No Distributed Quantum Advantage for Approximate Graph Coloring

Xavier Coiteux-Roy
TU Munich
Munich, Germany
Munich Center for Quantum Science and Technology
Munich, Germany
xavier.coiteux-roy@tum.de

Francesco d'Amore
Aalto University
Espoo, Finland
Bocconi University, BIDSA
Milan, Italy
francesco.damore2@unibocconi.it

Rishikesh Gajjala
Indian Institute of Science
Bangalore, India
Aalto University
Espoo, Finland
rishikeshg@iisc.ac.in

Fabian Kuhn
University of Freiburg
Freiburg, Germany
kuhn@cs.uni-freiburg.de

François Le Gall
Nagoya University
Nagoya, Japan
legall@math.nagoya-u.ac.jp

Henrik Lievonen
Aalto University
Espoo, Finland
henrik.lievonen@aalto.fi

Augusto Modanese
Aalto University
Espoo, Finland
augusto.modanese@aalto.fi

Marc-Olivier Renou
Inria
Paris, France
Université Paris-Saclay
Paris, France
Institut Polytechnique de Paris
Palaiseau, France
marc-olivier.renou@inria.fr

Gustav Schmid
University of Freiburg
Freiburg, Germany
schmidg@informatik.uni-freiburg.de

Jukka Suomela
Aalto University
Espoo, Finland
jukka.suomela@aalto.fi

ABSTRACT

We give an almost complete characterization of the hardness of c-coloring χ-chromatic graphs with distributed algorithms, for a wide range of models of distributed computing. In particular, we show that these problems do not admit any distributed quantum advantage. To do that:

(1) We give a new distributed algorithm that finds a c-coloring in χ-chromatic graphs in $\tilde{O}(n^{\frac{1}{\alpha}})$ rounds, with $\alpha = \left\lfloor \frac{c-1}{\chi-1} \right\rfloor$.
(2) We prove that any distributed algorithm for this problem requires $\Omega(n^{\frac{1}{\alpha}})$ rounds.

Our upper bound holds in the classical, deterministic LOCAL model, while the near-matching lower bound holds in the *non-signaling* model. This model, introduced by Arfaoui and Fraigniaud in 2014, captures all models of distributed graph algorithms that obey physical causality; this includes not only classical deterministic LOCAL and randomized LOCAL but also quantum-LOCAL, even with a pre-shared quantum state.

We also show that similar arguments can be used to prove that, e.g., 3-coloring 2-dimensional grids or c-coloring trees remain hard problems even for the non-signaling model, and in particular do not admit any quantum advantage. Our lower-bound arguments are purely graph-theoretic at heart; no background on quantum information theory is needed to establish the proofs.

CCS CONCEPTS

• **Theory of computation** → **Distributed computing models**; *Quantum computation theory*; *Distributed algorithms.*

KEYWORDS

distributed computing, graph coloring, non-signaling model, quantum advantage

ACM Reference Format:
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. 2024. No Distributed Quantum Advantage for Approximate Graph Coloring. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3618260.3649679

This work is licensed under a Creative Commons Attribution 4.0 International License.

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0383-6/24/06
https://doi.org/10.1145/3618260.3649679

1 INTRODUCTION

In this work, we settle the distributed computational complexity of approximate graph coloring, for deterministic, randomized, and quantum versions of the LOCAL model of distributed computing.

In brief, the setting is this: We have an input graph G with n nodes. Each node is a computer and each edge represents a communication link. Computation proceeds in synchronous rounds: each node sends a message to each of its neighbors, receives a message from each of its neighbors, and updates its own state. After T rounds, each node has to stop and announce its own output, and the outputs have to form a proper c-coloring of the input graph G. If the chromatic number of G is χ, in this setting it is trivial to find a χ-coloring in $T = O(n)$ rounds, as in $O(n)$ rounds all nodes can learn the full topology of their own connected component, and they can locally find an optimal coloring by brute force without any further communication. But the key questions are: How well can we color graphs in $T \ll n$ rounds? And how much does it help if we use quantum computers that can exchange quantum information, possibly with a pre-shared entangled state?

1.1 Main Result

We show that for all constants c, χ, and α, it is possible to find a c-coloring of a χ-colorable graph in $T = \tilde{O}(n^{1/\alpha})$ communication rounds if and only if

$$\alpha \leq \left\lfloor \frac{c-1}{\chi-1} \right\rfloor.$$

For example, if the graph is bipartite ($\chi = 2$), this means that the complexity of 2-coloring is $\tilde{\Theta}(n)$ rounds, 3-coloring is $\tilde{\Theta}(\sqrt{n})$ rounds, and 4-coloring is $\tilde{\Theta}(n^{1/3})$ rounds. Here we use $\tilde{O}$ and $\tilde{\Theta}$ to hide polylogarithmic factors, that is, our results are tight up to polylogarithmic factors.

Perhaps the biggest surprise is that this result holds for a wide range of models of distributed computing: the answer is the same for deterministic, randomized, and quantum versions of the LOCAL model, and it holds even if the algorithm has access to shared randomness or pre-shared quantum state (as long as the quantum state is prepared before we reveal the structure of graph G).

In particular, we show that there is *no distributed quantum advantage for approximate graph coloring* in the context of the LOCAL model, at least up to polylogarithmic factors.

1.2 Significance and Motivation

Our work is directly linked to two lines of research: understanding the quantum advantage in distributed settings, and the complexity of distributed graph coloring in classical settings.

Distributed quantum advantage. There is a long line of work [16, 22, 32–34, 40, 47, 48, 50] on quantum advantage in the CONGEST model—this is a bandwidth-limited version of the LOCAL model. However, much less is known about quantum advantage in the LOCAL model.

Earlier work by Gavoille et al. [25] and Arfaoui and Fraigniaud [4] on quantum-LOCAL brought primarily bad news: they showed that many classical LOCAL model lower bounds still hold in the quantum-LOCAL model. The quantum advantage demonstrated by [25] was limited to constant factors or required pre-shared quantum resources. The major breakthrough was the recent work by Le Gall et al. [35] that demonstrated that there is a problem that can be solved in only 2 rounds using quantum communication, whereas solving it in the classical setting requires $\Omega(n)$ rounds.

However, the problem from Le Gall et al. [35] is very different from the classical problems commonly studied in the field of distributed graph algorithms, and most importantly, it is not a *locally checkable* problem. Locally checkable problems are graph problems in which the task is to find a feasible solution subject to local constraints—perhaps the best-known example of such a problem is graph coloring. A lot of recent work on the classical LOCAL model has focused on locally checkable problems, and there is nowadays a solid understanding of the landscape of the distributed computational complexity of such problems for the classical models—see, e.g., [6–9, 13, 14, 17, 19, 23, 27, 28, 45]. However, what is wide open is how quantum-LOCAL changes the picture.

A major open problem is whether there is *any* locally checkable graph problem that can be solved asymptotically faster in quantum-LOCAL in comparison with the classical randomized LOCAL model, and it has been conjectured that no such problem exists [46]. In this work we provide more evidence in support of this conjecture: we show that various problems related to graph coloring do not admit any significant quantum advantage.

Hardness of distributed coloring. In a very recent work [2], the notion of locality was studied in three different settings: distributed, dynamic, and online graph algorithms. The authors showed that for locally checkable problems in rooted regular trees the three notions of locality coincide, but more generally the notions are distinct. The prime example of a problem that separates the models is 3-coloring bipartite graphs: the distributed locality (i.e., round complexity) of the problem is $\Omega(\sqrt{n})$ [15], but the online locality is $O(\log n)$ [2]. While this demonstrates that there is large gap between distributed and online settings, this also highlights a blind spot in our understanding of seemingly elementary questions in the classical LOCAL model: What, exactly, is the distributed complexity of 3-coloring bipartite graphs? Can we solve it in $\tilde{O}(\sqrt{n})$ rounds? And, more generally, what is the distributed complexity of c-coloring χ-colorable graphs?

