# Polylogarithmic-Time Deterministic Network Decomposition and Distributed Derandomization

Václav Rozhoň
ETH Zurich
rozhonv@student.ethz.ch

Mohsen Ghaffari
ETH Zurich
ghaffari@inf.ethz.ch

## Abstract

We present a simple polylogarithmic-time deterministic distributed algorithm for network decomposition. This improves on a celebrated $2^{O(\sqrt{\log n})}$-time algorithm of Panconesi and Srinivasan [STOC'92] and settles a central and long-standing question in distributed graph algorithms. It also leads to the first polylogarithmic-time deterministic distributed algorithms for numerous other problems, hence resolving several well-known and decades-old open problems, including Linial's question about the deterministic complexity of maximal independent set [FOCS'87; SICOMP'92]—which had been called the most outstanding problem in the area.

The main implication is a more general distributed derandomization theorem: Put together with the results of Ghaffari, Kuhn, and Maus [STOC'17] and Ghaffari, Harris, and Kuhn [FOCS'18], our network decomposition implies that

P-RLOCAL = P-LOCAL.

That is, for any problem whose solution can be checked deterministically in polylogarithmic-time, any polylogarithmic-time randomized algorithm can be derandomized to a polylogarithmic-time deterministic algorithm. Informally, for the standard first-order interpretation of efficiency as polylogarithmic-time, distributed algorithms do not need randomness for efficiency.

By known connections, our result leads also to substantially faster *randomized* distributed algorithms for a number of well-studied problems including $(\Delta+1)$-coloring, maximal independent set, and Lovász Local Lemma, as well as *massively parallel* algorithms for $(\Delta + 1)$-coloring.

# 1 Introduction

We present a polylogarithmic-time deterministic distributed algorithm for network decomposition. This leads to substantially faster *deterministic* and *randomized* algorithms for many well-studied problems in distributed graph algorithms, as well as a general and efficient distributed derandomization theorem. These resolve several central open problems in the area.

## 1.1 Background and State of the Art

**Model**: We work with the standard model of distributed computing called LOCAL [Lin87, Lin92]: The communication network is abstracted as an $n$-node graph $G = (V, E)$, with one processor on each node $v \in V$ which has a unique $\Theta(\log n)$-bit identifier. Communication happens in synchronous rounds, where per round each node can send one message, of potentially unbounded size, to each neighbor. In the CONGEST variant of the model [Pel00], each message can have $O(\log n)$ bits. At the beginning, each processor knows only its neighbors, and some estimates of global parameters, e.g., a polynomial upper bound on $n$. At the end, each processor should know its own part of the output, e.g., its color in the vertex coloring problem.

**State of the art**: Prior to this work, the state of the art in distributed graph algorithms exhibited a significant (often nearly-exponential) gap between randomized and deterministic distributed algorithms. This gap constituted one of the foundational and long-standing questions in distributed algorithms. A well-known special case is an open question of Linial [Lin87, Lin92] about the maximal independent set (MIS) problem:

> "*can it* [MIS] *always be found* [deterministically] *in polylogarithmic time?*"

This has been described as "probably the most outstanding open problem in the area" [BE13, Open Problem 11.2]. Prior to our work, the best known deterministic algorithm had a round complexity of $2^{O(\sqrt{\log n})}$, by Panconesi and Srinivasan [PS92]. This should be contrasted with the beautiful $O(\log n)$-time randomized algorithms of Luby [Lub86] and Alon, Babai, and Itai [ABI86].

There is an abundance of similar open questions about obtaining polylogarithmic-time deterministic algorithms for other graph problems that admit polylogarithmic-time randomized algorithms; this includes $(\Delta + 1)$-coloring, Lovász Local Lemma, defective colorings, hypergraph matching, sparse neighborhood covers, etc. Indeed, in the Conclusion and Open Problems chapter of their 2013 book, Barenboim and Elkin [BE13, Chapter 11] write:

> "*Perhaps the most fundamental open problem in this field is to understand the power and limitations of randomization.*"

They then continue to ask for a general derandomization technique:

> ***Open Problem 11.1*** *Develop a general derandomization technique for the distributed message-passing model.*

This generic open problem is followed by 16 concrete open problems, 7 of which ask for polylogarithmic-time (sometimes just called efficient) deterministic algorithms for various graphs problems that are known to admit efficient randomized algorithms. We note that a few of these concrete open problems were well-known, and they had been mentioned throughout the literature since the 1990s.

## 1.2 Our Contribution

In this paper, we answer all the concrete questions mentioned above by providing the first polylogarithmic-time deterministic algorithms for them. In fact, we show a more general distributed derandomization theorem, which proves the following:

**Theorem 1.1** (**LOCAL Derandomization Theorem**)**.** *We have*

$$\textsf{P-LOCAL} = \textsf{P-RLOCAL}.$$

*Here,* P-LOCAL *denotes the family of locally checkable problems[1] that can be solved by deterministic algorithms in* $\mathrm{poly}(\log n)$ *rounds of the* LOCAL *model in* $n$-*node graphs and* P-RLOCAL *denotes the family of* locally checkable problems *that can be solved by randomized algorithms in* $\mathrm{poly}(\log n)$ *rounds of the* LOCAL *model, with success probability* $1 - 1/n$.

Informally, if we follow the standard of viewing a $\mathrm{poly}(\log n)$-round algorithm as *efficient*[2] (see e.g. [BE13, Lin92, PS92]), Theorem 1.1 tells us that distributed algorithms in the LOCAL model do not need randomness for efficiency. This holds for any *locally checkable problem*, i.e., any problem for which the solution can be checked efficiently deterministically[3].

At the heart of our derandomization result, and as the main novelty of this paper, we provide the first $\mathrm{poly}(\log n)$-round deterministic algorithm for network decomposition:

**Theorem 1.2** (**Network Decomposition Algorithm**)**.** *There is a deterministic distributed algorithm that in any* $n$-*node network* $G = (V, E)$, *in* $\mathrm{poly}(\log n)$ *rounds of the* LOCAL *model, partitions the vertices into* $O(\log n)$ *disjoint color classes* $V_1, \ldots, V_{O(\log n)}$, *such that in the subgraph* $G[V_i]$ *induced by the vertices of each color* $i$, *each connected component has diameter* $O(\log n)$.

We prove Theorem 1.2 in Section 2. We note that prior to our work, the best known deterministic network decomposition had a round complexity of $2^{O(\sqrt{\log n})}$, due to a celebrated work of Panconesi and Srinivasan [PS92]. This itself was an improvement on a $2^{O(\sqrt{\log n \log \log n})}$-round distributed algorithm, presented by Awerbuch et al. [AGLP89], in their pioneering work that defined network decomposition and showed its applications for distributed algorithms.

Our derandomization result, stated in Theorem 1.1, follows by putting our new network decomposition, as stated in Theorem 1.2, together with the derandomization framework developed by Ghaffari, Harris, and Kuhn [GHK18] and Ghaffari, Kuhn, and Maus [GKM17].

**Implications**: Through known connections, this derandomization leads to better *deterministic* and *randomized* distributed algorithms for a long list of well-studied problems. A sampling of the end-results includes (I) $\mathrm{poly}(\log n)$-round deterministic algorithms for maximal independent set, $\Delta + 1$ coloring, the Lovász Local Lemma, and defective coloring, as well as (II) a $\mathrm{poly}(\log \log n)$-time randomized $\Delta + 1$ coloring [CLP18], a $\mathrm{poly}(\log \log n)$-time randomized algorithm for Lovász

---

[1]To make our derandomization theorem stronger and more widely applicable, we use a relaxed version of local checkability: we call a problem *locally checkable* if its solution can be checked deterministically in $poly(\log n)$ rounds, such that if the solution is incorrect, at least one node knows. Thus, each constraint of the problem spans a neighborhood of at most $poly(\log n)$ rounds. Notice that this readily includes problems such as MIS, coloring, etc. For a precise definition of locally checkable problems (but bounded to constant radius), we refer to [NS95].

[2]This is similar to viewing a centralized algorithm with $\mathrm{poly}(n)$ time complexity or a parallel (PRAM model) algorithm with $\mathrm{poly}(\log n)$ time complexity as efficient.

