots, x_m)}} p_j \geq p.$$

For any subset of players $S = \{i_1, \dots, i_s\} \subseteq [m]$ (with $i_1 < i_2 < \dots < i_s$), let us define the restriction of $\rho^{(\mathbf{g})}$ to S as $\rho_S^{(\mathbf{g})} : (x_{i_1}, \dots, x_{i_s}) \mapsto \{(\lambda_{S,j} : S \rightarrow \Sigma^s, p_{S,j}) : j \in J_S\}$ where

$$p_{S,j} = \sum_{\substack{i \in J: \\ \lambda_{S,j} = \lambda_i \upharpoonright_S}} p_i.$$

Definition 5.5 (Non-signaling game). A game $\mathbf{g} \subseteq \Sigma^m \times \Sigma^m$ is said to be *non-signaling* if there exists a strategy $\rho^{(\mathbf{g})} : (x_1, \dots, x_m) \mapsto \{(y_j : [m] \rightarrow \Sigma^m, p_j) : j \in J\}$ solving \mathbf{g} with probability 1 such that, for every pair of inputs $(x_1^{(1)}, \dots, x_m^{(1)}), (x_1^{(2)}, \dots, x_m^{(2)}) \in \Sigma^m \times \Sigma^m$ to \mathbf{g} , and any subset of the players $S = \{i_1, \dots, i_s\} \subseteq [m]$ (with $i_1 < i_2 < \dots < i_s$), if $x_i^{(1)} = x_i^{(2)}$ for all $i \in S$, then $\rho_S^{(\mathbf{g})}(x_{i_1}^{(1)}, \dots, x_{i_s}^{(1)}) = \rho_S^{(\mathbf{g})}(x_{i_1}^{(2)}, \dots, x_{i_s}^{(2)})$. Also, $\rho^{(\mathbf{g})}$ is said to be a *non-signaling strategy* for \mathbf{g} .

Interestingly, for every non-signaling game, the existence of a non-signaling strategy solving \mathbf{g} guarantees a strong completability property, which we exploit for our major result of this section. Let us now describe *strongly completable games*.

Informally, a game is strongly completable if every move can be “sequentialized” by splitting it into *plies* where each player P_i does the following in their respective ply:

1. P_i sees the inputs and outputs of all players that completed their ply before it as well as its own input x_i .
2. With *only* this information, P_i can produce an output y_i such that, no matter what the rest of the inputs are and in what order the rest of the players complete their plies, the other players are each able to find outputs that make the whole move a valid one using only information described in Item 1.

Note we are assuming that the process determining the order in which players must complete their plies is non-adaptive (i.e., it follows a fixed permutation).

Definition 5.6 (Strongly completable game). A game $\mathbf{g} \subseteq \Sigma^m \times \Sigma^m$ is said to be *strongly completable* if, for every permutation $\sigma : [m] \rightarrow [m]$, we have that

$$\forall x_{\sigma(1)} \in \Sigma : \exists y_{\sigma(1)} \in \Sigma : \dots \forall x_{\sigma(m)} \in \Sigma : \exists y_{\sigma(m)} \in \Sigma : (y_1 \dots y_m) \in \mathbf{g}(x_1 \dots x_m).$$

As mentioned before, non-signaling games are strongly completable.

Lemma 5.7. *Any non-signaling game $\mathbf{g} \subseteq \Sigma^m \times \Sigma^m$ is strongly completable.*

Proof. Since \mathbf{g} is non-signaling, there exists a non-signaling strategy $\rho^{(\mathbf{g})}$ solving \mathbf{g} with probability 1, which implies that all outputs sampled according to $\rho^{(\mathbf{g})}$ are global solutions. Fix any permutation $\sigma : [m] \rightarrow [m]$. Now, for the first player $\sigma(1)$ with input $x_{\sigma(1)}$, sample any output $y_{\sigma(1)}$ from the restriction $\rho_{\{\sigma(1)\}}^{(\mathbf{g})}(x_{\sigma(1)})$. Now we proceed recursively. Let $i > 1$ and suppose that all inputs and outputs for players $\sigma(1), \dots, \sigma(i-1)$ have been fixed. For player $\sigma(i)$ with input $x_{\sigma(i)}$, sample any

output $y_{\sigma(i)}$ using $\rho_{\{\sigma(1), \dots, \sigma(i)\}}^{(\mathfrak{g})}(x_{\sigma(1)}, \dots, x_{\sigma(i)})$ as follows: For any $j \in [m]$, let $p(y_{\sigma(1)}, \dots, y_{\sigma(j)})$ the probability of sampling $(y_{\sigma(1)}, \dots, y_{\sigma(j)})$ from $\rho_{\{\sigma(1), \dots, \sigma(j)\}}^{(\mathfrak{g})}(x_{\sigma(1)}, \dots, x_{\sigma(j)})$. Then, any $y_{\sigma(i)}$ is sampled with probability

$$\frac{p(y_{\sigma(1)}, \dots, y_{\sigma(i)})}{p(y_{\sigma(1)}, \dots, y_{\sigma(i-1)})}.$$

This procedure is well-defined because at any step there is always some output to choose with non-zero probability, due to the non-signaling property of $\rho^{(\mathfrak{g})}$ (which solves \mathfrak{g} with probability 1). This is because, at any step i , the probability of having sampled $(y_{\sigma(1)}, \dots, y_{\sigma(i)})$ is exactly the unconditional probability associated to the output labeling $\lambda_i : \{\sigma(1), \dots, \sigma(i)\} \rightarrow \Sigma^i$ defined by $\lambda_i(\sigma(j)) = y_{\sigma(j)}$: such distribution is given by $\rho_{\{\sigma(1), \dots, \sigma(i)\}}^{(\mathfrak{g})}(x_{\sigma(1)}, \dots, x_{\sigma(i)})$. Let us denote this probability by p_i . We prove this statement by induction. If $i = 1$, the thesis is trivial by construction. Assume $i > 1$ and the thesis to be true for $i - 1$. Let Y_1, \dots, Y_i be the random variables describing the outputs for players $\sigma(1), \dots, \sigma(i)$, respectively. Then, the probability of sampling $(y_{\sigma(1)}, \dots, y_{\sigma(i)})$ can be expressed as follows:

$$\begin{aligned} & \Pr [(Y_1, \dots, Y_i) = (y_{\sigma(1)}, \dots, y_{\sigma(i)})] \\ &= \Pr [(Y_1, \dots, Y_i) = (y_{\sigma(1)}, \dots, y_{\sigma(i)}) \mid (Y_1, \dots, Y_{i-1}) = (y_{\sigma(1)}, \dots, y_{\sigma(i-1)})] \\ & \quad \cdot \Pr [(Y_1, \dots, Y_{i-1}) = (y_{\sigma(1)}, \dots, y_{\sigma(i-1)})] \\ &= \frac{p_i}{p_{i-1}} \cdot p_{i-1} \\ &= p_i, \end{aligned} \tag{1}$$

where Equation (1) holds by construction of the sampling procedure. Finally, we conclude that $\Pr [(Y_1, \dots, Y_m) = (y_{\sigma(1)}, \dots, y_{\sigma(m)})] = p_m > 0$. Since all outputs sampled with non-zero probability according to $\rho^{(\mathfrak{g})}$ are actual solutions, then $(y_{\sigma(1)}, \dots, y_{\sigma(m)})$ is a global solution, completing the lemma. \square

We finally present the main result of this section:

Theorem 5.8. *For any collection of strongly completable games $\mathcal{G} = \{\mathfrak{g}_i \subseteq \Sigma^m \times \Sigma^m\}_i$ and any $d \in \mathbb{N}_+$, there is a deterministic LOCAL algorithm that solves $\text{GAMES}_d[\mathcal{G}]$ in $O(d)$ rounds.*

Theorem 5.8 and Lemma 5.7 imply the following corollary.

Corollary 5.9. *For any collection of non-signaling games $\mathcal{G} = \{\mathfrak{g}_i \subseteq \Sigma^m \times \Sigma^m\}_i$ and any $d \in \mathbb{N}_+$, there is a deterministic LOCAL algorithm that solves $\text{GAMES}_d[\mathcal{G}]$ in $O(d)$ rounds.*

To make the proof simpler, we make the following trivial observation:

Observation 5.10. *Let C be a degree- d circuit of half-games. Without loss of generality, we may assume that half-games H_1, \dots, H_d appear in a topological order in circuit C .*

With the help of this observation, we notice that the half-games of a circuit partition the circuit into connected components. In particular, we can consider the circuit formed by only the first k games.