Given the prominent role graph coloring plays in distributed graph algorithms, the state of the art is highly unsatisfactory—the upper and lower bounds are far from each other, even if we consider the seemingly elementary question of coloring bipartite graphs:

- As mentioned above, the complexity of 3-coloring bipartite graphs is known to be somewhere between $\Omega(\sqrt{n})$ and $O(n)$. In [15], the authors show that 3-coloring 2-dimensional grids requires $\Omega(\sqrt{n})$ rounds, and even though they study toroidal grids (which are not necessarily bipartite), the same result can be adapted to also show that 3-coloring bipartite graphs requires $\Omega(\sqrt{n})$ rounds. It is not known if this is tight; to the best of our knowledge, there is no upper bound other than the trivial $O(n)$-round algorithm.
- The complexity of 4-coloring bipartite graphs is only known to be somewhere between $\Omega(\log n)$ and $O(n)$. Linial's [37] lower bound for coloring trees applies, so we know that the complexity has to be at least $\Omega(\log n)$, but beyond that very little is known. The lower bound construction from [15]

cannot be used here since it is easy to 4-color grids. To come up with a nontrivial upper bound, it would be tempting to use network decompositions in the spirit of Barenboim [10], but we are lacking network decomposition algorithms with suitable parameters, and in any case this approach cannot produce 4-colorings or 5-colorings in $o(\sqrt{n})$ rounds.

In this work we solve all these open questions, up to polylogarithmic factors, for the general task of c-coloring χ-colorable graphs. We show that there is plenty of room for improvement in both upper and lower bounds. For example, in the case of 4-coloring bipartite graphs, the right bound turns out to be $\tilde{\Theta}(n^{1/3})$, which is far from what can be achieved with the state of the art outlined above.

2 CONTRIBUTIONS IN MORE DETAIL

We will now describe all of our results and contributions in more detail; we refer to Section 3 for an overview of the proof ideas and to the full version [20] for the proofs.

2.1 Classical Upper Bound

Let us start with the upper bound. We design new distributed algorithms with the following properties:

Theorem 2.1. *There exists a* det-LOCAL *algorithm* $\mathcal{A}_{\mathsf{det}}$ *and a* rand-LOCAL *algorithm* $\mathcal{A}_{\mathsf{rand}}$ *that, given a parameter* $\alpha \in \mathbb{N}$*, find a proper vertex coloring with* $\alpha(\chi-1)+1$ *colors in any graph with chromatic number* χ*, as follows:*

- $\mathcal{A}_{\mathsf{det}}$ *runs in* $O\big(n^{1/\alpha}\log^{3-1/\alpha} n\big)\cdot(\log\log n)^{O(1)}$ *rounds.*
- $\mathcal{A}_{\mathsf{rand}}$ *runs in* $O\big(n^{1/\alpha}\log^{2-1/\alpha} n\big)$ *rounds and succeeds with probability* $1-1/\mathrm{poly}(n)$.

We note that the algorithms do not need to know χ; it is sufficient to know α and n. As a corollary, we can, e.g., 3-color bipartite graphs in $\tilde{O}(\sqrt{n})$ rounds by setting $\alpha = 2$.

The number of colors $c = \alpha(\chi-1)+1$ may look like a rather unnatural expression, and there does not seem to be a priori any reason to expect that this would be tight—however, as we will see, this is indeed exactly the right number.

2.2 Non-signaling Model

Our main goal is to show that the algorithms in Theorem 2.1 are optimal (up to polylogarithmic factors), not only in the classical models but also in, e.g., all reasonable variants of the quantum-LOCAL model. To this end, we work in the *non-signaling model*, as defined by Arfaoui and Fraigniaud [4]; this is essentially equivalent to the φ-LOCAL model defined earlier by Gavoille et al. [25].

The non-signaling model is a characterization of output distributions that do not violate the *no-signaling from the future* principle or, equivalently, the *causality* principle [21]. To better understand this, suppose we have some classical rand-LOCAL algorithm $\mathcal{A}$ that runs in T rounds and outputs a vertex coloring. Let $p(G)$ be the output distribution of $\mathcal{A}$ when run on a graph G. The key observation is that this distribution is not arbitrary—in particular, it must satisfy the following property:

Definition 2.2 (Non-signaling distribution, informal version). The output distribution $p(G)$ of $\mathcal{A}$ is non-signaling beyond distance T if the following holds: Let V be a set of nodes of a graph G with $|V| = n$.

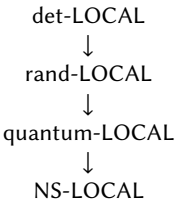

Figure 1: Four models of distributed computing that we study in this work: we prove upper bounds in det-LOCAL and rand-LOCAL, and we prove lower bounds in NS-LOCAL; quantum-LOCAL is sandwiched between them.

Fix a subset of nodes $U \subseteq V$ and consider $p(G)\restriction_U$, the restriction of $p(G)$ to U. Let $G[U,T]$ be the graph induced by the radius-T neighborhood of U in G. Now modify G outside $G[U,T]$ to obtain a different n-node graph G' while preserving $G[U,T] = G'[U,T]$. Then $p(G)\restriction_U = p(G')\restriction_U$.

Put otherwise, changes more than distance T away from U cannot influence the output distribution of U. It is not hard to see that this holds for det-LOCAL and rand-LOCAL, even if the algorithm has access to shared randomness. But what makes this notion particularly useful is that it is satisfied also by the quantum-LOCAL model, even with a pre-shared quantum state [4, 25]; see Fig. 1. Informally, a system that violates the non-signaling property would violate causality and enable faster-than-light communication, which is something that quantum physics does not allow.

We use NS-LOCAL to refer to the non-signaling model. We say that $\mathcal{A}$ is an NS-LOCAL algorithm that runs in T rounds if it produces an output distribution that is non-signaling beyond distance T. We will then prove statements of the form "any NS-LOCAL algorithm for this problem requires at least T rounds." As a corollary, this gives a T-round lower bound for det-LOCAL, rand-LOCAL, and quantum-LOCAL, even if we have access to shared randomness and pre-shared quantum states. This also puts limits on the existence of so-called finitely dependent colorings [29, 30].

2.3 Non-signaling Lower Bounds

The precise version of our lower bound result states that, for every large enough number of nodes n, there exists a χ-chromatic graph on n nodes that is hard to color in NS-LOCAL.

Theorem 2.3. *Let* $\chi \geq 2$*,* $c \geq \chi$ *be integers, and let* $\alpha = \lfloor \frac{c-1}{\chi-1} \rfloor$*. Let* $\varepsilon \in (0, \frac{\alpha-1}{\alpha})$ *be a real value, and let* $n \in \mathbb{N}$ *with*

$$n \geq \left\lceil \frac{\log \varepsilon^{-1}}{\log(1+\frac{1}{\alpha})} \right\rceil \cdot \frac{(6\chi+1)^{\alpha+1}-1}{6}.$$

Suppose $\mathcal{A}$ *is an* NS-LOCAL *algorithm for* c*-coloring graphs in the family* $\mathcal{F}$ *of* χ*-chromatic graphs of* n *nodes with success probability* $q > \varepsilon$*. Then the running time of* $\mathcal{A}$ *is at least*

$$T = \Omega\left(\frac{1}{\chi^{1+\frac{1}{\alpha}}} \cdot \left(\frac{n}{\log \varepsilon^{-1}}\right)^{\frac{1}{\alpha}}\right).$$

A key observation is that, if the parameters χ, c, and ε in Theorem 2.3 are constants, then $T = \Omega(n^{\frac{1}{\alpha}})$, which matches the upper bound in Theorem 2.1 up to polylogarithmic factors. In particular,

there is at best polylogarithmic room for any distributed quantum advantage.