[3]We remark that the vast majority of the problems studied in the LOCAL model throughout the literature are locally checkable. Moreover, such a restriction to locally checkable problems is necessary and the statement cannot hold for arbitrary problems, for trivial reasons: e.g., marking arbitrary $\Theta(\sqrt{n})$ nodes can be done in zero rounds by randomized algorithms but can be shown to require $\Omega(\sqrt{n})$ rounds for any deterministic algorithm [GHK18].

Local Lemma in constant degree graphs [GHK18], and an automatic complexity speed-up theorem from $o(\log n)$ to $\text{poly}(\log \log n)$ in constant-degree graphs, for *any* problem whose solution can be checked in $O(1)$ rounds [CP17]. We discuss the implications in Section 3.

## 1.3 An Overview of Our Network Decomposition Method

Our network decomposition algorithm is surprisingly simple. Next, we briefly recall the previous methods [PS92, AGLP89] and then give a quick outline of our construction:

**A recap on the previous constructions**: In the paper that introduced the concept of network decomposition, Awerbuch et al. [AGLP89] provide a deterministic algorithm that computes a network decomposition with clusters of diameter $2^{O(\sqrt{\log n \log \log n})}$, which are colored with $2^{O(\sqrt{\log n \log \log n})}$ colors, in $2^{O(\sqrt{\log n \log \log n})}$ rounds. In a nutshell, their algorithm is based on a hierarchical clustering. We start with each node being its own cluster. Over time, iteratively, we merge clusters together, in a manner that each final cluster has $2^{O(\sqrt{\log n \log \log n})}$ neighboring clusters, and thus the clusters can be easily colored with $2^{O(\sqrt{\log n \log \log n})}$ colors. Per iteration, we locally *group* clusters that have a "high" degree — more than $2^{O(\sqrt{\log n \log \log n})}$ neighboring clusters — around some selected clusters (chosen using a ruling set algorithm). Then, in each group, we merge all the clusters into one cluster. The center clusters are chosen using a *ruling set* procedure that ensures that the center clusters are somewhat far apart (concretely, at least 3 hops, in the *cluster graph* that connects any two clusters that have adjacent nodes), while any high-degree cluster has a center within a small distance (concretely, $O(\log n)$ hops, in the cluster graph). Due the separation and the high degrees, each merge is formed by grouping together at least $2^{\Theta(\sqrt{\log n \log \log n})}$ clusters. Hence, we finish in $O(\sqrt{\log n / \log \log n})$ iterations. Per iteration, each cluster has diameter at most $O(\log n)$ times the diameter of the previous clusters, and thus within $O(\sqrt{\log n / \log \log n})$ iteration, each cluster diameter grows to be at most $2^{O(\sqrt{\log n \log \log n})}$. The algorithm of Panconesi and Srinivasan [PS92] follows the same outline but replaces the ruling set procedure with a maximal independent set procedure (of a constant power of the cluster graph), computed by a clever and careful recursive idea. This replaces the $O(\log n)$ growth factor in the diameter per iteration with $O(1)$. Then, re-optimizing the parameters to take advantage of this change improves the bounds to give a network decomposition with clusters of diameter $2^{O(\sqrt{\log n})}$, which are colored with $2^{O(\sqrt{\log n})}$ colors, in $2^{O(\sqrt{\log n})}$ rounds.

In Appendix A, we provide a new method for constructing a network decomposition, which also achieves such a round complexity of $2^{O(\sqrt{\log n \log \log n})}$. We then also discuss how both of these twp approaches, which appear to be fundamentally different, cannot go below the complexity of $2^{O(\sqrt{\log n})}$, even on a particular simple and well-structured graph.

**Our construction, in a nutshell**: The main part of our result is to obtain a network decomposition with clusters of diameter $\text{poly}(\log n)$, which are colored with $\text{poly}(\log n)$ colors, in $\text{poly}(\log n)$ rounds. We provide a surprisingly simple algorithm for this. We can later transform this construction to improve the first two parameters to $O(\log n)$. Similar to the previously outlined methods, our algorithm also forms the clusters iteratively. However, unlike the hierarchical clusterings of [AGLP89, PS92]—where per iteration each new cluster is formed by merging a few of the nearby clusters of the previous iterations—during our construction, we *release* some clusters and allow each of their individual nodes to make an independent decision on which adjacent cluster to join; some of these nodes can also remain in their initial cluster, or *die*. Throughout the process, we ensure that at most a constant fraction of vertices die. Thus, via $O(\log n)$ repetition, each time by resurrecting the dead vertices and repeating the process on them, we can cluster all vertices. The decision of joining a neighboring cluster or dying is done in a manner that balances a few

desirable properties, as we outline next.

The clustering process has $b$ phases, where $b = O(\log n)$ denotes the number of bits in the identifiers. We start with a trivial clustering where each (remaining) vertex is one cluster, on its own. Each cluster is identified with the node identifier of its center vertex. We ensure that by the end of the $i^{th}$ phase, each two neighboring clusters have identifiers that agree in the $i$ least significant bits. In the $(i + 1)^{th}$ phase, clusters are categorized into *red* or *blue* clusters, based on the $(i + 1)^{th}$ least significant bit (while all clusters of each connected component agree on the $i$ least significant bits, by the construction's induction). Then, we *release* red clusters: their vertices might join one of the neighboring blue clusters, die, or remain in this red cluster if they have no neighboring blue cluster. On the other hand, each blue cluster retains all of its vertices and can also grow by accepting some of neighboring red vertices. This growth happens step by step, and hop by hop. Per step, each red node arbitrarily chooses a neighboring blue cluster to join, and each blue cluster checks the number of directly neighboring red vertices that want to join it. If they are at least a $1/(2b)$ fraction of the size of this blue cluster, they are accepted to join and they become blue. In this case, the cluster grows considerably in size, but also at most one hop in radius. But we cannot have more than $O(b \log n)$ such growth steps; beyond that the cluster would have more than $n$ vertices. On the other hand, if the fraction is less than a $1/(2b)$ fraction of the size of the blue cluster, all those red vertices die, and this blue cluster stops its growth for this iteration. This way, at the end of the steps of this phase, no edge remains between a blue and a red cluster, and at most a $1/(2b)$ fraction of all vertices die during the phase.

At the end of $b$ phases, one for each bit in the identifiers, at most a $b/(2b) = 1/2$ fraction of the vertices died, while each connected component of living vertices agrees on all the $b$ bits of the cluster identifier, i.e., is just one cluster. Since each cluster grows by at most one hop per each step of each phase, the cluster radii remain in $\text{poly}(\log n)$.

## 1.4 Other Related Work

We obtained the results in this paper after the second-named author learned about the statement of the main result of Kowalski and Krysta [KK19], which claims to provide a $\text{poly}(\log n)$ round algorithm for the splitting problem[4]. Due to results of Ghaffari et al. [GKM17], this statement would imply an alternative proof for P-LOCAL = P-RLOCAL. Hence, that would be effectively equivalent to the main result in our paper (modulo aspects such as the exact polynomial in the round complexity, message size, local computational complexity, simplicity, etc). However, two remarks are in order: (1) The proof in [KK19] has a fundamental flaw[5]. As of the time of preparing this version of our paper (i.e., May 12, 2020), we are not aware of any fix to that proof. (2) The methods in the two papers are completely different, in terms of both the general approach and the proof ingredients.

---

[4]Concretely, the second-named author received a request from the Program Committee of SODA 2020 to write a review on [KK19]. That review request was declined due to the conflict of interest.

[5]Here is a brief explanation: In page 11 of [KK19] (the version from 31 July 2019), the inequality $Pr[A_2|A_1] \leq 2^{-c'\alpha\delta}$ is incorrect. The provided argument says that this is derived by using the same arguments as done for $Pr[A_1]$ in Lemma 2. However, that argument cannot be applied in the second phase and further on. For the second phase, we have hyperedges left each with potentially only $\alpha\delta$ vertices that are left uncolored. Notice that this is much lower than the $\delta$ vertices that was assumed when proving Lemma 2. Hence, recalling that an edge is called biased if it has at most $\alpha\delta$ red edges or at most $\alpha\delta$ blue edges, the probability of an edge being biased in the new coloring (i.e., second phase) is effectively 1. If we change the definition of biased edges and allow the bias to go down by a constant factor per phase, this issue would be resolved superficially but then we can continue the argument for only $O(\log \log n)$ phases, as after that the hyperedges that initially had $\text{poly}(\log n)$ vertices might have no uncolored vertex left.