Definition 5.11 (Prefix of a circuit of half-games). Let C be a degree- d circuit of half-games, and let ξ and y_1, \dots, y_d be the inputs and let x_1, \dots, x_d be the outputs of C . Let $k \in [d]$. Define C^k to be the circuit consisting of only components connected to half-games H_1, \dots, H_k . In particular,

circuit C^k has only inputs ξ and (some subset of) y_1, \dots, y_{k-1} , and fully defines the outputs x_1, \dots, x_k . This is because, by the above observation, the games H_1, \dots, H_k are in a topological order in C . In particular, the input for game H_i depends only on ξ and the outputs of games H_j for $j < i$. See Figure 9 for a visualization.

We can now present our algorithm \mathcal{A} . The algorithm proceeds in d phases, each of which consists of two communication rounds: one where information flows from white to black nodes and another where it does so in the reverse direction. We refer to these as the *white* and *black rounds*, accordingly. Let us now describe what occurs in phase number $i \in [d]$.

1. The white nodes keep track of the results y_1, \dots, y_{i-1} of the half-games they have determined already. In the white phase i , each white node w evaluates $C^i(\xi, y_1, \dots, y_{i-1})$ to get value x_i . Node w then sends this value to the black node representing the said game according to the input permutation.
2. Let $j \in [m]$ denote the number of messages black node b has received by black round i , and let $s \in [m]^j$ be player numbers of the senders those messages in the order of reception, breaking ties arbitrarily among simultaneously-arrived messages. The black node b can map the ports to player numbers using the local input permutation. Let $x \in \Sigma^j$ be the values of those messages, in the same order.

Now the node b now plays game \mathbf{g} using s as the playing order and x as the inputs, producing partial output $y \in \Sigma^j$ for those nodes. Because \mathbf{g} is a strongly completable game, it can pick the outputs y one-by-one in such way that the game is still solvable for every possible future inputs. Moreover, the node plays the game in a deterministic manner in order not to change the values of previously-assigned outputs. This is possible by the definition of strongly completable games, see Definition 5.6. Finally, node b sends the player output y_k to each neighbor k it received a message in this round.

The round complexity of \mathcal{A} is evidently $2d = O(d)$. The white constraint is satisfied by construction. Hence for the correctness we need only argue that the black constraint is satisfied, which we do in the following lemma:

Lemma 5.12. *Algorithm \mathcal{A} is correct, that is, for every black node $b \in B$ with input game \mathbf{g} , the edges around it are labeled with (x_i, y_i) such that $y \in \mathbf{g}(x)$ (where $x = x_1 \cdots x_m$ and $y = y_1 \cdots y_m$).*

Proof. Fix such a black node $b \in B$. First notice that, since the algorithm is deterministic, for a fixed input to the algorithm we will have a fixed order in which b receives inputs x_i from white nodes and, hence, also a fixed order in which it sets the outputs y_i . Hence there is a permutation σ as in Definition 5.6 that corresponds to the order in which b sets the y_i . Having observed this, note that there is no difference (from the point of view of b) if the inputs x_i arrive in separate phases or all in the same one. Hence we can assume that the latter is the case. Having done so, then what b does is exactly what is guaranteed to be possible by Definition 5.6, and thus the claim follows. \square

This concludes the proof of Theorem 5.8.

Acknowledgments

Augusto Modanese is supported by the Helsinki Institute for Information Technology (HIIT). Henrik Lievonen is supported by the Research Council of Finland, Grants 333837 and 359104. Alkida Balliu and Dennis Olivetti are supported by MUR (Italy) Department of Excellence 2023 - 2027. Francesco

d'Amore is supported by MUR FARE 2020 - Project PARECoDi CUP J43C22000970001. Xavier Coiteux-Roy acknowledges funding from the BMW endowment fund and from the Swiss National Science Foundation. François Le Gall is supported by JSPS KAKENHI grants Nos. JP20H05966, 20H00579, 24H00071, MEXT Q-LEAP grant No. JPMXS0120319794 and JST CREST grant No. JP-MJCR24I4. Lucas Tendick, Marc-Olivier Renou and Isadora Veeren acknowledge funding from INRIA and CIEDS through the Action Exploratoire project DEPARTURE, and from the ANR for the JCJC grant LINKS (ANR-23-CE47-0003).

References

- [1] Amirreza Akbari, Xavier Coiteux-Roy, Francesco D'Amore, François Le Gall, Henrik Lievonen, Darya Melnyk, Augusto Modanese, Shreyas Pai, Marc-Olivier Renou, Václav Rozhon, and Jukka Suomela. Online locality meets distributed quantum computing. *CoRR*, abs/2403.01903, 2024. [arXiv:2403.01903](https://arxiv.org/abs/2403.01903), [doi:10.48550/ARXIV.2403.01903](https://doi.org/10.48550/ARXIV.2403.01903).
- [2] Amirreza Akbari, Navid Eslami, Henrik Lievonen, Darya Melnyk, Joonas Särkijärvi, and Jukka Suomela. Locality in online, dynamic, sequential, and distributed graph algorithms. In Kousha Etessami, Uriel Feige, and Gabriele Puppis, editors, *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023, July 10-14, 2023, Paderborn, Germany*, volume 261 of *LIPICs*, pages 10:1–10:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. [doi:10.4230/LIPICs.ICALP.2023.10](https://doi.org/10.4230/LIPICs.ICALP.2023.10).
- [3] Heger Arfaoui and Pierre Fraigniaud. What can be computed without communications? *SIGACT News*, 45(3):82–104, 2014. [doi:10.1145/2670418.2670440](https://doi.org/10.1145/2670418.2670440).
- [4] Alkida Balliu, Thomas Boudier, Sebastian Brandt, and Dennis Olivetti. Tight lower bounds in the supported LOCAL model. In Ran Gelles, Dennis Olivetti, and Petr Kuznetsov, editors, *Proceedings of the 43rd ACM Symposium on Principles of Distributed Computing, PODC 2024, Nantes, France, June 17-21, 2024*, pages 95–105. ACM, 2024. [doi:10.1145/3662158.3662798](https://doi.org/10.1145/3662158.3662798).
- [5] Alkida Balliu, Sebastian Brandt, Yi-Jun Chang, Dennis Olivetti, Mikael Rabie, and Jukka Suomela. The distributed complexity of locally checkable problems on paths is decidable. In Peter Robinson and Faith Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 262–271. ACM, 2019. [doi:10.1145/3293611.3331606](https://doi.org/10.1145/3293611.3331606).
- [6] Alkida Balliu, Sebastian Brandt, Yi-Jun Chang, Dennis Olivetti, Jan Studený, and Jukka Suomela. Efficient classification of locally checkable problems in regular trees. In Christian Scheideler, editor, *36th International Symposium on Distributed Computing, DISC 2022, October 25-27, 2022, Augusta, Georgia, USA*, volume 246 of *LIPICs*, pages 8:1–8:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. [doi:10.4230/LIPICs.DISC.2022.8](https://doi.org/10.4230/LIPICs.DISC.2022.8).
- [7] Alkida Balliu, Sebastian Brandt, Yuval Efron, Juho Hirvonen, Yannic Maus, Dennis Olivetti, and Jukka Suomela. Classification of distributed binary labeling problems. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing, DISC 2020, October 12-16, 2020, Virtual Conference*, volume 179 of *LIPICs*, pages 17:1–17:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. [doi:10.4230/LIPICs.DISC.2020.17](https://doi.org/10.4230/LIPICs.DISC.2020.17).
- [8] Alkida Balliu, Sebastian Brandt, Manuela Fischer, Rustam Latypov, Yannic Maus, Dennis Olivetti, and Jara Uitto. Exponential speedup over locality in MPC with optimal memory. In