While Theorem 2.3 implies bounds for coloring bipartite graphs in general, we will also use our techniques to prove bounds for specific bipartite graphs. By prior work, it is known that 3-coloring 2-dimensional grids is hard in the det-LOCAL model [1, 15]. We show that this also holds for the non-signaling model:

Theorem 2.4. *Let $\varepsilon \in (0, \frac{3}{4})$ and $N = \lceil \log(\varepsilon^{-1})/\log(\frac{4}{3}) \rceil$. Let $n_1, n_2 \in \mathbb{N}$ with $\lfloor \frac{n_1}{N} \rfloor \geq 5$ and $\lfloor \frac{n_2}{N} \rfloor \geq 5$. Suppose $\mathcal{A}$ is an* NS-LOCAL *algorithm that 3-colors $n_1 \times n_2$ grids with probability $q > \varepsilon$. Then, the running time of $\mathcal{A}$ is at least*

$$T = \Omega\left(\frac{\min(n_1, n_2)}{\log \varepsilon^{-1}}\right).$$

This result is easiest to interpret in the case of a square grid, i.e., $n_1 = n_2$. Then the lower bound (for constant ε) is simply $\Omega(\sqrt{n})$, where $n = n_1 \cdot n_2$, and this is trivially tight since the diameter of the grid is $O(\sqrt{n})$; hence the problem can be solved in $O(\sqrt{n})$ rounds with a det-LOCAL algorithm. In particular, there is no room for distributed quantum advantage (beyond possibly constant factors).

Finally, we revisit the classical result by Linial [37] about the hardness of coloring trees. We show that essentially the same lower bound holds in the non-signaling model:

Theorem 2.5. *Let $c \geq 2$ be an integer, and $\varepsilon \in (0, 1)$. Suppose $\mathcal{A}$ is an* NS-LOCAL *algorithm that c-colors trees of size $n \in \mathbb{N}$ with probability $q > \varepsilon$. Then, for infinitely many n, as long as $\varepsilon > e^{-n}$, the running time of $\mathcal{A}$ is at least*

$$T = \Omega\big(\log_c n - \log_c \log \varepsilon^{-1}\big).$$

3 KEY NEW IDEAS AND TECHNIQUES

In this section, we will give an informal overview of the key new ideas and techniques that we use to prove Theorems 2.1, 2.3, 2.4 and 2.5. We refer to the full version of the paper [20] for the formal proofs.

3.1 Classical Upper Bound

Background and prior work. The only existing distributed algorithm for solving the approximate coloring problem in general graphs that we are aware of is a folklore algorithm based on *network decompositions* [5, 38]. For parameters α and d, an (α, d)*-network decomposition* is a partition of the nodes V of a graph $G = (V, E)$ into clusters of (weak) diameter at most d together with a proper α-coloring of the cluster graph; recall that the *weak diameter* of a cluster $C \subseteq V$ is the maximum distance in G between any two nodes in C.

Given such a decomposition, it is not hard to see how to color graphs of chromatic number χ with $\alpha\chi$ colors in d rounds by using disjoint color palettes for clusters of different colors: For every $i \in [\alpha]$, the nodes in a cluster of color i use colors from the palette $\{i, \alpha + i, 2\alpha + i, \dots\}$. Each such cluster C is colored by having a leader node collect the entire topology of the cluster and then brute forcing an optimal coloring φ_C of the cluster. Since the graph has chromatic number χ, coloring φ_C uses at most χ colors. Finally, the leader broadcasts the coloring and assigns each node v in C the color $\alpha(\varphi_C(v) - 1) + i$. This results in a proper coloring with at most $\alpha\chi$ colors. In addition, the nodes do not require knowledge of χ in advance. This algorithm has for example been used by Barenboim [10] to compute a non-trivial approximate coloring in a constant number of rounds. Our algorithm is based on two new ideas that are outlined below.

New ingredient 1: New network decomposition algorithms. Network decomposition algorithms mostly focus on optimizing the product of α and d (as most applications of network decompositions require time proportional to αd) or on minimizing the number of cluster colors for a given maximum cluster diameter (e.g., [5, 10, 26, 38, 45]). We are interested in network decompositions with a *fixed* number of cluster colors α that is beyond our control and where we wish to optimize the value of d. By using existing clustering techniques [18, 26, 42, 45] with some minor adaptations, we give randomized and deterministic algorithms that, for any parameter $\varepsilon > 0$, compute a set of non-adjacent clusters such that the cluster diameter of each cluster is $\text{polylog}(n)/\varepsilon$ and the total number of unclustered nodes is at most εn. For every integer α, this can in turn be used to compute an (α, d)-network decomposition with $d = \tilde{O}(n^{1/\alpha})$.

Let us illustrate the idea for the case $\alpha = 2$. Setting $\varepsilon = 1/\sqrt{n}$, we compute a set of non-adjacent clusters of diameter $\tilde{O}(\sqrt{n})$ such that at most $O(\sqrt{n})$ nodes remain unclustered. The constructed clusters can all be colored with color 1 and the connected components of the unclustered nodes form the clusters of color 2. We thus obtain a $(2, \tilde{O}(\sqrt{n}))$-network decomposition in time $\tilde{O}(\sqrt{n})$ and one can therefore color graphs of chromatic number χ with 2χ colors in $\tilde{O}(\sqrt{n})$ rounds.

New ingredient 2: The hiding trick. We show how to reduce the number of colors used while keeping the round complexity the same. The main new idea is what we call the *hiding trick*: First we make sure that clusters of the same color are at least four hops apart; this can be achieved by running a network decomposition protocol on G^3. For simplicity, assume that we have a $(2, \tilde{O}(\sqrt{n}))$-network decomposition. Let C be a cluster of color 1. We first extend C to an extended cluster C_1 that includes all unclustered nodes that are adjacent to C. Next we find a proper χ-coloring of C_1 using colors $\{1, \dots, \chi\}$ by brute force. Finally, any boundary node of C_1 that has color χ is removed, thus yielding a cluster C_0 with $C \subseteq C_0 \subseteq C_1$. Note that C_0 is colored using at most χ colors and that and all boundary nodes of C_0 have a color different from χ. We have effectively *hidden the color χ inside the cluster*. Now we continue with the rest of the process and can safely use colors $\{\chi, \dots, 2\chi - 1\}$ to color the uncolored nodes in clusters of color 2. The end result is a proper $(2\chi - 1)$-coloring, and the running time is simply a constant times the cluster diameter $d = \tilde{O}(\sqrt{n})$. With this strategy, for instance, we can compute a 3-coloring of bipartite graphs in $\tilde{O}(\sqrt{n})$ rounds.

Theorem 2.1 is in essence a formalization and generalization of the hiding trick. We can hide one color in each cluster and therefore reuse one of the colors for all of the α cluster colors. This results in a coloring with $\alpha(\chi - 1) + 1$ colors. In addition, if one chooses the color palettes carefully, it is not necessary for the nodes to know χ beforehand. At first this may seem like an ad-hoc trick that cannot possibly be optimal—after all, we are saving only one color in each step. However, our matching lower bound shows that the hiding

trick is essentially the best that we can do in distributed coloring, even if we have access to quantum resources.

3.2 Non-signaling Lower Bounds

Prior work on classical lower bounds. As a warm-up, let us first see how one could prove a lower bound similar to Theorem 2.3 for classical models. For concreteness, let us focus on the case $\chi = 2$ and $c = 3$ in the det-LOCAL model. Then $\alpha = 2$, and we would like to show that 3-coloring bipartite graphs requires $\Omega(\sqrt{n})$ rounds.

Here we can make use of existential graph-theoretic arguments similar to the one already used by Linial [37]. Let $\mathcal{A}$ be a det-LOCAL algorithm that purportedly finds a 3-coloring in bipartite graphs in $T(n) = o(\sqrt{n})$ rounds. Now suppose that we are able to construct a graph G on n nodes with the following properties:

(1) G is locally bipartite: for any node v of G, the subgraph of G induced by the radius-$\Theta(\sqrt{n})$ neighborhood of v is bipartite.
(2) G is not 3-colorable: the chromatic number of G is at least 4.