4

# 2 The Network Decomposition Algorithm

In this section, we present a network decomposition algorithm that proves Theorem 1.2. We first describe in Section 2.1 an $O(\log^7 n)$-round deterministic distributed algorithm in the LOCAL model that computes a weak-diameter network decomposition for $n$-node graphs, with cluster weak-diameter $O(\log^3 n)$ and $O(\log n)$ colors. This algorithm can also be adapted to work in $O(\log^8 n)$ rounds of the CONGEST model. Then, in Section 2.3, we explain how the former can be transformed to an $O(\log^8 n)$-time deterministic algorithm in the LOCAL model for strong-diameter network decomposition, with cluster strong-diameter $O(\log n)$ and $O(\log n)$ colors. The distinction between weak-diameter and strong diameter is clarified in Section 2.1.

As a side remark, we note that all these constructions assume that nodes have unique $O(\log n)$-bit identifiers. As we will explain later in Remark 2.10), in the LOCAL model, these constructions can be turned into poly($\log n$)-round algorithms for the more general setting with identifiers from $[1, S]$, as long as $\log^* S = O(\log n)$.

## 2.1 Weak-Diameter Network Decomposition

Recall that for Theorem 1.2, we wish to construct a decomposition of the underlying graph in $O(\log n)$ color classes such that for each color class, each of its connected components has $O(\log n)$ diameter. Our initial algorithm will, however, provide only a weaker property, as we describe next. We will work with *clusters* of vertices, defined simply as a subset of vertices, such that any two vertices of a cluster are "close" in $G$, although the subgraph induced by the vertices of the cluster may have large diameter and may be even disconnected. This motivates the notion of weak-diameter and the corresponding relaxation of network decomposition:

**Definition 2.1.** *Given a graph $G$ and its subgraph $H$, we say that the weak-diameter of $H$ is at most $d$ if $G$ contains a path of length at most $d$ between any pair of vertices in $H$.*

**Definition 2.2.** *Given a graph $G$, we define a weak-diameter network decomposition of $G$ with $c$ colors and weak-diameter $d$ to be a coloring of the vertices with $c$ colors such that for each color $i \in [1, c]$, the subgraph $G_i$ induced by the vertices of color $i$ is partitioned into non-adjacent disjoint clusters, each of weak-diameter at most $d$ in graph $G$.*

Next we state the main technical contribution of this paper, which is a deterministic distributed algorithm that constructs a weak-diameter decomposition in poly($\log(n)$) rounds in the LOCAL model. With the known connection that transforms it to a strong-diameter decomposition algorithm, as we will later describe in Section 2.3, this implies Theorem 1.2.

Before stating the result, we recall another useful notion of Steiner trees. A Steiner trees is a tree with nodes labelled as *terminal* and *nonterminal*; the aim is to connect terminal nodes possibly via some nonterminal nodes. Here we use this notion to control the weak-diameter of each cluster.

**Theorem 2.3.** *Consider an arbitrary $n$-node network graph $G$ where each node has a unique $b = O(\log n)$-bit identifier. There is a deterministic distributed algorithm that computes a network decomposition $G$ with $O(\log n)$ colors and weak-diameter $O(\log^3 n)$, in $O(\log^7 n)$ rounds of the LOCAL model.*

*Moreover, for each color and each cluster $\mathcal{C}$ of vertices with this color, we have a Steiner tree $T_{\mathcal{C}}$ with radius $O(\log^3 n)$ in $G$, for which the set of terminal nodes is equal to $\mathcal{C}$. Furthermore, each edge in $G$ is in $O(\log^2 n)$ of these Steiner trees.*

5

The last part of the statement ensures that our algorithm can also be implemented and used in the more restrictive CONGEST model, as we will later discuss in Remark 2.11.

In the following lemma, we describe the process for constructing the clusters of one color of the network decomposition (e.g., the first color), in a way that it clusters at least half of the vertices. This last weakening of the guarantee is similar to the randomized network decomposition algorithm of [LS93]. Since after each application of this lemma only half of the vertices remain, by $\log n$ repetitions, we get a decomposition of all vertices, with $\log n$ colors.

**Lemma 2.4.** *Consider an arbitrary n-node network graph $G = (V, E)$ where each node has a unique $b = O(\log n)$-bit identifier, as well as a set $S \subseteq V$ of living vertices. There is a deterministic distributed algorithm that, in $O(\log^6 n)$ rounds in the LOCAL model, finds a subset $S' \subseteq S$ of living vertices, where $|S'| \geq |S|/2$, such that the subgraph $G[S']$ induced by set $S'$ is partitioned into non-adjacent disjoint clusters, each of weak-diameter $O(\log^3 n)$ in graph $G$.*

*Moreover, for each such cluster $\mathcal{C}$, we have a Steiner tree $T_{\mathcal{C}}$ with radius $O(\log^3 n)$ in $G$ where all nodes of $\mathcal{C}$ are exactly the terminal nodes of $T_{\mathcal{C}}$. Furthermore, each edge in $G$ is in $O(\log n)$ of these Steiner trees.*

We obtain Theorem 2.3 by $c = \log n$ iterations of applying Lemma 2.4, starting from $S = V$. For each iteration $j \in [1, \log n]$, the set $S'$ are exactly nodes of color $j$ in the network decomposition, and we then continue to the next iteration by setting $S \leftarrow S \setminus S'$.

**Construction outline for one color of the decomposition**: We now describe the construction outline of Lemma 2.4. The construction has $b = O(\log n)$ phases, corresponding to the number of bits in the identifiers. Initially, we think of all nodes of $S$ as *living*. During this construction, some living nodes *die*. We use $S'_i$ to denote the set of living vertices at the beginning of phase $i \in [0, b-1]$. Slightly abusing the notation, we let $S'_b$ denote the set of living vertices at the end of phase $b - 1$ and define $S'$ to be the final set of living nodes, i.e., $S' := S'_b$.

Moreover, we *label* each living node $v$ with a $b$-bit string $\ell(v)$, and we use these labels to define the clusters. At the beginning of the first phase, $\ell(v)$ is simply the unique identifier of node $v$. This label can change over time. For each $b$-bit label $L \in \{0, 1\}^b$, we define the corresponding *cluster* $S'_i(L) \subseteq S'_i$ in phase $i$ to be the set of all living vertices $v \in S'_i$ such that $\ell(v) = L$. We will maintain one Steiner tree $T_L$ for each cluster $S'_i(L)$ where all nodes $S'_i(L)$ are the terminal nodes of $T_L$. Initially, each cluster consists of only one vertex and this is also the only (terminal) node of its respective Steiner tree.

**Construction invariants**: The construction is such that, for each phase $i \in [0, b-1]$, we maintain the following invariants:

(I) For each $i$-bit string $Y$, the set $S'_i(Y) \subseteq S'_i$ of all living nodes whose label ends in suffix $Y$ has no edge to other living nodes $S'_i \setminus S'_i(Y)$. In other words, the set $S'_i(Y)$ is a union of some connected components of the subgraph $G[S'_i]$ induced by living nodes $S'_i$.

(II) For each label $L$ and the corresponding cluster $S'_i(L)$, the related Steiner tree $T_L$ has radius at most $iR$, where $R = O(\log^2 n)$. We emphasize that in the subgraph induced by living vertices a cluster can be disconnected.

(III) We have $|S'_{i+1}| \geq |S'_i|(1 - 1/2b)$.

These invariants, together with Observation 2.9 about the overlaps of the Steiner trees, prove Lemma 2.4. In particular, from the first invariant we conclude that at the end of $b$ phases, different clusters are non-adjacent. From the second invariant we conclude that each cluster has a Steiner

tree with radius $bR = O(\log^3 n)$. Finally, from the third invariant we conclude that for the final set of living nodes $S' = S'_b$, we have $|S'| \geq (1 - 1/2b)^b |S| \geq |S|/2$.