- Christian Scheideler, editor, *36th International Symposium on Distributed Computing, DISC 2022, October 25-27, 2022, Augusta, Georgia, USA*, volume 246 of *LIPICs*, pages 9:1–9:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.DISC.2022.9.
- [9] Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. Lower bounds for maximal matchings and maximal independent sets. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 481–497. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00037.
- [10] Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. Lower bounds for maximal matchings and maximal independent sets. *Journal of the ACM*, 68(5):39:1–39:30, 2021. doi:10.1145/3461458.
- [11] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. Improved distributed lower bounds for MIS and bounded (out-)degree dominating sets in trees. In Avery Miller, Keren Censor-Hillel, and Janne H. Korhonen, editors, *PODC '21: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, July 26-30, 2021*, pages 283–293. ACM, 2021. doi:10.1145/3465084.3467901.
- [12] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. Distributed Δ -coloring plays hide-and-seek. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 464–477. ACM, 2022. doi:10.1145/3519935.3520027.
- [13] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. Distributed maximal matching and maximal independent set on hypergraphs. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 2632–2676. SIAM, 2023. doi:10.1137/1.9781611977554.CH100.
- [14] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, Dennis Olivetti, and Joonatan Saarhelo. Towards fully automatic distributed lower bounds, 2024. URL: <https://arxiv.org/abs/2410.20224>, arXiv: 2410.20224.
- [15] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, Dennis Olivetti, and Gustav Schmid. On the node-averaged complexity of locally checkable problems on trees. In Rotem Oshman, editor, *37th International Symposium on Distributed Computing, DISC 2023, October 10-12, 2023, L'Aquila, Italy*, volume 281 of *LIPICs*, pages 7:1–7:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.DISC.2023.7.
- [16] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, Dennis Olivetti, and Gustav Schmid. Completing the node-averaged complexity landscape of lcls on trees. In Ran Gelles, Dennis Olivetti, and Petr Kuznetsov, editors, *Proceedings of the 43rd ACM Symposium on Principles of Distributed Computing, PODC 2024, Nantes, France, June 17-21, 2024*, pages 369–379. ACM, 2024. doi:10.1145/3662158.3662773.
- [17] Alkida Balliu, Sebastian Brandt, and Dennis Olivetti. Distributed lower bounds for ruling sets. *SIAM Journal on Computing*, 51(1):70–115, 2022. doi:10.1137/20M1381770.

- [18] Alkida Balliu, Sebastian Brandt, Dennis Olivetti, Jan Studený, Jukka Suomela, and Aleksandr Tereshchenko. Locally checkable problems in rooted trees. In Avery Miller, Keren Censor-Hillel, and Janne H. Korhonen, editors, *PODC '21: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, July 26-30, 2021*, pages 263–272. ACM, 2021. doi:10.1145/3465084.3467934.
- [19] Alkida Balliu, Sebastian Brandt, Dennis Olivetti, and Jukka Suomela. How much does randomness help with locally checkable problems? In Yuval Emek and Christian Cachin, editors, *PODC '20: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, August 3-7, 2020*, pages 299–308. ACM, 2020. doi:10.1145/3382734.3405715.
- [20] Alkida Balliu, Sebastian Brandt, Dennis Olivetti, and Jukka Suomela. Almost global problems in the LOCAL model. *Distributed Computing*, 34(4):259–281, 2021. doi:10.1007/S00446-020-00375-2.
- [21] Alkida Balliu, Keren Censor-Hillel, Yannic Maus, Dennis Olivetti, and Jukka Suomela. Locally checkable labelings with small messages. In Seth Gilbert, editor, *35th International Symposium on Distributed Computing, DISC 2021, October 4-8, 2021, Freiburg, Germany (Virtual Conference)*, volume 209 of *LIPICs*, pages 8:1–8:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.DISC.2021.8.
- [22] Alkida Balliu, Juho Hirvonen, Janne H. Korhonen, Tuomo Lempiäinen, Dennis Olivetti, and Jukka Suomela. New classes of distributed time complexity. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 1307–1318. ACM, 2018. doi:10.1145/3188745.3188860.
- [23] Alkida Balliu, Juho Hirvonen, Dennis Olivetti, and Jukka Suomela. Hardness of minimal symmetry breaking in distributed computing. In Peter Robinson and Faith Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 369–378. ACM, 2019. doi:10.1145/3293611.3331605.
- [24] Alkida Balliu, Janne H. Korhonen, Fabian Kuhn, Henrik Lievonen, Dennis Olivetti, Shreyas Pai, Ami Paz, Joel Rybicki, Stefan Schmid, Jan Studený, Jukka Suomela, and Jara Uitto. Sinkless orientation made simple. In Telikepalli Kavitha and Kurt Mehlhorn, editors, *2023 Symposium on Simplicity in Algorithms, SOSA 2023, Florence, Italy, January 23-25, 2023*, pages 175–191. SIAM, 2023. doi:10.1137/1.9781611977585.CH17.
- [25] Jonathan Barrett, Carlton M. Caves, Bryan Eastin, Matthew B. Elliott, and Stefano Pironio. Modeling Pauli measurements on graph states with nearest-neighbor classical communication. *Physical Review A*, 75(1), 2007. doi:10.1103/physreva.75.012103.
- [26] Sebastian Brandt. An automatic speedup theorem for distributed problems. In Peter Robinson and Faith Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 379–388. ACM, 2019. doi:10.1145/3293611.3331611.
- [27] 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 Daniel Wicks and Yishay Mansour, editors, *Proceedings of the 48th Annual ACM SIGACT*

- Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 479–488. ACM, 2016. doi:[10.1145/2897518.2897570](https://doi.org/10.1145/2897518.2897570).
- [28] Sebastian Brandt, Juho Hirvonen, Janne H. Korhonen, Tuomo Lempiäinen, Patric R. J. Östergård, Christopher Purcell, Joel Rybicki, Jukka Suomela, and Przemyslaw Uznanski. LCL problems on grids. In Elad Michael Schiller and Alexander A. Schwarzmann, editors, *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017, Washington, DC, USA, July 25-27, 2017*, pages 101–110. ACM, 2017. doi:[10.1145/3087801.3087833](https://doi.org/10.1145/3087801.3087833).
- [29] Sebastian Brandt and Dennis Olivetti. Truly tight-in- Δ bounds for bipartite maximal matching and variants. In Yuval Emek and Christian Cachin, editors, *PODC '20: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, August 3-7, 2020*, pages 69–78. ACM, 2020. doi:[10.1145/3382734.3405745](https://doi.org/10.1145/3382734.3405745).
- [30] Gilles Brassard, Anne Broadbent, and Alain Tapp. Quantum Pseudo-Telepathy. *Foundations of Physics*, 35(11):1877–1907, 2005. doi:[10.1007/s10701-005-7353-4](https://doi.org/10.1007/s10701-005-7353-4).
- [31] Sergey Bravyi, David Gosset, and Robert König. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, 2018. doi:[10.1126/science.aar3106](https://doi.org/10.1126/science.aar3106).
- [32] Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. Bell nonlocality. *Reviews of Modern Physics*, 86(2):419–478, 2014. doi:[10.1103/revmodphys.86.419](https://doi.org/10.1103/revmodphys.86.419).
- [33] Keren Censor-Hillel, Orr Fischer, François Le Gall, Dean Leitersdorf, and Rotem Oshman. Quantum distributed algorithms for detection of cliques. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPICs*, pages 35:1–35:25. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:[10.4230/LIPICs.ITCS.2022.35](https://doi.org/10.4230/LIPICs.ITCS.2022.35).
- [34] Yi-Jun Chang. The complexity landscape of distributed locally checkable problems on trees. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing, DISC 2020, October 12-16, 2020, Virtual Conference*, volume 179 of *LIPICs*, pages 18:1–18:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:[10.4230/LIPICs.DISC.2020.18](https://doi.org/10.4230/LIPICs.DISC.2020.18).
- [35] Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. An exponential separation between randomized and deterministic complexity in the LOCAL model. *SIAM Journal on Computing*, 48(1):122–143, 2019. doi:[10.1137/17M1117537](https://doi.org/10.1137/17M1117537).
- [36] Yi-Jun Chang and Seth Pettie. A time hierarchy theorem for the LOCAL model. *SIAM Journal on Computing*, 48(1):33–69, 2019. doi:[10.1137/17M1157957](https://doi.org/10.1137/17M1157957).
- [37] Yi-Jun Chang, Jan Studený, and Jukka Suomela. Distributed graph problems through an automata-theoretic lens. *Theoretical Computer Science*, 951:113710, 2023. doi:[10.1016/J.TCS.2023.113710](https://doi.org/10.1016/J.TCS.2023.113710).
- [38] B. S. Cirel’son. Quantum generalizations of Bell’s inequality. *Letters in Mathematical Physics*, 4(2):93–100, 1980. doi:[10.1007/bf00417500](https://doi.org/10.1007/bf00417500).
- [39] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed Experiment to Test Local Hidden-Variable Theories. *Physical Review Letters*, 23(15):880–884, 1969. doi:[10.1103/physrevlett.23.880](https://doi.org/10.1103/physrevlett.23.880).