The graph G is not bipartite, but (since $\mathcal{A}$ operates in the LOCAL model) we can apply $\mathcal{A}$ to G anyway and observe what happens. As the chromatic number of G is at least 4, certainly $\mathcal{A}$ cannot 3-color G; hence there has to be at least one node v such that in the local neighborhood X of v the output of $\mathcal{A}$ is not a valid 3-coloring. However, by assumption, the local neighborhood of v up to distance $\Theta(\sqrt{n})$ is bipartite, so with some cutting and pasting we can construct another graph $G' = (V, E')$ on the same set of nodes such that G' is bipartite and the graph induced by the radius-$T(n)$ neighborhood of X is the same in G and G'. Hence the output of $\mathcal{A}$ in X is the same in both graphs, which means $\mathcal{A}$ produces an invalid coloring in G' (which is bipartite) in the neighborhood X. The key point in this argument is the existence of the *cheating graph* G that "fools" $\mathcal{A}$ as $\mathcal{A}$ cannot locally tell the difference between G and the valid input graph G'.

New ingredient 1: Bogdanov's construction. To apply the proof strategy outlined above, we need a suitable construction of a cheating graph. It turns out we can make direct use of Bogdanov's [12] graph-theoretic work—this is a 10-year-old result, but so far this seems to have been a blind spot in the research community, and we are not aware of any lower bounds in the context of distributed graph algorithms that make use of it.

Together with our new algorithm from Theorem 2.1, this gives a near-complete characterization of the complexity of c-coloring χ-colorable graphs in det-LOCAL. A similar argument applies (with some adjustments) to the rand-LOCAL model as well; in particular, we can exploit the independence of the output distribution between well-separated subgraphs of the input graph to boost the failing probability of an algorithm (see the full version [20] for the details).

New ingredient 2: Defining cheating graphs for NS-LOCAL. Although proof techniques based on cheating graphs are commonly used in the context of det-LOCAL and rand-LOCAL, we stress the same line of reasoning *does not hold* in quantum-LOCAL or NS-LOCAL. In fact, in the context of NS-LOCAL new challenges emerge, which we discuss later in this section. Our new proof strategy overcomes these issues: it builds on the idea of cheating graphs, but it allows us to directly derive a lower bound for NS-LOCAL. To the best of our knowledge, this is the first work establishing that Linial's argument [37] can be adapted to more general non-signaling models. This is one of our main technical contributions.

We present a new definition of a cheating graph that is applicable in NS-LOCAL. Suppose we are interested in a locally checkable problem $\mathcal{P}$ in graph family $\mathcal{F}$.

Definition 3.1 (Cheating graph, informal version). Consider a sufficiently large $N > 0$. A graph G is a cheating graph for $(\mathcal{P}, \mathcal{F})$ if

(1) problem $\mathcal{P}$ is not solvable on G;
(2) for a suitable k, we can cover G with k subgraphs $G^{(1)}, \dots, G^{(k)}$ such that $\mathcal{P}$ is solvable over each of the graphs induced by their radius-$T(n)$ neighborhoods, where $n = |V(G)| \cdot N$;
(3) we can take any N subgraphs $G^{(x_1)}, \dots, G^{(x_N)}$ together with their radius-$T(n)$ neighborhoods, possibly with replacement, form their disjoint union $\tilde{G}$, and find a graph $H \in \mathcal{F}$ of size n that contains an induced subgraph isomorphic to $\tilde{G}$.

See [20] for the formal version. We show that the existence of graphs of arbitrarily large size with the above properties directly implies a lower bound equal to T to the problem $\mathcal{P}$ over the graph family $\mathcal{F}$ that holds also in NS-LOCAL.

We note that the precise requirements for k and N depend on T: to prove Theorem 2.5 we exploit the fact that when T is small we can afford a large k, while for Theorems 2.3 and 2.4 we deal with a large T, and then it will be important to construct a cheating graph with a constant k.

Our definition of a cheating graph is somewhat technical, but through three examples we demonstrate that this is indeed an effective way of proving lower bounds.

New ingredient 3: New analysis of Bogdanov's construction. While in the classical models we could directly apply Bogdanov [12] as a black box, this is no longer the case in NS-LOCAL. Nevertheless we show that the construction of [12] indeed gives a cheating graph for c-coloring χ-chromatic graph.

It is known that the graph constructed by [12] is locally χ-chromatic, but globally has chromatic number greater than c, implying property (1) in Definition 3.1. We further go through the details of the construction and prove that the graph also satisfies properties (2) and (3) from Definition 3.1—these are properties outside the scope of [12]. This then yields Theorem 2.3.

New ingredient 4: New analysis of quadrangulations of the Klein bottle. We make use of properties of quadrangulations of the Klein bottle [3, 43, 44] to construct a family of graphs that is locally grid-like but is not 3-colorable. We show that such quadrangulations give cheating graphs for 3-coloring grids, from which we obtain Theorem 2.4.

New ingredient 5: New analysis of Ramanujan graphs. We revisit the construction of Ramanujan graphs [39], that is, high-girth and high-chromatic regular graphs, which Linial [37] used in his lower-bound proof. Again, we show that it provides us with a cheating graph (Definition 3.1) for c-coloring trees, and Theorem 2.5 follows.

Discussion. While quadrangulations of the Klein bottle and Ramanujan graphs have been used in prior work to prove lower bounds for the classical models, by e.g. Aboulker et al. [1], Linial [37], we

remark that to our knowledge, this is the first time that the applicability of Bogdanov's graph-theoretic work [12] in the context of distributed computing and quantum computing lower bounds is recognized (in spite of it being a 10-year-old result).

We also note that the classical version of Theorem 2.4 by [15] uses fundamentally different proof techniques: the argument in [15] is primarily *algorithmic*, while our proof is primarily *graph-theoretic*. The algorithmic proof from the prior work seems to be fundamentally incompatible with the NS-LOCAL model, while the graph-theoretic proof also yields a lower bound for NS-LOCAL. This suggests a general blueprint for lifting prior lower bounds from det-LOCAL or rand-LOCAL to NS-LOCAL: (1) re-prove the previous result using existential graph-theoretic arguments, and (2) apply the cheating graph idea to lift it to NS-LOCAL.

While the proof technique developed in this work is applicable in many graph problems, there are also problems for which cheating graphs do not exist (e.g., sinkless orientation on 2-regular graphs). An open research direction is developing proof strategies that can be used to derive NS-LOCAL lower bounds for those cases.

4 LOWER BOUND TECHNIQUE IN MORE DETAIL

Here we give a more formal description of the lower bound technique illustrated above.

Fix a sufficiently large $N > 0$. Consider any locally checkable problem $\mathcal{P}$ over some graph family $\mathcal{F}$. We want to show that, whenever a cheating graph (Definition 3.1) for the pair $(\mathcal{P}, \mathcal{F})$ exists, any T-round algorithm solving the problem has failing probability at least $1 - (1 - 1/k)^N$, where k is the size of the subgraph cover of the cheating graph.

Let G be the cheating graph. For simplicity, we can think of $\mathcal{P}$ as the 3-coloring problem, and $\mathcal{F}$ to be the family of bipartite graphs. Provided that $\mathcal{F}$ respects some natural properties, properties (1) and (2) from Definition 3.1 ensure that we can get a lower bound in rand-LOCAL. Indeed, assume there is a T-round randomized algorithm $\mathcal{A}$ that 3-colors bipartite graphs. Clearly, $\mathcal{A}$ fails to 3-color G with probability 1. Hence, there is an $i^\star \in [k]$ such that the failing probability of $\mathcal{A}$ over $G^{(i^\star)}$ is at least $1/k$. Hence, $\mathcal{A}$ will fail on any bipartite graph of at most $n = |V(G)| \cdot N$ nodes containing an induced subgraph isomorphic to the radius-$T(n)$ neighborhood of $G^{(i^\star)}$ with probability $1/k$. If k is not small enough (e.g., $k = w(1)$), the failing probability tends to 0. In rand-LOCAL we can amplify the failure probability as follows: Suppose $\mathcal{F}$ contains a graph H_N of at most $n = |V(G)| \cdot N$ nodes that contains, as subgraphs, N disjoint copies of the radius-$T(n)$ neighborhood of $G^{(i^\star)}$ in G. This is always possible if $\mathcal{F}$ is the family of bipartite graphs. By independence, the failing probability of $\mathcal{A}$ over H_N is at least $1 - (1 - 1/k)^N$ (see Fig. 2a). Hence, such an algorithm cannot exist. The property that $\mathcal{F}$ contains H_N is reasonable for many natural problems (e.g., c-coloring χ-chromatic graphs for all combinations of c and χ) where, given a graph for which the problem is solvable, one can connect disjoint copies of the graphs and obtain a solvable instance of the problem.