**Outline of one phase of construction**: We now outline the construction of one phase and describe its goal (see also Figure 1). Let us think about some fixed phase $i$. We focus on one specific $i$-bit suffix $Y$ and the respective set $S'_i(Y)$. Let us categorize the nodes in $S'_i(Y)$ into two groups of *blue* and *red*, based on whether the $(i+1)^{th}$ least significant bit of their label is 0 or 1. Hence, all blue nodes have labels of the form $(* \ldots * 0Y)$ and all red nodes have labels of the form $(* \ldots * 1Y)$, where $*$ can be an arbitrary bit. During this phase, we make some small number of the red vertices die and we change the labels of some of the other red vertices to blue labels (and then the node is also colored blue). All blue nodes remain living and keep their label. The eventual goal is that, at the end of the phase, among the living nodes, there is no edge from a blue node to a red node. Hence, each connected component of the living nodes consists either entirely of blue nodes or entirely of red nodes. Therefore, the length of the common suffix in each connected component is incremented, which leads to invariant (I) for the next phase. The construction ensures that we kill at most $|S'_i(Y)|/2b$ red vertices of set $S'_i(Y)$, during this phase. We next describe this construction.

**Steps of one phase**: Each phase consists of $R = 10b \log n = O(\log^2 n)$ steps, each of which will be implemented in $O(\log^3 n)$ rounds. Hence, the overall round complexity of one phase is $O(\log^5 n)$ and over all the $O(\log n)$ phases, the round complexity of the whole construction of Lemma 2.4 is $O(\log^6 n)$ as advertised in its statement. Each step of the phase works as follows: each red node sends a request to an arbitrary neighboring blue cluster, if there is one, to join that blue cluster (by adopting the label). For each blue cluster $A$, we have two possibilities:

(1) If the number of adjacent red nodes that requested to join $A$ is less than or equal to $|A|/2b$, then $A$ does not accept any of them and all these requesting red nodes die (because of their request being denied by $A$). In that case, cluster $A$ *stops* for this whole phase and does not participate in any of the remaining steps of this phase.

(2) Otherwise — i.e., if the number of adjacent red nodes that requested to join $A$ is strictly greater than $|A|/2b$ — then $A$ accepts all these requests and each of these red nodes change their label to the blue label that is common among all nodes of $A$. In this case, we also grow the Steiner tree of cluster $A$ by one hop to include all these newly joined nodes.

We note that each step can be performed in $O(\log^3 n)$ rounds, because each blue cluster has a Steiner tree of depth $O(\log^3 n)$ and therefore can gather the number of vertices in the cluster, as well as the number of red vertices that would like to join this cluster. We also emphasize that in each step, each red node acts alone, independent of other nodes in the same red cluster. Hence, red clusters may shrink, disconnect, or even get dissolved over time. Once a red node adopts a blue label (or if a node had a blue label at the beginning), it will maintain that label throughout the phase. Therefore, blue clusters can only grow, and have more and more red nodes join them. We also emphasize that we can have blue clusters adjacent to each other, and red clusters adjacent to each other – the objective is to have no edge connecting a red cluster to a blue cluster.

Let us observe how the Steiner trees of the clusters evolve: For each blue cluster, the corresponding Steiner tree only grows. To have a similar property about the Steiner trees of red clusters, we do the following: Although for a red cluster, a terminal red node might become blue, we keep it in this tree as a nonterminal node. We note that although the Steiner tree of a red cluster is not used in the current phase, it may be used in the next phases.

**Analysis**: We next provide some simple observations about this construction in one phase, which allow us to argue that the construction maintains invariants (I) to (III), described above.
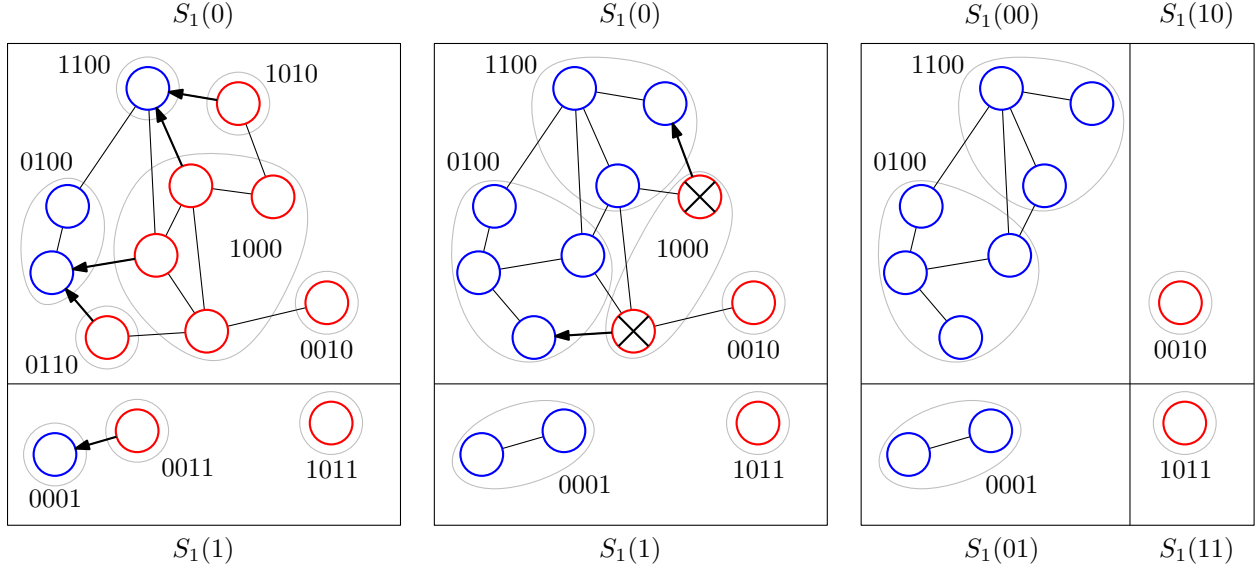
7

Figure 1: In this illustration, we consider the second phase of the algorithm, in a simple example graph. The three figures show the configuration in the beginning of three steps of this phase from left to right. Note that, at the beginning of this phase, the clusters are already separated according to their least significant bit (as a result of the first phase). When the second phase starts—i.e., in the left figure—the second least significant bit determines whether each cluster is blue or red. Adjacent red nodes are proposing to blue nodes (dark arrows) to join their clusters and blue clusters decide either to relabel them so that they join this cluster or to make them die (crossed red vertices). In the end, blue and red clusters are separated. Note that nothing will happen in the third phase, since the only two adjacent clusters share the same bit on the third least significant bit. Their boundary will be resolved only in the last phase.

**Observation 2.5.** *Any blue cluster stops after at most $4b \log n$ steps.*

*Proof.* In each step that a cluster $A$ does not stop, its size grows by a factor of at least $(1 + 1/2b)$, as it accepts at least $|A|/2b$ requests from neighboring red nodes. Hence, after $4b \log n$ steps of growth, the size would exceed $(1 + 1/2b)^{4b \log n} > n$, which is not possible. Therefore, cluster $A$ stops after at most $4b \log n$ steps. $\square$

**Observation 2.6.** *Once a blue cluster $A$ stops, it has no edge to a red node and it will never have one, during this phase. This implies invariant (I).*

*Proof.* By the observation above, cluster $A$ stops after at most $4b \log n$ steps. Consider the step in which cluster $A$ stops. In that step, each neighboring red node (if there is one) either requested to join $A$ or some other blue cluster. In the former case, that red node dies. In the latter case, the node adopts a blue label or dies. In either case, the node is not a living red node anymore (and it will never become one). From this point onward, this blue cluster $A$ never grows or shrinks. $\square$

**Observation 2.7.** *In each step, the radius of the Steiner tree of each blue cluster grows by at most 1, while the radius of the Steiner tree of each red cluster does not grow. This implies invariant (II).*

**Observation 2.8.** *The total number of red vertices in $S_i'(Y)$ that die during this phase is at most $|S_i'(Y)|/(2b)$. This implies invariant (III).*

*Proof.* From Observation 2.6, it follows that each blue cluster $A$ stops exactly once, and if it had $|A|$ vertices at that point, it makes at most $|A|/(2b)$ red vertices die. Hence, in total over the whole

8

phase, the number of red vertices that die is at most a $1/(2b)$ fraction of the number of nodes in blue clusters that stop, and thus at most $|S_i'(Y)|/(2b)$. $\qquad\square$

The above completes the description of our algorithm in the LOCAL model. As we will later remark about its applications in the CONGEST model, we finish the proof of Lemma 2.4 by adding the following observation about the overlaps of the constructed Steiner trees.