- [40] Xavier Coiteux-Roy, Francesco D’Amore, Rishikesh Gajjala, Fabian Kuhn, François Le Gall, Henrik Lievonen, Augusto Modanese, Marc-Olivier Renou, Gustav Schmid, and Jukka Suomela. No distributed quantum advantage for approximate graph coloring. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 1901–1910. ACM, 2024. doi:10.1145/3618260.3649679.
- [41] Sameep Dahal, Francesco D’Amore, Henrik Lievonen, Timoth e Picavet, and Jukka Suomela. Brief announcement: Distributed derandomization revisited. In Rotem Oshman, editor, *37th International Symposium on Distributed Computing, DISC 2023, October 10-12, 2023, L’Aquila, Italy*, volume 281 of *LIPICs*, pages 40:1–40:5. Schloss Dagstuhl - Leibniz-Zentrum f ur Informatik, 2023. doi:10.4230/LIPICs.DISC.2023.40.
- [42] Pierre Fraigniaud, Ma el Luce, Fr ed eric Magniez, and Ioan Todinca. Even-cycle detection in the randomized and quantum CONGEST model. In Ran Gelles, Dennis Olivetti, and Petr Kuznetsov, editors, *Proceedings of the 43rd ACM Symposium on Principles of Distributed Computing, PODC 2024, Nantes, France, June 17-21, 2024*, pages 209–219. ACM, 2024. doi:10.1145/3662158.3662767.
- [43] Cyril Gavoille, Adrian Kosowski, and Marcin Markiewicz. What can be observed locally? In Idit Keidar, editor, *Distributed Computing, 23rd International Symposium, DISC 2009, Elche, Spain, September 23-25, 2009. Proceedings*, volume 5805 of *Lecture Notes in Computer Science*, pages 243–257. Springer, 2009. doi:10.1007/978-3-642-04355-0_26.
- [44] Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 270–277. SIAM, 2016. doi:10.1137/1.9781611974331.CH20.
- [45] Mohsen Ghaffari, Juho Hirvonen, Fabian Kuhn, and Yannic Maus. Improved Distributed Delta-Coloring. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*, pages 427–436. ACM, 2018. doi:10.1145/3212734.3212764.
- [46] Mohsen Ghaffari and Fabian Kuhn. Deterministic distributed vertex coloring: Simpler, faster, and without network decomposition. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1009–1020. IEEE, 2021. doi:10.1109/FOCS52979.2021.00101.
- [47] Daniel M. Greenberger, Michael A. Horne, and Anton Zeilinger. Going beyond Bell’s theorem. In *Bell’s Theorem, Quantum Theory and Conceptions of the Universe*, pages 69–72. Springer Netherlands, 1989. doi:10.1007/978-94-017-0849-4_10.
- [48] Christoph Grunau, V aclav Rozhon, and Sebastian Brandt. The landscape of distributed complexities on trees and beyond. In Alessia Milani and Philipp Woelfel, editors, *PODC ’22: ACM Symposium on Principles of Distributed Computing, Salerno, Italy, July 25 - 29, 2022*, pages 37–47. ACM, 2022. doi:10.1145/3519270.3538452.
- [49] Juho Hirvonen and Jukka Suomela. Distributed algorithms. <https://jukkasuomela.fi/da2020/>, 2020.

- [50] Taisuke Izumi and François Le Gall. Quantum distributed algorithm for the all-pairs shortest path problem in the CONGEST-CLIQUE model. In Peter Robinson and Faith Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 84–93. ACM, 2019. doi:10.1145/3293611.3331628.
- [51] Taisuke Izumi, François Le Gall, and Frédéric Magniez. Quantum distributed algorithm for triangle finding in the CONGEST model. In Christophe Paul and Markus Bläser, editors, *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France*, volume 154 of *LIPICs*, pages 23:1–23:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.STACS.2020.23.
- [52] François Le Gall and Frédéric Magniez. Sublinear-time quantum computation of the diameter in CONGEST networks. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*, pages 337–346. ACM, 2018. doi:10.1145/3212734.3212744.
- [53] François Le Gall, Harumichi Nishimura, and Ansis Rosmanis. Quantum advantage for the LOCAL model in distributed computing. In Rolf Niedermeier and Christophe Paul, editors, *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, volume 126 of *LIPICs*, pages 49:1–49:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.STACS.2019.49.
- [54] Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992. doi:10.1137/0221015.
- [55] Frédéric Magniez and Ashwin Nayak. Quantum distributed complexity of set disjointness on a line. *ACM Transactions on Computation Theory*, 14(1):5:1–5:22, 2022. doi:10.1145/3512751.
- [56] N. David Mermin. Quantum mysteries revisited. *American Journal of Physics*, 58(8):731–734, 1990. doi:10.1119/1.16503.
- [57] Moni Naor and Larry J. Stockmeyer. What can be computed locally? *SIAM Journal on Computing*, 24(6):1259–1277, 1995. doi:10.1137/S0097539793254571.
- [58] Dennis Olivetti. Round Eliminator: a tool for automatic speedup simulation. <https://github.com/oliddennis/round-eliminator>, 2020.
- [59] Sandu Popescu and Daniel Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24(3):379–385, 1994. doi:10.1007/bf02058098.
- [60] Armin Tavakoli, Alejandro Pozas-Kerstjens, Ming-Xing Luo, and Marc-Olivier Renou. Bell nonlocality in networks. *Reports on Progress in Physics*, 85(5):056001, 2022. doi:10.1088/1361-6633/ac41bb.
- [61] Joran van Apeldoorn and Tijn de Vos. A framework for distributed quantum queries in the CONGEST model. In Alessia Milani and Philipp Woelfel, editors, *PODC '22: ACM Symposium on Principles of Distributed Computing, Salerno, Italy, July 25 - 29, 2022*, pages 109–119. ACM, 2022. doi:10.1145/3519270.3538413.
- [62] ChengSheng Wang, Xudong Wu, and Penghui Yao. Complexity of eccentricities and all-pairs shortest paths in the quantum CONGEST model. *CoRR*, abs/2206.02766, 2022. arXiv:2206.02766, doi:10.48550/ARXIV.2206.02766.

- [63] Elie Wolfe, Robert W. Spekkens, and Tobias Fritz. The Inflation Technique for Causal Inference with Latent Variables. *Journal of Causal Inference*, 7(2), 2019. [doi:10.1515/jci-2017-0020](https://doi.org/10.1515/jci-2017-0020).
- [64] Xudong Wu and Penghui Yao. Quantum complexity of weighted diameter and radius in CONGEST networks. In Alessia Milani and Philipp Woelfel, editors, *PODC '22: ACM Symposium on Principles of Distributed Computing, Salerno, Italy, July 25 - 29, 2022*, pages 120–130. ACM, 2022. [doi:10.1145/3519270.3538441](https://doi.org/10.1145/3519270.3538441).

A Omitted proofs for the first round elimination step

We first observe that, a necessary condition for a label to be at least as strong as another label, is for the two labels to be of the same color.

Observation A.1. *Let L_1 and L_2 be two labels of different colors. Then, $L_1 \not\leq L_2$ and $L_2 \not\leq L_1$.*

Then, by the definition of strength relation, the following also holds.

Observation A.2. *Let L_1 and L_2 be two labels. Let C be the set of configurations described by a set S of condensed configurations. Assume that, for all condensed configurations of S it holds that, if a disjunction contains L_1 , then the disjunction also contains L_2 . Then, w.r.t. the constraint C , it holds that $L_1 \leq L_2$.*

By transitivity of the strength relation, the following also holds.