However, as we anticipated, in the NS-LOCAL model some issues arise:

(i) If two graphs G and H have different sizes, then even if they share two identical subgraphs G' and H' with isomorphic radius-T neighborhoods, a non-signaling output distribution is not guaranteed to be identical over G' and H'; this is due to the no-cloning principle [21, 41, 49].

(ii) If we look at the output distributions for two subsets of nodes X and Y, then even if X and Y are far from each other, we cannot assume that the outputs of these subsets are independent.

Issue (i) puts some limits on the choice of the graph used to "fool" the algorithm, while issue (ii) makes it necessary to deal with possible dependencies among different parts of the input graph. Such issues are the reason why we require the cheating graph to satisfy property (3) in Definition 3.1.

To solve issue (i), we consider a graph that is the disjoint union of N copies $G_1, \ldots, G_n$ of the cheating graph G: such graph has $|V(G)| \cdot N$ vertices, exactly the same as the graph of property (3) from Definition 3.1. Consider now any NS-LOCAL algorithm $\mathcal{A}$ that 3-colors bipartite graphs with locality T, and apply it to the graph $\bigsqcup_{i \in [N]} G_i$. Clearly, $\mathcal{A}$ will fail to solve the problem in each G_j with probability 1. At this point, we cannot continue as before: while it is true that in each G_j we can find an index $i^\star$ such that the probability of $\mathcal{A}$ failing in $G_j^{(i^\star)}$ is at least $1/k$, we cannot use independence to increase the failing probability.

Property (3) ensures that, for a sufficiently large N, and for any sequence of indices $\mathbf{x} = (x_1, \ldots, x_N) \in [k]^N$, there exists a graph $H_{\mathbf{x}}$ of size $|V(G)| \cdot N$ that contains an induced subgraph isomorphic to the disjoint union of the radius-T neighborhoods of $G^{(x_1)}, \ldots, G^{(x_N)}$. However, correlations among these subgraphs of $H_{\mathbf{x}}$ might hold. To overcome this issue, we need to consider all possible sequences of subgraphs $G_1^{(x_1)}, \ldots, G_N^{(x_N)}$ at the same time, where $\mathbf{x} = (x_1, \ldots, x_N) \in [k]^N$ (see Fig. 2b). Fix a total order for the elements in $[k]^N$, and let its ordered elements be $\mathbf{x}_1, \ldots, \mathbf{x}_{k^N}$. Let $\mathcal{F}_{\mathbf{x}_j}$ be the event that $\mathcal{A}$ fails in $G_i^{(z_i)}$, where z_i is the i-th element of $\mathbf{x}_j$, for each $i = 1, \ldots, N$. Furthermore, for each index $\mathbf{x} \in [k]^N$, let $\mathcal{I}_{\mathbf{x}}$ be the set of all indices $\mathbf{y} \in [k]^N$ such that $\mathbf{y}$ and $\mathbf{x}$ share the same element at the i-th position, for some i, i.e., $\mathbf{x}(i) = \mathbf{y}(i)$. Notice that, for $\mathbf{x} = (x_1, \ldots, x_N)$, $\bigcup_{\mathbf{y} \in \mathcal{I}_{\mathbf{x}}} \mathcal{F}_{\mathbf{y}}$ describes the event that there is an $i \in [N]$ such that $\mathcal{A}$ fails on $G_i^{(x_i)}$.

We claim that there exists an $\mathbf{x}^\star \in [k]^N$ such that

$$\Pr\left[\bigcup_{\mathbf{y} \in \mathcal{I}_{\mathbf{x}^\star}} \mathcal{F}_{\mathbf{y}}\right] \geq 1 - \left(1 - \frac{1}{k}\right)^N,$$

implying that the dependencies behave "well enough", hence solving issue (ii). We proceed by contradiction: we assume that, for all $\mathbf{x} \in [k]^N$,

$$\Pr\left[\bigcup_{\mathbf{y} \in \mathcal{I}_{\mathbf{x}}} \mathcal{F}_{\mathbf{y}}\right] < 1 - \left(1 - \frac{1}{k}\right)^N.$$

While

$$\Pr\left[\bigcup_{\mathbf{x} \in [k]^N} \mathcal{F}_{\mathbf{x}}\right] = 1,$$

the events in $\{\mathcal{F}_{\mathbf{x}}\}_{\mathbf{x} \in [k]^N}$ are not disjoint and the sum of their probability is not 1. To better deal with the math, we define $\mathcal{E}_{\mathbf{x}_1} =$

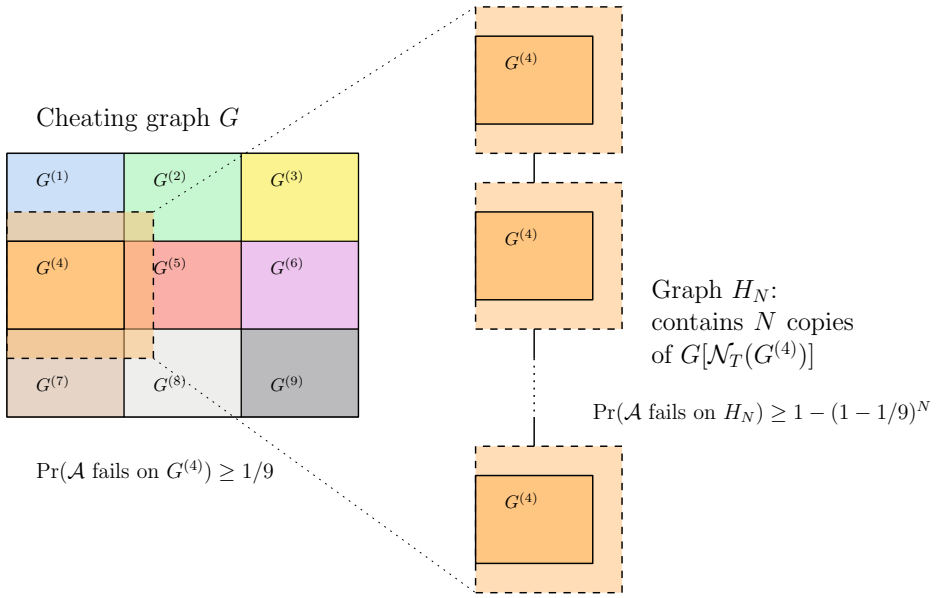

(a) Construction for the rand-LOCAL model. For any $T(n)$-round algorithm $\mathcal{A}$ solving the problem, there is an $i^\star \in [9]$ (in the figure, $i^\star = 4$) such that $\Pr[\mathcal{A}$ fails in $G^{(i^\star)}] \geq 1/9$. Then, $\Pr[\mathcal{A}$ fails on $H_N] \geq 1-(1-1/9)^N$ by independence, where H_N is an admissible instance. As long as $|V(H_N)| \leq n$, this gives the lower bound.