**Observation 2.9.** *Eeach edge is used in $O(\log n)$ Steiner trees.*

*Proof.* Each edge can be in the Steiner tree of a cluster only if that cluster at some point included one of the two endpoints of this edge. Throughout the construction, each node changes its label at most $b = O(\log n)$ times, i.e., at most once per label bit. Hence, each edge is used in $O(\log n)$ Steiner trees. $\qquad\square$

Below, just to help with the intuition, we discuss an idealized global view of the process in one phase. We then state some remarks about extensions of the result to the CONGEST model and the settings with larger identifiers.

**An intuitive global view of one (ideal) phase**: We next describe a different global view for an idealized version of the process in one phase. We hope that this view helps in understanding the process; concretely, the above process can be seen as a *localized* version of the idealized global view, where some decisions are performed locally (thus, the colors of nodes might differ in the two processes, but the growth of blue nodes and the number of red nodes that die, when the growth stops, behave similarly).

The process described above for one phase intends to make sure that there is no edge between red and blue living nodes, while the number of (red) nodes that die is kept small. For that, we grow the blue clusters locally (i.e., relabeling some red nodes to adopt blue labels, while keeping each blue label for the entire phase), each by $O(\log^2 n)$ hops, while making some red nodes die in the meantime. The process also guarantees that only a $1/2b$ fraction of nodes die. If we were to ignore the exact labels of the blue nodes and red nodes, and we would just remember whether a node is red or blue, the quantitative aspects of this process — namely the number of steps of growth and the number of red nodes that die — would resemble a simpler global ball carving argument: we would start from the initial "ball" of all blue nodes being together, and would grow this "ball" hop by hop, as long as in each step we grow by at least a $(1 + 1/2b)$ factor. In the first step that there is no such rapid growth — which will happen within $4b\log n$ steps — we would carve all the neighboring red nodes and call them dead. That would be at most a $1/2b$ fraction of the blue nodes (and hence all nodes). Once these boundary nodes are dead, there is no edge between living red and blue nodes.

**Remark about the length of identifiers**: For the construction in the LOCAL model, the requirement on the size of the identifiers of each node can be substantially weakened; this is important for applications when we use the algorithm in the shattering framework, e.g., [Gha16, BEPS16].

**Remark 2.10.** *In the construction provided above, we assumed that the nodes of the $n$-node graph have $O(\log n)$-bit unique identifiers. This construction can be extended to an $O(\log^* S \cdot \log^6 n)$-round algorithm in the LOCAL model for the setting where the identifiers are from $[1, S]$.*

*Proof.* Let $T(n) := O(\log^7 n)$ be the complexity of the algorithm in $n$-node graphs with $(3\log n)$-bit labels. We compute an $O(n^2)$ coloring of $G^{T(n)}$—the graph on the same set of vertices as $G$ but where we connect every two vertices $v$ and $u$ that have distance at most $T(n)$ in $G$—using the coloring algorithm of Linial [Lin92], in $T(n) \cdot \log^*(S) = O(\log^7 n) \cdot \log^*(S)$ rounds. We recall that

9

Linial's algorithm provides a $O(\Delta^2)$-coloring of any graph with maximum degree at most $\Delta$ where nodes have identifiers from $[1, S]$ in $O(\log^* S)$ rounds of the LOCAL model. Once we compute a coloring of $G^{T(n)}$ with $O(n^2)$ colors, we can then adopt these colors as "unique" identifiers with no more than $(3 \log n)$-bits. Since each node sees unique identifiers in its $(T(n))$-hop neighborhood, the LOCAL algorithm works as if nodes had unique identifiers. $\qquad\square$

## 2.2 Construction in the CONGEST Model and Extension to Graph Powers

Although we formulated the algorithm in the LOCAL model of computation, it can be easily observed that it also runs in the more restrictive CONGEST model.

**Remark 2.11.** *The whole network decomposition construction described in Lemma 2.4 can be performed in $O(\log^8 n)$ rounds of the CONGEST model.*

*Proof.* Recall from above that in the construction of clusters of one color, each edge is used in $O(\log n)$ Steiner trees. Moreover, whenever we add an edge to a particular Steiner tree, we can think of it as being oriented from a newly added node towards a node that was already in the tree. This gives a natural orientation of its edges that points to its root, which is the vertex whose original identifier is equal to the label of the cluster, and that was initially the only member of this cluster.

The construction only uses convergecast and broadcast along these rooted trees (to decide whether the cluster should continue growing or it should stop). Hence, by using every $O(\log n)$ rounds of CONGEST model as one *big-round*, we can perform the construction of one color in $O(\log^6 n)$ big-rounds, i.e., $O(\log^7 n)$ rounds of the CONGEST model. Over all the $O(\log n)$ colors, this leads to a round complexity of $O(\log^8 n)$ rounds of the CONGEST model. $\qquad\square$

In the CONGEST model it is particularly helpful that Lemma 2.4 gives us the underlying Steiner tree for each cluster, with the property that each appears in only $O(\log n)$ trees per color. These Steiner trees can later be used for simultaneous broadcast or convergecast in the clusters.

**Extending the CONGEST-model construction to graph powers**: When solving a distributed problem for an underlying graph $G$, it is often helpful to simulate its power $G^k$ and run certain algorithm on this simulated graph (for example, this will be the case in Theorem 2.14 as well as all the applications in Section 3.3.5). Obtaining a network decomposition for $G^k$ is straightforward in the LOCAL model, where each node can start by collecting its $k$-hop neighbourhood and then simulate each step of an algorithm for $G^k$ with additional slowdown proportional to $k$. However, this cannot be done easily in the CONGEST model[6]. That being said, our algorithm can be adapted to provide a weak-diameter network decomposition for $G^k$ in the CONGEST model without the need of an explicit construction of $G^k$.

A weak-diameter decomposition of $G^k$ with $c$ color classes of weak-diameter $d$ can be also interpreted as a weak-diameter decomposition of $G$ with $c$ color classes of weak-diameter $k \cdot d$, where any two clusters of the same color class are at least $k + 1$ hops apart.

**Theorem 2.12.** *There is an algorithm in the CONGEST model that, given a value $k$ that is known to all nodes, in $O(k \log^8 n \cdot \min(k, \log^2 n))$ communication rounds outputs a weak-diameter network decomposition of $G^k$ with $O(\log n)$ color classes, each one with $O(k \cdot \log^3 n)$ weak-diameter in $G$.*

---

[6]Even the task of collecting 2-hop neighbourhood of a given node $u$ cannot be generally solved in $\mathrm{poly}(\log n)$ rounds, since the number of vertices in the 2-hop neighbourhood of $u$ can be much larger than the number of connections of $u$ to its neighbours that can be used to collect information.

*Moreover, for each color and each cluster $\mathcal{C}$ of vertices with this color, we have a Steiner tree $T_{\mathcal{C}}$ with radius $O(k \cdot \log^3 n)$ in $G$, for which the set of terminal nodes is equal to $\mathcal{C}$. Furthermore, each edge in $G$ is in $O(\log^2 n \cdot \min(k, \log^2 n))$ of these Steiner trees.*

The proof idea is to run the algorithm from Lemma 2.4, with one change: vertices that will propose in one step will not be just red vertices bordering with a blue vertex, but all red vertices in a $k$-hop neighbourhood of some blue vertex. This idea by itself readily gives us a poly$(\log n)$ round algorithm for $k = \text{poly}(\log n)$. As we will show, with some more work, we can also get an algorithm with the same round complexity as the algorithm from Remark 2.11 whenever $k$ is constant. Below, we provide a concrete proof sketch. The full proof is deferred to Appendix B.