Observation A.3. *Let L_1, L_2, L_3 be labels satisfying $L_1 \leq L_2$ and $L_3 \not\leq L_2$. Then, $L_3 \not\leq L_1$.*

Proof of Lemma 4.15. By Observation A.1, labels of color j can only be in relation with labels of color j . By Observation A.2, the relation depicted in the diagram is present. Hence, we only need to prove that $(?)_j \not\leq (!)_j$, which holds since, while the configuration $(!)_j^2 (?)_j$ is allowed, the configuration $(!)_j^3$ is not allowed. \square

Proof of Lemma 4.12. By Observation A.1, labels of color 1 can only be in relation with labels of color 1. By Observation A.2, the relations depicted in the diagram are present. Hence, we only need to prove that no additional relations than the ones of Figure 1 hold. We consider all possible cases:

- $(--)_1 \not\leq (-0)_1$ because the configuration $(--)_1 (-0)_1^2$ is allowed but $(-0)_1^3$ is not.
- $(--)_1 \not\leq (-1)_1$ because the configuration $(-1)_1 (--)_1^2$ is allowed but $(--)_1 (-1)_1^2$ is not.
- $(--)_1 \not\leq (!!)_1$ holds by Observation A.3.
- $(-0)_1 \not\leq (-1)_1$ because the configuration $(-0)_1 (-1)_1 (--)_1$ is allowed but $(--)_1 (-1)_1^2$ is not.
- $(-1)_1 \not\leq (-0)_1$ because the configuration $(-1)_1 (-0)_1^2$ is allowed but $(-0)_1^3$ is not.
- $(-0)_1 \not\leq (!!)_1$ and $(-1)_1 \not\leq (!!)_1$ hold by Observation A.3. \square

Proof of Lemma 4.13. By Observation A.1, labels of color j can only be in relation with labels of color j . By Observation A.2, the relations depicted in the diagram are present. Hence, we only need to prove that no additional relations than the ones of Figure 2 hold. We consider all possible cases:

- $(--)_j \not\leq (00)_j$ because $(00)_j (01)_j (--)_j$ is allowed but $(01)_j (00)_j^2$ is not.
- $(--)_j \not\leq (01)_j$ because $(--)_j (00)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(--)_j \not\leq (10)_j$ because $(00)_j (10)_j (--)_j$ is allowed but $(00)_j (10)_j^2$ is not.
- $(--)_j \not\leq (11)_j$ because $(00)_j (11)_j (--)_j$ is allowed but $(00)_j (11)_j^2$ is not.
- $(--)_j \not\leq (0-)_j$, $(--)_j \not\leq (1-)_j$, $(--)_j \not\leq (-0)_j$, $(--)_j \not\leq (-1)_j$, and $(--)_j \not\leq (!!)_j$ by Observation A.3.
- $(00)_j \not\leq (01)_j$ because $(00)_j^3$ is allowed but $(01)_j (00)_j^2$ is not.
- $(00)_j \not\leq (10)_j$ because $(10)_j (00)_j^2$ is allowed but $(00)_j (10)_j^2$ is not.
- $(00)_j \not\leq (11)_j$ because $(11)_j (00)_j^2$ is allowed but $(00)_j (11)_j^2$ is not.
- $(00)_j \not\leq (0-)_j$, $(00)_j \not\leq (1-)_j$, $(00)_j \not\leq (-0)_j$, $(00)_j \not\leq (-1)_j$, and $(00)_j \not\leq (!!)_j$ by Observation A.3.
- $(01)_j \not\leq (00)_j$ because $(00)_j (01)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(01)_j \not\leq (10)_j$ because $(00)_j (01)_j (10)_j$ is allowed but $(00)_j (10)_j^2$ is not.

- $(01)_j \not\leq (11)_j$ because $(00)_j (01)_j (11)_j$ is allowed but $(00)_j (11)_j^2$ is not.
- $(01)_j \not\leq (0-)_j$, $(01)_j \not\leq (1-)_j$, $(01)_j \not\leq (-0)_j$, $(01)_j \not\leq (-1)_j$, and $(01)_j \not\leq (!!)_j$ by Observation A.3.
- $(10)_j \not\leq (00)_j$ because $(00)_j (01)_j (10)_j$ is allowed but $(01)_j (00)_j^2$ is not.
- $(10)_j \not\leq (01)_j$ because $(10)_j (00)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(10)_j \not\leq (11)_j$ because $(00)_j (10)_j (11)_j$ is allowed but $(00)_j (11)_j^2$ is not.
- $(10)_j \not\leq (0-)_j$, $(10)_j \not\leq (1-)_j$, $(10)_j \not\leq (-0)_j$, $(10)_j \not\leq (-1)_j$, $(10)_j \not\leq (!!)_j$ by Observation A.3.
- $(11)_j \not\leq (00)_j$ because $(00)_j (01)_j (11)_j$ is allowed but $(01)_j (00)_j^2$ is not.
- $(11)_j \not\leq (01)_j$ because $(11)_j (00)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(11)_j \not\leq (10)_j$ because $(00)_j (10)_j (11)_j$ is allowed but $(00)_j (10)_j^2$ is not.
- $(11)_j \not\leq (0-)_j$, $(11)_j \not\leq (1-)_j$, $(11)_j \not\leq (-0)_j$, $(11)_j \not\leq (-1)_j$, and $(11)_j \not\leq (!!)_j$ by Observation A.3.
- $(0-)_j \not\leq (10)_j$ because $(00)_j (10)_j (0-)_j$ is allowed but $(00)_j (10)_j^2$ is not.
- $(0-)_j \not\leq (11)_j$ because $(00)_j (11)_j (0-)_j$ is allowed but $(00)_j (11)_j^2$ is not.
- $(0-)_j \not\leq (1-)_j$, $(0-)_j \not\leq (-0)_j$, $(0-)_j \not\leq (-1)_j$, and $(0-)_j \not\leq (!!)_j$ by Observation A.3.
- $(1-)_j \not\leq (00)_j$ because $(00)_j (01)_j (1-)_j$ is allowed but $(01)_j (00)_j^2$ is not.
- $(1-)_j \not\leq (01)_j$ because $(1-)_j (00)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(1-)_j \not\leq (0-)_j$, $(1-)_j \not\leq (-0)_j$, $(1-)_j \not\leq (-1)_j$, and $(1-)_j \not\leq (!!)_j$ by Observation A.3.
- $(-0)_j \not\leq (01)_j$ because $(-0)_j (00)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(-0)_j \not\leq (11)_j$ because $(00)_j (11)_j (-0)_j$ is allowed but $(00)_j (11)_j^2$ is not.
- $(-0)_j \not\leq (0-)_j$, $(-0)_j \not\leq (1-)_j$, $(-0)_j \not\leq (-1)_j$, and $(-0)_j \not\leq (!!)_j$ by Observation A.3.
- $(-1)_j \not\leq (00)_j$ because $(00)_j (01)_j (-1)_j$ is allowed but $(01)_j (00)_j^2$ is not.
- $(-1)_j \not\leq (10)_j$ because $(00)_j (10)_j (-1)_j$ is allowed but $(00)_j (10)_j^2$ is not.
- $(-1)_j \not\leq (0-)_j$, $(-1)_j \not\leq (1-)_j$, $(-1)_j \not\leq (-0)_j$, and $(-1)_j \not\leq (!!)_j$ by Observation A.3. □

Proof of Lemma 4.14. By Observation A.1, labels of color j can only be in relation with labels of color j . By Observation A.2, the relations depicted in the diagram are present. Hence, we only need to prove that no additional relations than the ones of Figure 3 hold. We consider all possible cases:

- $(--)_j \not\leq (00)_j$ because $(00)_j (01)_j (--)_j$ is allowed but $(01)_j (00)_j^2$ is not.
- $(--)_j \not\leq (01)_j$ because $(--)_j (00)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(--)_j \not\leq (10)_j$ because $(00)_j (10)_j (--)_j$ is allowed but $(00)_j (10)_j^2$ is not.
- $(--)_j \not\leq (11)_j$ because $(00)_j (11)_j (--)_j$ is allowed but $(00)_j (11)_j^2$ is not.
- $(--)_j \not\leq (-0)_j$ and $(--)_j \not\leq (-1)_j$ by Observation A.3.
- $(00)_j \not\leq (01)_j$ because $(00)_j^3$ is allowed but $(01)_j (00)_j^2$ is not.
- $(00)_j \not\leq (10)_j$ because $(10)_j (00)_j^2$ is allowed but $(00)_j (10)_j^2$ is not.
- $(00)_j \not\leq (11)_j$ because $(11)_j (00)_j^2$ is allowed but $(00)_j (11)_j^2$ is not.
- $(00)_j \not\leq (-0)_j$ and $(00)_j \not\leq (-1)_j$ by Observation A.3.
- $(01)_j \not\leq (00)_j$ because $(00)_j (01)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(01)_j \not\leq (10)_j$ because $(00)_j (01)_j (10)_j$ is allowed but $(00)_j (10)_j^2$ is not.
- $(01)_j \not\leq (11)_j$ because $(00)_j (01)_j (11)_j$ is allowed but $(00)_j (11)_j^2$ is not.
- $(01)_j \not\leq (-0)_j$ and $(01)_j \not\leq (-1)_j$ by Observation A.3.
- $(10)_j \not\leq (00)_j$ because $(00)_j (01)_j (10)_j$ is allowed but $(01)_j (00)_j^2$ is not.
- $(10)_j \not\leq (01)_j$ because $(10)_j (00)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(10)_j \not\leq (11)_j$ because $(00)_j (10)_j (11)_j$ is allowed but $(00)_j (11)_j^2$ is not.
- $(10)_j \not\leq (-0)_j$ and $(10)_j \not\leq (-1)_j$ by Observation A.3.
- $(11)_j \not\leq (00)_j$ because $(00)_j (01)_j (11)_j$ is allowed but $(01)_j (00)_j^2$ is not.