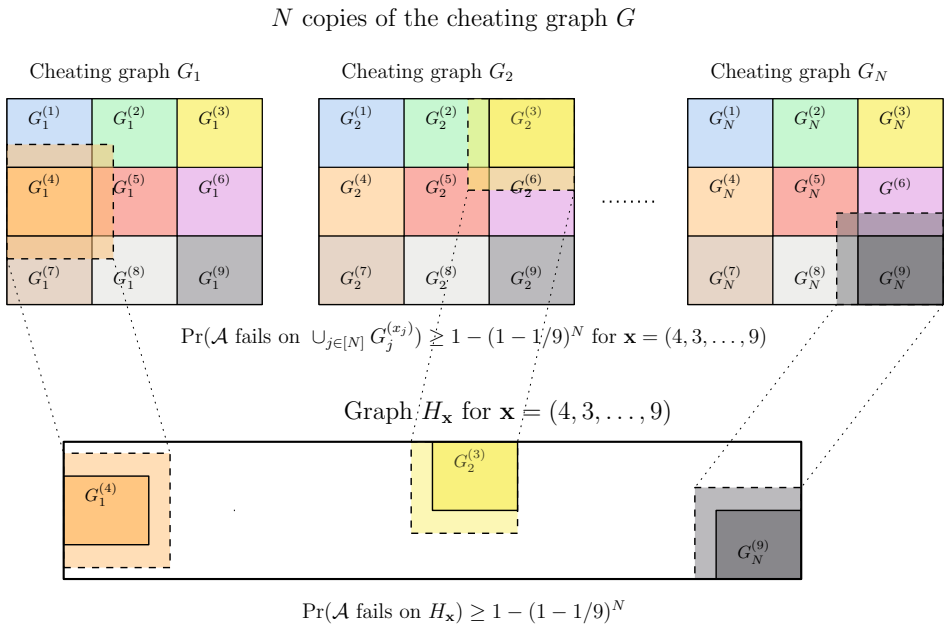

(b) Construction for the NS-LOCAL model. We start with N copies $G_1, \ldots, G_N$ of G and consider their disjoint union. We prove that, in this specific graph, there is already a combination of indices $\mathbf{x} = (x_1, \ldots, x_N) \in [9]^N$ (in the figure, $\mathbf{x} = (4,3,\ldots,9)$) for which $\Pr[\mathcal{A}$ fails on $\bigcup_{j\in[N]} G_j^{(x_j)}] \geq 1-(1-1/9)^N$. Then, property (3) of Definition 3.1 ensures that we can construct an admissible instance $H_{\mathbf{x}}$ as shown in the figure, with $|V(H_{\mathbf{x}})| = n$. By the properties of the NS-LOCAL model, since $H_{\mathbf{x}}$ and $\bigsqcup_{i\in[N]} G_i$ share the same local view around $\bigcup_{j\in[N]} G_j^{(x_j)}$, $\mathcal{A}$ fails on $H_{\mathbf{x}}$ too with at least the same probability.

Figure 2: Illustration of the lower-bound argument based on the cheating graph G. For $n = |V(G)| \cdot N$, the problem is solvable in each $T(n)$-radius neighborhood of $G^{(i)}$, $i \in [9]$, but not on G.

$\mathcal{F}_{\mathbf{x}_1}$ and, recursively, we define

$$\mathcal{E}_{\mathbf{x}_j} = \mathcal{F}_{\mathbf{x}_j} \setminus \bigcup_{i=1}^{j-1} \mathcal{E}_{\mathbf{x}_i}.$$

Clearly the events in $\{\mathcal{E}_{\mathbf{x}}\}_{\mathbf{x}\in[k]^N}$ are pairwise disjoint: furthermore, it holds that $\sum_{\mathbf{x}\in[k]^N} \Pr[\mathcal{E}_{\mathbf{x}}] = 1$ as

$$\bigcup_{\mathbf{x}\in[k]^N} \mathcal{E}_{\mathbf{x}} = \bigcup_{\mathbf{x}\in[k]^N} \mathcal{F}_{\mathbf{x}}.$$

For each $\mathbf{x} \in [k]^N$, it trivially holds that

$$\sum_{\mathbf{y}\in\mathcal{I}_{\mathbf{x}}} \Pr[\mathcal{E}_{\mathbf{y}}] + \sum_{\mathbf{y}\in[k]^N\setminus\mathcal{I}_{\mathbf{x}}} \Pr[\mathcal{E}_{\mathbf{y}}] = 1.$$

Moreover, for each $\mathbf{x} \in [k]^N$, we have $\Pr[\mathcal{E}_{\mathbf{x}}] \leq \Pr[\mathcal{F}_{\mathbf{x}}]$ as $\mathcal{E}_{\mathbf{x}} \subseteq \mathcal{F}_{\mathbf{x}}$, hence

$$\sum_{\mathbf{y}\in\mathcal{I}_{\mathbf{x}}} \Pr[\mathcal{E}_{\mathbf{y}}] < 1 - \left(1 - \frac{1}{k}\right)^N.$$

Thus,

$$\sum_{\mathbf{y}\in[k]^N\setminus\mathcal{I}_{\mathbf{x}}} \Pr[\mathcal{E}_{\mathbf{y}}] > \left(1 - \frac{1}{k}\right)^N.$$

It follows that

$$\sum_{\mathbf{x}\in[k]^N} \sum_{\mathbf{y}\in[k]^N\setminus\mathcal{I}_{\mathbf{x}}} \Pr[\mathcal{E}_{\mathbf{y}}] > k^N \left(1 - \frac{1}{k}\right)^N = (k-1)^N.$$

Also, notice that for each $\mathbf{y} \in [k]^N$ the cardinality of the set $\{\mathbf{x} \in [k]^N \mid \mathbf{y} \in [k]^N \setminus \mathcal{I}_{\mathbf{x}}\}$ is $(k-1)^N$. Hence,

$$\begin{aligned}\sum_{\mathbf{x}\in[k]^N} \sum_{y\in[k]^N\setminus\mathcal{I}_{\mathbf{x}}} \Pr[\mathcal{E}_{\mathbf{y}}] &= \sum_{\mathbf{y}\in[k]^N} \sum_{\substack{\mathbf{x}\in[k]^N:\\ \mathbf{y}\in[k]^N\setminus\mathcal{I}_{\mathbf{x}}}} \Pr[\mathcal{E}_{\mathbf{y}}] \\ &= (k-1)^N \sum_{\mathbf{y}\in[k]^N} \Pr[\mathcal{E}_{\mathbf{y}}] \\ &= (k-1)^N,\end{aligned}$$

reaching a contradiction. Thus, there exists an $\mathbf{x}^\star \in [k]^N$ such that

$$\Pr\left[\bigcup_{\mathbf{y}\in\mathcal{I}_{\mathbf{x}^\star}} \mathcal{F}_{\mathbf{y}}\right] \geq 1 - \left(1 - \frac{1}{k}\right)^N.$$

Property (3) of Definition 3.1 ensures that there is a graph $H_{\mathbf{x}^\star} \in \mathcal{F}$ of n nodes such that $H_{\mathbf{x}^\star}$ contains

$$\mathbf{G}_{\mathbf{x}^\star} = \bigcup_{i\in[N]} G_i^{(\mathbf{x}_i^\star)}$$

as an induced subgraph, and $H_{\mathbf{x}^\star}$ and $\bigsqcup_{i\in[N]} G_i$ share the same radius-$T(n)$ neighborhood around $\mathbf{G}_{\mathbf{x}^\star}$. By the property of the NS-LOCAL model, we get that the failing probability of $\mathcal{A}$ on $H_{\mathbf{x}^\star}$ is at least $1 - (1 - 1/k)^N$.

ACKNOWLEDGMENTS

We would like to thank Dennis Olivetti for pointing out the work by Bogdanov [12] on highly chromatic graphs with small local chromatic number. We are grateful to Sebastian Brandt and Michele Pernice for investigating and contributing to the understanding of topological properties of some structures used in preliminary versions of our results. Furthermore, we thank Sameep Dahal for helping to deal with dependencies arisen in the lower bound technique for the non-signaling model, and we thank Darya Melnyk, Shreyas Pai, Chetan Gupta, Alkida Balliu, Amirreza Akbari, Yannic Maus, and all participants of the Distributed Graph Algorithms Workshop (February 2023, Freiburg, Germany) for the helpful discussions during the long development of this work. We are also thankful to Armin Biere, Austin Buchanan, Bill Cook, Stefano Gualandi, and Bernardo Subercaseaux for discussions and advice related to our computational investigations of local and global chromatic numbers. We also thank Elie Wolfe, Harry Buhrman and Paolo Perinotti for helpful discussions on the relation between the non-signaling LOCAL model and physical principles. We made use of e.g. Plingeling [11], Treengeling [11], Gimsatul [24], MapleSAT [36], and PySAT [31] in our computational experiments, and we also wish to acknowledge CSC – IT Center for Science, Finland, for computational resources.