*Proof sketch of Theorem 2.12.* Consider the algorithm from Theorem 2.3 with the following change. We generalize each *step* — red nodes with a blue neighbour are proposing to an arbitrary blue neighbour— so that all the nodes with at least one blue neighbour in the their $k$-hop neighbourhood are proposing. We can mark all such red nodes in $k$ rounds by running a breadth first search (BFS) from all blue nodes simultaneously. This can be implemented in such a way that each edge is used in at most one round, hence it can be done in $k$ rounds of the CONGEST model. Note that running BFS also naturally outputs an oriented forest with blue nodes as the roots. This forest will also contain dead nodes. Now it is simple to implement proposals of red nodes to blue clusters: each red node in a forest will propose to its root and, if accepted by the respective blue cluster, the whole path to the root (that potentially also contain dead nodes) is added to the corresponding Steiner tree.

It is easy to see that this adaptation of our algorithm returns a network decomposition of $G^k$.

Before we analyze the running time, let us add an additional optimization to ensure that each edge will be added to at most $k$ Steiner trees during one phase (it is added to at most one Steiner tree during one phase in the original algorithm): During a phase and in between steps, a dead node $v$ that was used in a BFS tree rooted at $r$ will change to a different BFS tree only if the blue root $r'$ of the new BFS tree is closer to the node $v$ than its current root $r$. This property will hold in our implementation of the algorithm because of the following two properties: (A) we break symmetry during BFS in the same manner in all the steps (e.g., toward the one with smaller ID) and (B) even blue clusters that stopped growing are running BFS in every step, to make sure that a dead node changes a BFS tree only when a closer blue cluster appeared. With this optimization, each edge is used by at most $O(\min(k, \log^2(n)))$ different Steiner trees in a phase and, hence, by at most $O(\log^2(n) \cdot \min(k, \log^2(n)))$ Steiner trees in total.

We are now ready to bound the running time. The algorithm constructs $O(\log n)$ color classes and each one is constructed in $O(\log n)$ phases with each phase containing $O(\log^2 n)$ steps. For each step, we need to run breadth first search for $O(k)$ steps and broadcast information to root via Steiner trees of depth $O(k \log^3 n)$, which dominates the $O(k)$ steps of the breadth first search. Moreover, each edge is used by $O(\log n \cdot \min(k + \log^2 n))$ of the Steiner trees. This implies a round complexity $O(k \log^8 n \cdot \min(k + \log^2 n))$. □

## 2.3 Strong-Diameter Network Decomposition

We now explain that by a known method, first presented by Awerbuch et al. [ABCP96], in the LOCAL model, we can transform the algorithm of Theorem 2.3 for weak-diameter network decomposition to an algorithm for strong-diameter network decomposition, which thus proves Theorem 1.2. Since this is a known connection, we provide only a sketch of the proof.

**Definition 2.13.** *Given a graph $G = (V, E)$, we define a network decomposition of $G$ with $c$ colors and strong-diameter $d$ to be a partitioning of the vertices into $c$ disjoint color classes $V_1, \ldots, V_c,$*

such that in the subgraph $G[V_i]$ induced by the vertices of each color $i$, each connected component has diameter at most $d$. Each of these connected component of the subgraph $G[V_i]$ is called a *cluster*.

The following theorem statement is a rephrased and formalized version of Theorem 1.2:

**Theorem 2.14.** *Consider an arbitrary $n$-node network graph where each node has a unique $b = O(\log n)$-bit identifier. There is a deterministic distributed algorithm that computes a network decomposition of this graph with $O(\log n)$ colors and strong-diameter $O(\log n)$, in $O(\log^8 n)$ rounds in the* LOCAL *model.*

*Proof Sketch.* We first recall the standard sequential algorithm for building a network decomposition with $\log n$ colors and strong-diameter $O(\log n)$ per cluster, and then we explain how the weak-diameter network decomposition algorithm of Theorem 2.3 helps us build such a strong-diameter decomposition in a distributed manner.

The standard sequential algorithm algorithm for building a network decomposition with $\log n$ colors and strong-diameter $O(\log n)$ per cluster works as follows. We describe the process for determining the nodes of the first color. The other colors are obtained similarly, by applying the same construction repeatedly for $\log n$ times, each to the graph induced by the remaining nodes. For the first color, starting with an arbitrary node $v$, we do a *sequential ball carving*. That is, we grow a *ball* around this vertex, hop by hop. A ball of radius $r$ is simply all nodes that are within distance $r$ of node $v$. We increment the radius $r$ of this ball gradually, one by one, as long as the number of the nodes outside the ball that are adjacent to the ball is at least equal to the number of nodes in the ball. Once this growth condition is not satisfied, which will happen before $r \leq \log n$ as otherwise the ball has more than $n$ nodes, we consider all nodes in the ball as one cluster (in the first color) of the decomposition, and we consider all nodes outside but adjacent to it as *dead*. All dead nodes are removed from the construction of the first color of the network decomposition. Then, repeatedly, if any node remains that is not dead but also not clustered, we continue a similar sequential ball carving process starting from that node, among nodes that are not dead. This gives the clusters of the first color of the decomposition. Then, we bring all the dead nodes back to life and repeat this process among them, getting the clusters of the second color, and so on. Since each time we cluster at least $1/2$ of the vertices, we finish after at most $\log n$ repetitions, i.e., at most $\log n$ colors.

We now explain how Theorem 2.3 allows us to get an efficient distributed simulation of this sequential construction, thus proving Theorem 2.14. Let $G' := G^{10 \log n}$, i.e., the graph on the same set of vertices as $G$ but where we connect every two vertices $v$ and $u$ that have distance at most $10 \log n$ in $G$. We apply the algorithm of Theorem 2.3 to obtain a weak-diameter network decomposition of $G'$, in $O(\log^8 n)$ rounds of communication on $G$. The resulting network decomposition is a coloring of vertices with $q = O(\log n)$ colors where the clusters in each color have weak-diameter $O(\log^3 n)$ in $G'$, and thus weak-diameter $O(\log^4 n)$ in $G$. We next use this *helper* network decomposition to build our *output* strong-diameter network decomposition with $O(\log n)$ colors and $O(\log n)$ diameter.

We describe the process for determining the nodes of the first color in the output network decomposition. The other colors are obtained similarly, by applying the same construction repeatedly for $O(\log n)$ times, each to the graph induced by the remaining nodes.

To determine the nodes of the first color of the output decomposition, we process the colors of the helper network decomposition one by one, in $q$ stages. Let us fix one stage (and thus one color of the helper network decomposition, and its clusters). For each cluster, we elect a leader for it and we gather the topology of the subgraph of all remaining nodes within $\log n$ hops of the nodes of this cluster. Notice that since the cluster has weak-diameter $O(\log^4 n)$, this can be done

12

in $O(\log^4 n)$ rounds. Moreover, the topologies gathered by different clusters are disjoint. This is because different clusters of this color of the helper decomposition have distance at least $10 \log n$, since otherwise $G'$ would contain an edge connecting the two clusters.

Each cluster $\mathcal{C}$ will perform a sequential ball carving process, on the topology that it has gathered, as follows: We start from an arbitrary node $v$ of color $i \in [1, q]$ in cluster $\mathcal{C}$, and grow a *ball* around it, hop by hop, in the subgraph induced by the remaining nodes. A ball is simply all remaining nodes that are within a certain distance of node $v$, in the subgraph induced by the remaining nodes. We grow the radius of this ball gradually, as long as the number of the nodes outside the ball that are adjacent to the ball is at least equal to the number of nodes in the ball. Once this growth condition is not satisfied, we consider all nodes in the ball as one cluster of the output decomposition, and we consider all nodes outside but adjacent to it as *dead*. All dead nodes are removed from the construction of this color of the output network decomposition. Then, if any node $v'$ of cluster $\mathcal{C}$ remains unclustered (for the output decomposition), we start a similar ball growing process from $v'$, but only on the graph induced by the remaining nodes. We continue similarly until all nodes of cluster $\mathcal{C}$ are clustered for the output decomposition.