- $(11)_j \not\leq (01)_j$ because $(11)_j (00)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(11)_j \not\leq (10)_j$ because $(00)_j (10)_j (11)_j$ is allowed but $(00)_j (10)_j^2$ is not.
- $(11)_j \not\leq (-0)_j$ and $(11)_j \not\leq (-1)_j$ by Observation A.3.
- $(-0)_j \not\leq (01)_j$ because $(-0)_j (00)_j^2$ is allowed but $(01)_j (00)_j^2$ is not.
- $(-0)_j \not\leq (11)_j$ because $(00)_j (11)_j (-0)_j$ is allowed but $(00)_j (11)_j^2$ is not.
- $(-1)_j \not\leq (00)_j$ because $(00)_j (01)_j (-1)_j$ is allowed but $(01)_j (00)_j^2$ is not.
- $(-1)_j \not\leq (10)_j$ because $(00)_j (10)_j (-1)_j$ is allowed but $(00)_j (10)_j^2$ is not.
- $(-0)_j \not\leq (-1)_j$ and $(-1)_j \not\leq (-0)_j$ by Observation A.3. □

B Omitted proofs for the second round elimination step

We observe that, as in the case of the first round elimination step, a necessary condition for a label to be at least as strong as another label, is for the two labels to be of the same color.

Observation B.1. *Let L_1 and L_2 be two labels of different colors. Then, $L_1 \not\leq L_2$ and $L_2 \not\leq L_1$.*

Again, by the definition of strength relation, the following also holds.

Observation B.2. *Let L_1 and L_2 be two labels. Let C be the set of configurations described by a set S of condensed configurations. Assume that, for all condensed configurations of S it holds that, if a disjunction contains L_1 , then the disjunction also contains L_2 . Then, w.r.t. the constraint C , it holds that $L_1 \leq L_2$.*

Again, by transitivity of the strength relation, the following also holds.

Observation B.3. *Let L_1, L_2, L_3 be labels satisfying $L_1 \leq L_2$ and $L_3 \not\leq L_2$. Then, $L_3 \not\leq L_1$.*

Proof of Lemma 4.24. By Observation B.1, labels of color j can only be in relation with labels of color j . By Observation B.2, the relation depicted in the diagram is present. Hence, we only need to prove that $\langle (!)_j \rangle \not\leq \langle (?)_j \rangle$, which holds since the configuration $L_1 \dots L_\Delta$ is allowed, where:

- For $1 \leq k \leq \Delta - i - 1$, $L_k = \langle (--)_k \rangle$;
- $L_{\Delta-i} = \langle (00)_{\Delta-i}, (01)_{\Delta-i} \rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$ such that $j \neq k$, $L_k = \langle (?)_k \rangle$;
- $L_j = \langle (!)_j \rangle$,

but by replacing $\langle (!)_j \rangle$ with $\langle (?)_j \rangle$ we obtain a configuration that is not allowed. □

Proof of Lemma 4.21. By Observation B.1, labels of color 1 can only be in relation with labels of color 1. By Observation B.2, the relations depicted in the diagram are present. Hence, we only need to prove that no additional relations than the ones of Figure 5 hold. We consider all possible cases:

- $\langle (!!)_1 \rangle \not\leq \langle (-0)_1 \rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:
 - $L_1 = \langle (!!)_1 \rangle$;
 - For $2 \leq k \leq \Delta - i - 1$, $L_k = \langle (11)_k \rangle$;
 - $L_{\Delta-i} = \langle (10)_{\Delta-i}, (11)_{\Delta-i} \rangle$;
 - For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle (?)_k \rangle$;

but by replacing $\langle (!!)_1 \rangle$ with $\langle (-0)_1 \rangle$ we obtain a configuration that is not allowed.

- $\langle (!!)_1 \rangle \not\leq \langle (-1)_1 \rangle$ for symmetric reasons as in the case $\langle (!!)_1 \rangle \not\leq \langle (-0)_1 \rangle$.
- $\langle (!!)_1 \rangle \not\leq \langle (--)_1 \rangle$ by Observation B.3.
- $\langle (-0)_1 \rangle \not\leq \langle (-1)_1 \rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:
 - $L_1 = \langle (-0)_1 \rangle$;
 - For $2 \leq k \leq \Delta - i - 1$, $L_k = \langle (00)_k \rangle$;
 - $L_{\Delta-i} = \langle (00)_{\Delta-i}, (01)_{\Delta-i} \rangle$;
 - For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle (?)_k \rangle$;

but by replacing $\langle (-0)_1 \rangle$ with $\langle (-1)_1 \rangle$ we obtain a configuration that is not allowed.

- $\langle (-1)_1 \rangle \not\leq \langle (-0)_1 \rangle$ for symmetric reasons as in the case $\langle (-0)_1 \rangle \not\leq \langle (-1)_1 \rangle$.
- $\langle (-1)_1 \rangle \not\leq \langle (--)_1 \rangle$ and $\langle (-0)_1 \rangle \not\leq \langle (--)_1 \rangle$ hold by Observation B.3. \square

Proof of Lemma 4.22. By Observation B.1, labels of color j can only be in relation with labels of color j . By Observation B.2, the relations depicted in the diagram are present. Hence, we only need to prove that no additional relations than the ones of Figure 6 hold. We consider all possible cases.

- $\langle (00)_j \rangle \not\leq \langle (1-j), (-1)_j \rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:
 - $L_1 = \langle (-0)_1 \rangle$;
 - For $2 \leq k \leq \Delta - i - 1$, $L_k = \langle (00)_k \rangle$;
 - $L_{\Delta-i} = \langle (00)_{\Delta-i}, (01)_{\Delta-i} \rangle$;
 - For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle (?)_k \rangle$;

but by replacing $\langle (00)_j \rangle$ with $\langle (1-j), (-1)_j \rangle$ we obtain a configuration that is not allowed.

- $\langle (00)_j \rangle \not\leq \langle (--)_j \rangle$, $\langle (00)_j \rangle \not\leq \langle (01)_j \rangle$, $\langle (00)_j \rangle \not\leq \langle (10)_j \rangle$, $\langle (00)_j \rangle \not\leq \langle (11)_j \rangle$,
 $\langle (00)_j \rangle \not\leq \langle (01)_j, (10)_j \rangle$, $\langle (00)_j \rangle \not\leq \langle (1-j)_j \rangle$, $\langle (00)_j \rangle \not\leq \langle (-1)_j \rangle$ hold by Observation B.3.
- $\langle (01)_j \rangle \not\leq \langle (--)_j \rangle$, $\langle (01)_j \rangle \not\leq \langle (00)_j \rangle$, $\langle (01)_j \rangle \not\leq \langle (10)_j \rangle$, $\langle (01)_j \rangle \not\leq \langle (11)_j \rangle$,
 $\langle (01)_j \rangle \not\leq \langle (00)_j, (11)_j \rangle$, $\langle (01)_j \rangle \not\leq \langle (1-j)_j \rangle$, $\langle (01)_j \rangle \not\leq \langle (-0)_j \rangle$, $\langle (01)_j \rangle \not\leq \langle (1-j), (-0)_j \rangle$,
 $\langle (10)_j \rangle \not\leq \langle (--)_j \rangle$, $\langle (10)_j \rangle \not\leq \langle (00)_j \rangle$, $\langle (10)_j \rangle \not\leq \langle (01)_j \rangle$, $\langle (10)_j \rangle \not\leq \langle (11)_j \rangle$,
 $\langle (10)_j \rangle \not\leq \langle (00)_j, (11)_j \rangle$, $\langle (10)_j \rangle \not\leq \langle (0-j)_j \rangle$, $\langle (10)_j \rangle \not\leq \langle (-1)_j \rangle$, $\langle (10)_j \rangle \not\leq \langle (0-j), (-1)_j \rangle$,
 $\langle (11)_j \rangle \not\leq \langle (--)_j \rangle$, $\langle (11)_j \rangle \not\leq \langle (00)_j \rangle$, $\langle (11)_j \rangle \not\leq \langle (01)_j \rangle$, $\langle (11)_j \rangle \not\leq \langle (10)_j \rangle$,
 $\langle (11)_j \rangle \not\leq \langle (01)_j, (10)_j \rangle$, $\langle (11)_j \rangle \not\leq \langle (0-j)_j \rangle$, $\langle (11)_j \rangle \not\leq \langle (-0)_j \rangle$, $\langle (11)_j \rangle \not\leq \langle (0-j), (-0)_j \rangle$ for symmetric reasons.