This work was supported in part by the Research Council of Finland, Grant 333837 and by the German Research Foundation (DFG), Grant 491819048. Francesco d'Amore was supported by MUR FARE 2020 Project PAReCoDi CUP J43C22000970001. François Le Gall was supported by JSPS KAKENHI grants Nos. JP20H05966, JP20H04139 and MEXT Q-LEAP grant No. JPMXS0120319794. Augusto Modanese was supported by the Helsinki Institute of Information Technology (HIIT). Marc-Olivier Renou was supported by INRIA and CIEDS in the Action Exploratoire project DEPARTURE. Xavier Coiteux-Roy was supported by a Postdoc.Mobility fellowship from the Swiss National Science Foundation (SNSF).

REFERENCES

[1] Pierre Aboulker, Marthe Bonamy, Nicolas Bousquet, and Louis Esperet. 2019. Distributed Coloring in Sparse Graphs with Fewer Colors. *The Electronic Journal of Combinatorics* 26, 4 (2019), 4. https://doi.org/10.37236/8395

[2] Amirreza Akbari, Navid Eslami, Henrik Lievonen, Darya Melnyk, Joona Särkijärvi, and Jukka Suomela. 2023. Locality in Online, Dynamic, Sequential, and Distributed Graph Algorithms. In *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023, July 10-14, 2023, Paderborn, Germany (LIPIcs, Vol. 261)*, Kousha Etessami, Uriel Feige, and Gabriele Puppis (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 10:1–10:20. https://doi.org/10.4230/LIPICS.ICALP.2023.10

[3] Dan Archdeacon, Joan P. Hutchinson, Atsuhiro Nakamoto, Seiya Negami, and Katsuhiro Ota. 2001. Chromatic numbers of quadrangulations on closed surfaces. *Journal of Graph Theory* 37, 2 (2001), 100–114. https://doi.org/10.1002/JGT.1005

[4] Heger Arfaoui and Pierre Fraigniaud. 2014. What can be computed without communications? *SIGACT News* 45, 3 (2014), 82–104. https://doi.org/10.1145/2670418.2670440

[5] Baruch Awerbuch, Andrew V. Goldberg, Michael Luby, and Serge A. Plotkin. 1989. Network Decomposition and Locality in Distributed Computation. In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*. IEEE Computer Society, 364–369. https://doi.org/10.1109/SFCS.1989.63504

[6] Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. 2019. Lower Bounds for Maximal Matchings and Maximal Independent Sets. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, David Zuckerman (Ed.). IEEE Computer Society, 481–497. https://doi.org/10.1109/FOCS.2019.00037

[7] Alkida Balliu, Sebastian Brandt, Dennis Olivetti, and Jukka Suomela. 2018. Almost Global Problems in the LOCAL Model. In *32nd International Symposium on Distributed Computing, DISC 2018, New Orleans, LA, USA, October 15-19, 2018 (LIPIcs, Vol. 121)*, Ulrich Schmid and Josef Widder (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 9:1–9:16. https://doi.org/10.4230/LIPICS.DISC.2018.9

[8] Alkida Balliu, Sebastian Brandt, Dennis Olivetti, and Jukka Suomela. 2020. How much does randomness help with locally checkable problems?. In *PODC '20: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, August 3-7, 2020*, Yuval Emek and Christian Cachin (Eds.). ACM, 299–308. https://doi.org/10.1145/3382734.3405715

[9] Alkida Balliu, Juho Hirvonen, Janne H. Korhonen, Tuomo Lempiäinen, Dennis Olivetti, and Jukka Suomela. 2018. New classes of distributed time complexity. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, Ilias Diakonikolas, David Kempe, and Monika Henzinger (Eds.). ACM, 1307–1318. https://doi.org/10.1145/3188745.3188860

[10] Leonid Barenboim. 2012. On the Locality of Some NP-Complete Problems. In *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 7392)*, Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer (Eds.). Springer, 403–415. https://doi.org/10.1007/978-3-642-31585-5_37

[11] Armin Biere. 2017. CaDiCaL, Lingeling, Plingeling, Treengeling, YalSAT Entering the SAT Competition 2017. In *Proc. of SAT Competition 2017 – Solver and Benchmark Descriptions (Department of Computer Science Series of Publications B, Vol. B-2017-1)*, Tomáš Balyo, Marijn Heule, and Matti Järvisalo (Eds.). University of Helsinki, 14–15.

[12] Ilya I. Bogdanov. 2013. Examples of topologically highly chromatic graphs with locally small chromatic number. *CoRR* abs/1311.2844 (2013). https://doi.org/10.48550/arXiv.1311.2844 arXiv:1311.2844

[13] Sebastian Brandt. 2019. An Automatic Speedup Theorem for Distributed Problems. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, Peter Robinson and Faith Ellen (Eds.). ACM, 379–388. https://doi.org/10.1145/3293611.3331611

[14] Sebastian Brandt, Orr Fischer, Juho Hirvonen, Barbara Keller, Tuomo Lempiäinen, Joel Rybicki, Jukka Suomela, and Jara Uitto. 2016. A lower bound for the distributed Lovász local lemma. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, Daniel Wichs and Yishay Mansour (Eds.). ACM, 479–488. https://doi.org/10.1145/2897518.2897570

[15] Sebastian Brandt, Juho Hirvonen, Janne H. Korhonen, Tuomo Lempiäinen, Patric R. J. Östergård, Christopher Purcell, Joel Rybicki, Jukka Suomela, and Przemyslaw Uznanski. 2017. LCL Problems on Grids. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017, Washington, DC, USA, July 25-27, 2017*, Elad Michael Schiller and Alexander A. Schwarzmann (Eds.). ACM, 101–110. https://doi.org/10.1145/3087801.3087833

[16] Keren Censor-Hillel, Orr Fischer, François Le Gall, Dean Leitersdorf, and Rotem Oshman. 2022. Quantum Distributed Algorithms for Detection of Cliques. In *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA (LIPIcs, Vol. 215)*, Mark Braverman (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 35:1–35:25. https://doi.org/10.4230/LIPICS.ITCS.2022.35

[17] Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. 2016. An Exponential Separation between Randomized and Deterministic Complexity in the LOCAL Model. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, Irit Dinur (Ed.). IEEE Computer Society, 615–624. https://doi.org/10.1109/FOCS.2016.72

[18] Yi-Jun Chang and Zeyong Li. 2023. The Complexity of Distributed Approximation of Packing and Covering Integer Linear Programs. In *Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing, PODC 2023, Orlando, FL, USA, June 19-23, 2023*, Rotem Oshman, Alexandre Nolin, Magnús M. Halldórsson, and Alkida Balliu (Eds.). ACM, 32–43. https://doi.org/10.1145/3583668.3594562

[19] Yi-Jun Chang and Seth Pettie. 2019. A Time Hierarchy Theorem for the LOCAL Model. *SIAM J. Comput.* 48, 1 (2019), 33–69. https://doi.org/10.1137/17M1157957

[20] 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. 2023. No distributed quantum advantage for approximate graph coloring. *CoRR* abs/2307.09444 (2023). https://doi.org/10.48550/ARXIV.2307.09444 arXiv:2307.09444

[21] Giacomo Mauro D'Ariano, Giulio Chiribella, and Paolo Perinotti. 2016. *Quantum Theory from First Principles: An Informational Approach*. Cambridge University Press. https://doi.org/10.1017/9781107338340

[22] Michael Elkin, Hartmut Klauck, Danupon Nanongkai, and Gopal Pandurangan. 2014. Can quantum communication speed up distributed computation?. In *ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, France, July 15-18, 2014*, Magnús M. Halldórsson and Shlomi Dolev (Eds.). ACM, 166–175. https://doi.org/10.1145/2611462.2611488

[23] Manuela Fischer and Mohsen Ghaffari. 2017. Sublogarithmic Distributed Algorithms for Lovász Local Lemma, and the Complexity Hierarchy. In *31st International Symposium on Distributed Computing, DISC 2017, October 16-20, 2017, Vienna, Austria (LIPIcs, Vol. 91)*, Andréa W. Richa (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 18:1–18:16. https://doi.org/10.4230/LIPICS.DISC.2017.18