In each step of growing a ball, the number of nodes grows by a 2 factor. Hence, any ball can grow by at most $\log n$ hops. This implies that the ball growing processes from cluster $\mathcal{C}$ will never reach the ball growing processes from any other cluster $\mathcal{C}'$ of color $i$ of the helper decomposition. Furthermore, each time that we stop a ball's growth, the number of nodes on the boundary of it that die is less than the number of nodes inside the ball (which get clustered for the output network decomposition). Hence, after going through all the $q$ stages, at least $1/2$ of living nodes get clustered, and at most $1/2$ of living nodes die.

Then, we bring all dead nodes back to life and proceed to build the next color of the output network decomposition, only on the subgraph induced by these remaining nodes. As per repetition the number of remaining nodes reduces by a 2 factor, we finish in $\log n$ repetitions. $\qquad\square$

# 3 Implications and Applications

As mentioned before, despite its simplicity, our efficient deterministic network decomposition has far-reaching implications, leading to a general efficient distributed derandomization theorem and better *deterministic* and *randomized* distributed algorithms for a range of problems, as well as some improvements in massively parallel computation (aka, the MapReduce algorithms). We next overview these implications. We start in Section 3.1 with the well-studied problems of maximal independent set and coloring, which were among the most well-known open problems in distributed graph algorithms and get settled immediately by our network decomposition. This also serves as a warm up for the standard method of using network decomposition. Then, in Section 3.2, we present our general derandomization result for the LOCAL model, thus proving Theorem 1.1. Finally, in Section 3.3, we overview a list of other well-studied problems for which we get substantial (deterministic or randomized) improvements.

## 3.1 Maximal Independent Set and Coloring

### 3.1.1 MIS

The Maximal Independent Set (MIS) problem is one of the central problems in the study of distributed graph algorithms. As mentioned before, there have been well-known $O(\log n)$-round randomized algorithm for this problem since the 1980s [Lub86, ABI86] but obtaining a deterministic algorithm for it had remained open.

**Deterministic MIS**: We next explain how the efficient network decomposition of Theorem 2.3 directly gives a poly($\log n$)-round deterministic MIS algorithm. This already answers Linial's long-standing open question and settles Open Problem 11.2 in the book of Barenboim and Elkin [BE13]. Weaker forms of this problem appear as Open Problems 11.5 and 11.8 in the same book [BE13] and they are now resolved. The method is fairly standard and thus we provide a proof sketch. It also allows us to recall the usual method of using network decomposition to solve problems such as maximal independent set and coloring [AGLP89].

**Corollary 3.1.** *There is a deterministic distributed algorithm, in the* LOCAL *model, that computes a maximal independent set in* poly($\log n$) *rounds.*

*Proof.* First, we compute a network decomposition with $O(\log n)$ colors and clusters of diameter $O(\log^3 n)$, in $O(\log^7 n)$ rounds, using Theorem 2.3. Then, we process the clusters color by color. In each color $i$, the center node of each cluster aggregates at the center the topology of the cluster as well as the information of which nodes adjacent to the cluster have already been added to the maximal independent set, when processing the previous colors 1 to $i-1$. Since the cluster diameter is $O(\log^3 n)$, this information can be gathered in $O(\log^3 n)$ rounds. Then, the center simulates a greedy process of adding the vertices of this cluster to the MIS, one by one, for any node that does not already have a neighbor in the MIS. Since any two cluster of the same color are non-adjacent, the computations of different clusters can happen simultaneously. Processing each color takes $O(\log^3 n)$ rounds, which means that we finish processing all the $O(\log n)$ colors in $O(\log^4 n)$ rounds. Together with the $O(\log^7 n)$ rounds used for computing the network decomposition, this is a deterministic maximal independent set algorithm that runs in $O(\log^7 n)$ rounds. □

We note that, due to a very recent breakthrough of Balliu et al. [BBH+19], any deterministic algorithm for MIS needs a round complexity of $\Omega(\log n / \log \log n)$.

**Randomized MIS**: Plugging the above deterministic MIS algorithm into the shattering framework of the algorithm of [Gha16] improves also the randomized complexity of MIS:

**Corollary 3.2.** *There is a randomized distributed algorithm, in the* LOCAL *model, that computes a maximal independent set in* $O(\log \Delta) + $poly($\log \log n$) *rounds, with probability at least* $1 - 1/$poly($n$).

We note that due to a celebrated lower bound of Kuhn, Moscibroda and Wattenhofer [KMW16], any (randomized) algorithm for MIS needs a round complexity of $\Omega(\frac{\log \Delta}{\log \log \Delta})$, which means the $\Delta$ dependency in the above algorithm is nearly optimal. Moreover, regarding the dependency on $n$, due to another result of Balliu et al. [BBH+19], any randomized algorithm for MIS needs a round complexity of $\Omega(\frac{\log \log n}{\log \log \log n})$, on some graphs with $\Delta = \Omega(\frac{\log \log n}{\log \log \log n})$. Thus, one cannot hope for an algorithm with round complexity $O(\log \Delta) + o(\frac{\log \log n}{\log \log \log n})$, or even $o(\Delta) + o(\frac{\log \log n}{\log \log \log n})$.

**MIS with small messages**: The algorithm described in the proof of Theorem 3.1 works in the LOCAL model, where message sizes are unbounded. We can also obtain an algorithm for the CONGEST model, where message sizes are bounded to $O(\log n)$:

**Corollary 3.3.** *There is a deterministic distributed algorithm, in the* CONGEST *model, that computes a maximal independent set in* poly($\log n$) *rounds.*

*Proof Sketch.* The method outline is similar to the LOCAL model algorithm, with two exceptions: (1) we use the CONGEST-model variant of our network decomposition, which runs in $O(\log^8 n)$ rounds, (2) when processing each cluster, we use a CONGEST-model MIS algorithm of Censor-Hillel, Parter, and Shwartzman [CHPS17], instead of the naive topology gathering step. Concretely,

Censor-Hillel et al. give an $O(D \log^2 n)$-round MIS algorithm in the CONGEST model, $D$ denotes the graph diameter. When processing the colors of network decomposition, for each cluster of the color, we can run the algorithm of Censor-Hillel et al. on the cluster (ignoring nodes that already have a neighbor in the MIS). Recall from Lemma 2.4 that per color, each edge of the graph is used by the Steiner trees of $O(\log n)$ clusters. Hence, we can run the algorithm of Censor-Hillel et al. for all the clusters of the same color, in parallel, in $O(\log^3 n \cdot \log^2 n \cdot \log n) = O(\log^6 n)$ rounds. Over all the $O(\log n)$ colors, this MIS computation runs in $O(\log^7 n)$ rounds of the CONGEST model, besides the initial $O(\log^8 n)$ rounds spent for computing a network decomposition. $\qquad\square$

### 3.1.2 Coloring

**Deterministic coloring**: One can apply the standard method for using network decomposition, as done above when proving Theorem 3.1, to also obtain an $O(\log^7 n)$ round algorithm for $\Delta + 1$ vertex coloring, where $\Delta$ denotes the maximum degree, or its generalization to list-coloring. This efficient coloring resolves Open Problem 11.3 in the book of Barenboim and Elkin [BE13] and gives an alternative, and more systematic, solution for Open Problem 11.4, which asked for an efficient deterministic $(2\Delta - 1)$-edge coloring (that problem was settled first in [FGK17]).

**Corollary 3.4.** *There is a deterministic distributed algorithm, in the* LOCAL *model, that computes a* $(\Delta+1)$ *vertex coloring, where* $\Delta$ *denotes the maximum degree in the graph, in* $\mathrm{poly}(\log n)$ *rounds. The algorithm can also be generalized to list-coloring where each vertex* $v$ *should choose its color from a list* $L_v$ *of colors, where* $|L_v| \geq deg(v) + 1$.