- $\langle (0-j)_j \rangle \not\leq \langle (1-j), (-0)_j \rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:
 - $L_1 = \langle (-0)_1 \rangle$;
 - For $2 \leq k < j$, $L_k = \langle (00)_k \rangle$;
 - $L_j = \langle (0-j)_j \rangle$;
 - For $j < k \leq \Delta - i - 1$, $L_k = \langle (11)_k \rangle$;
 - $L_{\Delta-i} = \langle (10)_{\Delta-i}, (11)_{\Delta-i} \rangle$;
 - For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle (?)_k \rangle$;

but by replacing $\langle (0-j)_j \rangle$ with $\langle (1-j), (-0)_j \rangle$ we obtain a configuration that is not allowed.

- $\langle (0-j)_j \rangle \not\leq \langle (1-j), (-1)_j \rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-0)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(00)_k\rangle$;
- $L_j = \langle(0-)_j\rangle$
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
- $L_{\Delta-i} = \langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(0-)_j\rangle$ with $\langle(1-)_j, (-1)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(0-)_j\rangle \not\leq \langle(-)_j\rangle$, $\langle(0-)_j\rangle \not\leq \langle(00)_j\rangle$, $\langle(0-)_j\rangle \not\leq \langle(01)_j\rangle$, $\langle(0-)_j\rangle \not\leq \langle(10)_j\rangle$,
 $\langle(0-)_j\rangle \not\leq \langle(01)_j, (10)_j\rangle$, $\langle(0-)_j\rangle \not\leq \langle(11)_j\rangle$, $\langle(0-)_j\rangle \not\leq \langle(00)_j, (11)_j\rangle$, $\langle(0-)_j\rangle \not\leq \langle(1-)_j\rangle$,
 $\langle(0-)_j\rangle \not\leq \langle(-0)_j\rangle$, $\langle(0-)_j\rangle \not\leq \langle(-1)_j\rangle$ hold by Observation B.3.
- $\langle(1-)_j\rangle \not\leq \langle(-)_j\rangle$, $\langle(1-)_j\rangle \not\leq \langle(00)_j\rangle$, $\langle(1-)_j\rangle \not\leq \langle(01)_j\rangle$, $\langle(1-)_j\rangle \not\leq \langle(10)_j\rangle$,
 $\langle(1-)_j\rangle \not\leq \langle(01)_j, (10)_j\rangle$, $\langle(1-)_j\rangle \not\leq \langle(11)_j\rangle$, $\langle(1-)_j\rangle \not\leq \langle(00)_j, (11)_j\rangle$, $\langle(1-)_j\rangle \not\leq \langle(0-)_j\rangle$,
 $\langle(1-)_j\rangle \not\leq \langle(-0)_j\rangle$, $\langle(1-)_j\rangle \not\leq \langle(0-)_j, (-0)_j\rangle$, $\langle(1-)_j\rangle \not\leq \langle(-1)_j\rangle$, $\langle(1-)_j\rangle \not\leq \langle(0-)_j, (-1)_j\rangle$ for
symmetric reasons.
- $\langle(-0)_j\rangle \not\leq \langle(0-)_j, (-1)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:
 - $L_1 = \langle(-1)_1\rangle$;
 - For $2 \leq k < j$, $L_k = \langle(11)_k\rangle$;
 - $L_j = \langle(-0)_j\rangle$
 - For $j < k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
 - $L_{\Delta-i} = \langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle$;
 - For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(-0)_j\rangle$ with $\langle(0-)_j, (-1)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(-0)_j\rangle \not\leq \langle(1-)_j, (-1)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:
 - $L_1 = \langle(-0)_1\rangle$;
 - For $2 \leq k < j$, $L_k = \langle(00)_k\rangle$;
 - $L_j = \langle(-0)_j\rangle$
 - For $j < k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
 - $L_{\Delta-i} = \langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle$;
 - For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;
- but by replacing $\langle(-0)_j\rangle$ with $\langle(1-)_j, (-1)_j\rangle$ we obtain a configuration that is not allowed.
- $\langle(-0)_j\rangle \not\leq \langle(-)_j\rangle$, $\langle(-0)_j\rangle \not\leq \langle(00)_j\rangle$, $\langle(-0)_j\rangle \not\leq \langle(01)_j\rangle$, $\langle(-0)_j\rangle \not\leq \langle(10)_j\rangle$,
 $\langle(-0)_j\rangle \not\leq \langle(01)_j, (10)_j\rangle$, $\langle(-0)_j\rangle \not\leq \langle(11)_j\rangle$, $\langle(-0)_j\rangle \not\leq \langle(00)_j, (11)_j\rangle$, $\langle(-0)_j\rangle \not\leq \langle(0-)_j\rangle$,
 $\langle(-0)_j\rangle \not\leq \langle(1-)_j\rangle$, $\langle(-0)_j\rangle \not\leq \langle(-1)_j\rangle$ hold by Observation B.3.
 - $\langle(-1)_j\rangle \not\leq \langle(-)_j\rangle$, $\langle(-1)_j\rangle \not\leq \langle(00)_j\rangle$, $\langle(-1)_j\rangle \not\leq \langle(01)_j\rangle$, $\langle(-1)_j\rangle \not\leq \langle(10)_j\rangle$,
 $\langle(-1)_j\rangle \not\leq \langle(01)_j, (10)_j\rangle$, $\langle(-1)_j\rangle \not\leq \langle(11)_j\rangle$, $\langle(-1)_j\rangle \not\leq \langle(00)_j, (11)_j\rangle$, $\langle(-1)_j\rangle \not\leq \langle(0-)_j\rangle$,
 $\langle(-1)_j\rangle \not\leq \langle(1-)_j\rangle$, $\langle(-1)_j\rangle \not\leq \langle(0-)_j, (-0)_j\rangle$, $\langle(-1)_j\rangle \not\leq \langle(1-)_j, (-0)_j\rangle$ for
symmetric reasons.
 - $\langle(01)_j, (10)_j\rangle \not\leq \langle(1-)_j, (-0)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:
 - $L_1 = \langle(-0)_1\rangle$;

- For $2 \leq k < j$, $L_k = \langle(00)_k\rangle$;
- $L_j = \langle(01)_j, (10)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(11)_k\rangle$;
- $L_{\Delta-i} = \langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(01)_j, (10)_j\rangle$ with $\langle(1-)_j, (-0)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(01)_j, (10)_j\rangle \not\leq \langle(0-)_j, (-1)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-1)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(11)_k\rangle$;
- $L_j = \langle(01)_j, (10)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
- $L_{\Delta-i} = \langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(01)_j, (10)_j\rangle$ with $\langle(0-)_j, (-1)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(01)_j, (10)_j\rangle \not\leq \langle(--)_j\rangle$, $\langle(01)_j, (10)_j\rangle \not\leq \langle(00)_j\rangle$, $\langle(01)_j, (10)_j\rangle \not\leq \langle(01)_j\rangle$,
 $\langle(01)_j, (10)_j\rangle \not\leq \langle(10)_j\rangle$, $\langle(01)_j, (10)_j\rangle \not\leq \langle(11)_j\rangle$, $\langle(01)_j, (10)_j\rangle \not\leq \langle(00)_j, (11)_j\rangle$,
 $\langle(01)_j, (10)_j\rangle \not\leq \langle(0-)_j\rangle$, $\langle(01)_j, (10)_j\rangle \not\leq \langle(1-)_j\rangle$, $\langle(01)_j, (10)_j\rangle \not\leq \langle(-0)_j\rangle$,
 $\langle(01)_j, (10)_j\rangle \not\leq \langle(-1)_j\rangle$ hold by Observation B.3.