[24] Mathias Fleury and Armin Biere. 2022. Scalable Proof Producing Multi-Threaded SAT Solving with Gimsatul through Sharing instead of Copying Clauses. *CoRR* abs/2207.13577 (2022). https://doi.org/10.48550/ARXIV.2207.13577 arXiv:2207.13577

[25] Cyril Gavoille, Adrian Kosowski, and Marcin Markiewicz. 2009. What Can Be Observed Locally?. In *Distributed Computing, 23rd International Symposium, DISC 2009, Elche, Spain, September 23-25, 2009. Proceedings (Lecture Notes in Computer Science, Vol. 5805)*, Idit Keidar (Ed.). Springer, 243–257. https://doi.org/10.1007/978-3-642-04355-0_26

[26] Mohsen Ghaffari, Christoph Grunau, Bernhard Haeupler, Saeed Ilchi, and Václav Rozhon. 2023. Improved Distributed Network Decomposition, Hitting Sets, and Spanners, via Derandomization. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, Nikhil Bansal and Viswanath Nagarajan (Eds.). SIAM, 2532–2566. https://doi.org/10.1137/1.9781611977554.CH97

[27] Mohsen Ghaffari, David G. Harris, and Fabian Kuhn. 2018. On Derandomizing Local Distributed Algorithms. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, Mikkel Thorup (Ed.). IEEE Computer Society, 662–673. https://doi.org/10.1109/FOCS.2018.00069

[28] Mohsen Ghaffari and Hsin-Hao Su. 2017. Distributed Degree Splitting, Edge Coloring, and Orientations. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, Philip N. Klein (Ed.). SIAM, 2505–2523. https://doi.org/10.1137/1.9781611974782.166

[29] Alexander E. Holroyd, Tom Hutchcroft, and Avi Levy. 2018. Finitely dependent cycle coloring. *Electronic Communications in Probability* 23 (2018). https://doi.org/10.1214/18-ecp118

[30] Alexander E. Holroyd and Thomas M. Liggett. 2016. Finitely Dependent Coloring. *Forum of Mathematics, Pi* 4, Article e9 (2016). https://doi.org/10.1017/fmp.2016.7

[31] Alexey Ignatiev, António Morgado, and João Marques-Silva. 2018. PySAT: A Python Toolkit for Prototyping with SAT Oracles. In *Theory and Applications of Satisfiability Testing - SAT 2018 - 21st International Conference, SAT 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 9-12, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 10929)*, Olaf Beyersdorff and Christoph M. Wintersteiger (Eds.). Springer, 428–437. https://doi.org/10.1007/978-3-319-94144-8_26

[32] Taisuke Izumi and François Le Gall. 2019. Quantum Distributed Algorithm for the All-Pairs Shortest Path Problem in the CONGEST-CLIQUE Model. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, Peter Robinson and Faith Ellen (Eds.). ACM, 84–93. https://doi.org/10.1145/3293611.3331628

[33] Taisuke Izumi, François Le Gall, and Frédéric Magniez. 2020. Quantum Distributed Algorithm for Triangle Finding in the CONGEST Model. In *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France (LIPIcs, Vol. 154)*, Christophe Paul and Markus Bläser (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 23:1–23:13. https://doi.org/10.4230/LIPICS.STACS.2020.23

[34] François Le Gall and Frédéric Magniez. 2018. Sublinear-Time Quantum Computation of the Diameter in CONGEST Networks. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*. ACM, 337–346. https://doi.org/10.1145/3212734.3212744

[35] François Le Gall, Harumichi Nishimura, and Ansis Rosmanis. 2019. Quantum Advantage for the LOCAL Model in Distributed Computing. In *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany (LIPIcs, Vol. 126)*, Rolf Niedermeier and Christophe Paul (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 49:1–49:14. https://doi.org/10.4230/LIPICS.STACS.2019.49

[36] Jia Hui Liang, Vijay Ganesh, Pascal Poupart, and Krzysztof Czarnecki. 2016. Learning Rate Based Branching Heuristic for SAT Solvers. In *Theory and Applications of Satisfiability Testing - SAT 2016 - 19th International Conference, Bordeaux, France, July 5-8, 2016, Proceedings (Lecture Notes in Computer Science, Vol. 9710)*, Nadia Creignou and Daniel Le Berre (Eds.). Springer, 123–140. https://doi.org/10.1007/978-3-319-40970-2_9

[37] Nathan Linial. 1992. Locality in Distributed Graph Algorithms. *SIAM J. Comput.* 21, 1 (1992), 193–201. https://doi.org/10.1137/0221015

[38] Nathan Linial and Michael E. Saks. 1993. Low diameter graph decompositions. *Combinatorica* 13, 4 (1993), 441–454. https://doi.org/10.1007/BF01303516

[39] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. 1988. Ramanujan graphs. *Combinatorica* 8, 3 (1988), 261–277. https://doi.org/10.1007/BF02126799

[40] Frédéric Magniez and Ashwin Nayak. 2022. Quantum Distributed Complexity of Set Disjointness on a Line. *ACM Transactions on Computation Theory* 14, 1 (2022), 5:1–5:22. https://doi.org/10.1145/3512751

[41] Ll. Masanes, A. Acin, and N. Gisin. 2006. General properties of nonsignaling theories. *Physical Review A* 73, 1, Article 012112 (2006). https://doi.org/10.1103/physreva.73.012112

[42] Gary L. Miller, Richard Peng, and Shen Chen Xu. 2013. Parallel graph decompositions using random shifts. In *25th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '13, Montreal, QC, Canada - July 23 - 25, 2013*, Guy E. Blelloch and Berthold Vöcking (Eds.). ACM, 196–203. https://doi.org/10.1145/2486159.2486180

[43] Bojan Mohar and Paul D. Seymour. 2002. Coloring Locally Bipartite Graphs on Surfaces. *Journal of Combinatorial Theory, Series B* 84, 2 (2002), 301–310. https://doi.org/10.1006/JCTB.2001.2086

[44] Bojan Mohar, Gábor Simonyi, and Gábor Tardos. 2013. Local chromatic number of quadrangulations of surfaces. *Combinatorica* 33, 4 (2013), 467–495. https://doi.org/10.1007/S00493-013-2771-Y

[45] Václav Rozhon and Mohsen Ghaffari. 2020. Polylogarithmic-time deterministic network decomposition and distributed derandomization. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy (Eds.). ACM, 350–363. https://doi.org/10.1145/3357713.3384298

[46] Jukka Suomela. 2023. *Open problems related to locality in distributed graph algorithms.* https://jukkasuomela.fi/open/

[47] Joran van Apeldoorn and Tijn de Vos. 2022. A Framework for Distributed Quantum Queries in the CONGEST Model. In *PODC '22: ACM Symposium on Principles of Distributed Computing, Salerno, Italy, July 25 - 29, 2022*, Alessia Milani and Philipp Woelfel (Eds.). ACM, 109–119. https://doi.org/10.1145/3519270.3538413

[48] ChengSheng Wang, Xudong Wu, and Penghui Yao. 2022. Complexity of Eccentricities and All-Pairs Shortest Paths in the Quantum CONGEST Model. *CoRR* abs/2206.02766 (2022). https://doi.org/10.48550/ARXIV.2206.02766 arXiv:2206.02766

[49] W. K. Wootters and W. H. Zurek. 1982. A single quantum cannot be cloned. *Nature* 299, 5886 (1982), 802–803. https://doi.org/10.1038/299802a0

[50] Xudong Wu and Penghui Yao. 2022. Quantum Complexity of Weighted Diameter and Radius in CONGEST Networks. In *PODC '22: ACM Symposium on Principles of Distributed Computing, Salerno, Italy, July 25 - 29, 2022*, Alessia Milani and Philipp Woelfel (Eds.). ACM, 120–130. https://doi.org/10.1145/3519270.3538441

Received 13-NOV-2023; accepted 2024-02-11