- $\langle(00)_j, (11)_j\rangle \not\leq \langle(1-)_j, (-1)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-0)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(00)_k\rangle$;
- $L_j = \langle(00)_j, (11)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
- $L_{\Delta-i} = \langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(00)_j, (11)_j\rangle$ with $\langle(1-)_j, (-1)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(00)_j, (11)_j\rangle \not\leq \langle(0-)_j, (-0)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-1)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(11)_k\rangle$;
- $L_j = \langle(00)_j, (11)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(11)_k\rangle$;
- $L_{\Delta-i} = \langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(00)_j, (11)_j\rangle$ with $\langle(0-)_j, (-0)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(00)_j, (11)_j\rangle \not\leq \langle(--)_j\rangle$, $\langle(00)_j, (11)_j\rangle \not\leq \langle(00)_j\rangle$, $\langle(00)_j, (11)_j\rangle \not\leq \langle(01)_j\rangle$,
 $\langle(00)_j, (11)_j\rangle \not\leq \langle(10)_j\rangle$, $\langle(00)_j, (11)_j\rangle \not\leq \langle(01)_j, (10)_j\rangle$, $\langle(00)_j, (11)_j\rangle \not\leq \langle(11)_j\rangle$,
 $\langle(00)_j, (11)_j\rangle \not\leq \langle(0-)_j\rangle$, $\langle(00)_j, (11)_j\rangle \not\leq \langle(1-)_j\rangle$, $\langle(00)_j, (11)_j\rangle \not\leq \langle(-0)_j\rangle$,
 $\langle(00)_j, (11)_j\rangle \not\leq \langle(-1)_j\rangle$ hold by Observation B.3.

- $\langle(0-)_j, (-0)_j\rangle \not\leq \langle(1-)_j, (-0)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-0)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(00)_k\rangle$;
- $L_j = \langle(0-)_j, (-0)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(11)_k\rangle$;
- $L_{\Delta-i} = \langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(0-)_j, (-0)_j\rangle$ with $\langle(1-)_j, (-0)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(0-)_j, (-1)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-1)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(11)_k\rangle$;
- $L_j = \langle(0-)_j, (-0)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
- $L_{\Delta-i} = \langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(0-)_j, (-0)_j\rangle$ with $\langle(0-)_j, (-1)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(1-)_j, (-1)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-0)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(00)_k\rangle$;
- $L_j = \langle(0-)_j, (-0)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
- $L_{\Delta-i} = \langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(0-)_j, (-0)_j\rangle$ with $\langle(1-)_j, (-1)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(-)_j\rangle$, $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(00)_j\rangle$, $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(01)_j\rangle$,
 $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(10)_j\rangle$, $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(01)_j, (10)_j\rangle$, $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(11)_j\rangle$,
 $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(00)_j, (11)_j\rangle$, $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(0-)_j\rangle$, $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(1-)_j\rangle$,
 $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(-0)_j\rangle$, $\langle(0-)_j, (-0)_j\rangle \not\preceq \langle(-1)_j\rangle$ hold by Observation B.3.

- $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(1-)_j, (-1)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-0)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(00)_k\rangle$;
- $L_j = \langle(0-)_j, (-1)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
- $L_{\Delta-i} = \langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(0-)_j, (-1)_j\rangle$ with $\langle(1-)_j, (-1)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(1-)_j, (-0)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-0)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(00)_k\rangle$;
- $L_j = \langle(0-)_j, (-1)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(11)_k\rangle$;
- $L_{\Delta-i} = \langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(0-)_j, (-1)_j\rangle$ with $\langle(1-)_j, (-0)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(0-)_j, (-0)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-1)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(11)_k\rangle$;
- $L_j = \langle(0-)_j, (-1)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(11)_k\rangle$;
- $L_{\Delta-i} = \langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(0-)_j, (-1)_j\rangle$ with $\langle(0-)_j, (-0)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(--)_j\rangle$, $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(00)_j\rangle$, $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(01)_j\rangle$,
 $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(10)_j\rangle$, $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(01)_j, (10)_j\rangle$, $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(11)_j\rangle$,
 $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(00)_j, (11)_j\rangle$, $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(0-)_j\rangle$, $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(1-)_j\rangle$,
 $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(-0)_j\rangle$, $\langle(0-)_j, (-1)_j\rangle \not\preceq \langle(-1)_j\rangle$ hold by Observation B.3.
- $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(--)_j\rangle$, $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(00)_j\rangle$, $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(01)_j\rangle$,
 $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(10)_j\rangle$, $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(01)_j, (10)_j\rangle$, $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(11)_j\rangle$,
 $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(00)_j, (11)_j\rangle$, $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(0-)_j\rangle$, $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(1-)_j\rangle$,
 $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(-0)_j\rangle$, $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(0-)_j, (-0)_j\rangle$, $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(-1)_j\rangle$,
 $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(0-)_j, (-1)_j\rangle$, $\langle(1-)_j, (-0)_j\rangle \not\preceq \langle(1-)_j, (-1)_j\rangle$ for symmetric reasons.
- $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(--)_j\rangle$, $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(00)_j\rangle$, $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(01)_j\rangle$,
 $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(10)_j\rangle$, $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(01)_j, (10)_j\rangle$, $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(11)_j\rangle$,
 $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(00)_j, (11)_j\rangle$, $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(0-)_j\rangle$, $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(1-)_j\rangle$,
 $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(-0)_j\rangle$, $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(0-)_j, (-0)_j\rangle$, $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(1-)_j, (-0)_j\rangle$,
 $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(-1)_j\rangle$, $\langle(1-)_j, (-1)_j\rangle \not\preceq \langle(0-)_j, (-1)_j\rangle$ for symmetric reasons.
- $\langle(!!)_j\rangle \not\preceq \langle(0-)_j, (-1)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-1)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(11)_k\rangle$;
- $L_j = \langle(!!)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
- $L_{\Delta-i} = \langle(00)_{\Delta-i}, (01)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(!!)_j\rangle$ with $\langle(0-)_j, (-1)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(!!)_j\rangle \not\preceq \langle(0-)_j, (-0)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:

- $L_1 = \langle(-1)_1\rangle$;
- For $2 \leq k < j$, $L_k = \langle(11)_k\rangle$;
- $L_j = \langle(!)_j\rangle$;
- For $j < k \leq \Delta - i - 1$, $L_k = \langle(11)_k\rangle$;
- $L_{\Delta-i} = \langle(10)_{\Delta-i}, (11)_{\Delta-i}\rangle$;
- For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(!)_j\rangle$ with $\langle(0-)_j, (-0)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(!)_j\rangle \not\leq \langle(1-)_j, (-1)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(1-)_j, (-0)_j\rangle$ for symmetric reasons.
- $\langle(!)_j\rangle \not\leq \langle(-)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(00)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(01)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(10)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(01)_j, (10)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(11)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(00)_j, (11)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(0-)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(1-)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(-0)_j\rangle$, $\langle(!)_j\rangle \not\leq \langle(-1)_j\rangle$ hold by Observation B.3. \square

Proof of Lemma 4.23. By Observation B.1, labels of the special color $j = \Delta - i$ can only be in relation with labels of color j . By Observation B.2, the relations depicted in the diagram are present. Hence, we only need to prove that no additional relations than the ones of Figure 7 hold. We consider all possible cases:

- $\langle(00)_j, (01)_j\rangle \not\leq \langle(10)_j, (11)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:
 - $L_1 = \langle(-0)_1\rangle$;
 - For $2 \leq k \leq \Delta - i - 1$, $L_k = \langle(00)_k\rangle$;
 - $L_j = \langle(00)_j, (01)_j\rangle$;
 - For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(00)_j, (01)_j\rangle$ with $\langle(10)_j, (11)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(10)_j, (11)_j\rangle \not\leq \langle(00)_j, (01)_j\rangle$ for symmetric reasons as in the case $\langle(00)_j, (01)_j\rangle \not\leq \langle(10)_j, (11)_j\rangle$.
- $\langle(-0)_j, (-1)_j\rangle \not\leq \langle(00)_j, (01)_j\rangle$ because the configuration $L_1 \dots L_\Delta$ is allowed, where:
 - $L_1 = \langle(-1)_1\rangle$;
 - For $2 \leq k \leq \Delta - i - 1$, $L_k = \langle(11)_k\rangle$;
 - $L_j = \langle(-0)_j, (-1)_j\rangle$;
 - For $\Delta - i + 1 \leq k \leq \Delta$, $L_k = \langle(?)_k\rangle$;

but by replacing $\langle(-0)_j, (-1)_j\rangle$ with $\langle(00)_j, (01)_j\rangle$ we obtain a configuration that is not allowed.

- $\langle(-0)_j, (-1)_j\rangle \not\leq \langle(10)_j, (11)_j\rangle$ for symmetric reasons as in the case $\langle(-0)_j, (-1)_j\rangle \not\leq \langle(00)_j, (01)_j\rangle$. \square