**Randomized coloring**: Moreover, plugging this deterministic list-coloring algorithm of Theorem 3.4 into the randomized coloring algorithm of Chang, Li, and Pettie [CLP18] improves the randomized complexity of $\Delta + 1$ coloring from $2^{O(\sqrt{\log \log n})}$ to $\mathrm{poly}(\log \log n)$:

**Corollary 3.5.** *There is a randomized distributed algorithm, in the* LOCAL *model, that computes a* $(\Delta + 1)$ *vertex coloring, where* $\Delta$ *denotes the maximum degree in the graph, in* $\mathrm{poly}(\log \log n)$ *rounds, with probability at least* $1 - 1/\mathrm{poly}(n)$.

*Proof Sketch.* Following the shattering framework [BEPS16], the randomized phase of the algorithm of [CLP18] works in $O(\log^* \Delta)$ rounds, and colors almost all nodes, except for some small components of nodes that remain uncolored. The guarantee is that, with probability at least $1 - 1/\mathrm{poly}(n)$, each remaining component has $\mathrm{poly}(\log n)$ vertices. After that, for the deterministic phase, we can invoke the deterministic list-coloring algorithm of Theorem 3.4 on each of these components separately, all in parallel. Since each component has $\mathrm{poly}(\log n)$ vertices, this would run in $\mathrm{poly}(\log(\mathrm{poly}(\log n))) = \mathrm{poly}(\log \log n)$ rounds, and would complete the partial coloring to a coloring for all vertices. $\qquad\square$

As another coloring result, by using Theorem 3.4 along with the method of [BE11], one can obtain an arboricity-dependent coloring:

**Corollary 3.6.** *There is a deterministic distributed algorithm that computes a* $(2 + o(1))a$*-coloring of any graph with arboricity at most* $a$*, in* $\mathrm{poly}(\log n)$ *rounds of the* LOCAL *model.*

**Massively Parallel Computation (MPC) of coloring**: we also get a nearly-exponential improvement for massively parallel (aka, MapReduce) algorithms [KSV10] for $\Delta + 1$ coloring. It is beyond the scope of this paper to explain the exact setting and review the related literature. For those, and particularly for the coloring problem, we refer the readers to [KSV10, CFG+19, GKU19].

We just briefly state that in the MPC model (with strongly sublinear memory per machine), the $n$-node graph is partitioned among a number of machines, each with memory $n^\alpha$ for a constant $\alpha < 1$, and per round each machine can send $n^\alpha$ bits to the other machines.

We obtain our improvement by plugging in the LOCAL-model deterministic list-coloring algorithm of Theorem 3.4 into the algorithm of [CFG$^+$19]. This gives a randomized MPC $\Delta+1$ coloring algorithm, with strongly sublinear memory per machine, with round complexity of $O(\log \log \log n)$, which improves on the previous bound of $O(\sqrt{\log \log n})$.

**Corollary 3.7.** *There is a randomized MPC algorithm, in the regime where each machine has memory $n^\alpha$ for any constant $\alpha < 1$, that computes a $\Delta + 1$ coloring of any $n$-node graph with maximum degree at most $\Delta$ in $O(\log \log \log n)$ rounds, with high probability.*

We also note that due to a conditional hardness result of [GKU19], conditioned on a standard hardness assumption of $\Omega(\log n)$-complexity for connectivity, improving this $O(\log \log \log n)$-round randomized MPC coloring algorithm would imply a deterministic $\log^{o(1)} n$-round deterministic distributed algorithm for $\Delta + 1$ coloring, in the LOCAL model, which would be a major improvement on the state of the art (Theorem 3.4).

## 3.2 Derandomization via Network Decomposition

We now explain how our network decomposition, when put together with the approach of [GHK18, GKM17], leads to an efficient derandomization method for the LOCAL model. We note that this result can be viewed as answering Open Problem 11.1 in the book of Barenboim and Elkin [BE13], which asked for developing "a general derandomization technique for the distributed message passing model" and was followed by several locally checkable problems that admit poly$(\log n)$-round randomized algorithms but no known poly$(\log n)$-round deterministic algorithm.

**Theorem 1.1 (LOCAL Derandomization Theorem)** *We have*

$$\text{P-LOCAL} = \text{P-RLOCAL}.$$

*Here,* P-LOCAL *denotes the family of locally checkable problems that can be solved by deterministic algorithms in* poly$(\log n)$ *rounds of the* LOCAL *model in $n$-node graphs and* P-RLOCAL *denotes the family of* locally checkable problems *that can be solved by randomized algorithms in* poly$(\log n)$ *rounds of the* LOCAL *model, with success probability $1 - 1/n$.*

*Proof Sketch.* A formal and precise description of this procedure can be found in [GHK18]. To keep this article self-contained and accessible to a broad audience, we provide a less formal sketch here, and without going through the language of the SLOCAL model of [GKM17].

Consider any locally checkable problem $\mathcal{P}$ that can be checked in $t(n)$ rounds by a deterministic LOCAL-model algorithm, and a randomized LOCAL-model algorithm $\mathcal{A}$ for $\mathcal{P}$ that runs in exactly $r(n)$ rounds and produces correct outputs with probability at least $1 - 1/\operatorname{poly}(n)$. Thus, composing these, we have an algorithm $\mathcal{B}$ that runs in $R = r(n) + t(n)$ rounds and computes the outputs for $\mathcal{P}$, as well as a correctness indicator flag $f_v$ for each node $v$ such that if a constraint of $\mathcal{P}$ involving node $v$ is not satisfied, then $f_v = 1$. In other words, if for all nodes $v \in V$ the indicator flags $f_v = 0$, the output is a valid solution for the problem. Moreover, the expected number of flags that equal to 1 is at most $1/\operatorname{poly}(n)$. We derandomize this algorithm $\mathcal{B}$ by working through the network decomposition, and fixing the randomness of different nodes, via a method of conditional expectation for the function $\sum_v f_v$.

We first take a network decomposition of $G^{2R+1}$ where each two nodes are connected if their distance is at most $2R + 1$. This can be computed deterministically in $R \operatorname{poly}(\log n)$ rounds of the

LOCAL model, using Theorem 2.3. We get a decomposition into clusters of radius $O(R \log^3 n)$, colored with $O(\log n)$ colors, such that any two clusters of the same color are more than $2R + 1$ hops apart.

Then, similar to the standard method explained in the proof of Theorem 3.1, we work through the colors of the network decomposition, one by one. Per color $i$, each cluster gathers the topology from $2R$-hop neighborhood of the cluster in the cluster center (this topology also includes the information of how randomness has been fixed, when processing previous colors), in $O(R \log^3 n)$ rounds. Then, each cluster center fixes the randomness of its vertices one by one, in a sequential manner, ensuring that the expectation of $\sum_v f_v$ conditioned on the fixed randomness does not increase. Notice that since $\mathcal{B}$ is an $R$ round algorithm, the randomness of each node $u$ influences only $f_v$ for nodes $v$ that are within distance $R$ of node $u$. Hence, the cluster center can compute the change in the expected value of $\sum_v f_v$ when fixing the randomness of each node $u$ in its cluster, and can fix the randomness in a way that does not increase the conditional expectation. Moreover, clusters of the same color can work in parallel as they are more than $2R + 1$ hops apart and hence they do not influence the same indicator flag $f_v$ for any node $v$. Once each cluster center fixes the randomness of the node's of its cluster, it reports these values back to the nodes, in $O(R \log^3 n)$ rounds. Then, we proceed to the next color and repeat a similar procedure. Once we finish processing all the $O(\log n)$ colors, all the randomness is fixed, and still the expected value of $\sum_v f_v$ is at most $1/\text{poly}(n) \ll 1$. Since $\sum_v f_v$ has to be a non-negative integer value, we must have $\sum_v f_v = 0$, which means all $f_v = 0$ and thus all the constraints are satisfied. Overall, we now have a deterministic algorithm that runs in $R \cdot \text{poly}(\log n)$ rounds. Hence, any locally checkable problem whose solution can be checked deterministically in $t(n) = \text{poly}(\log n)$ rounds and admits a randomized algorithm that runs in $r(n) = \text{poly}(\log n)$ rounds also has a deterministic algorithm that runs in $(r(n) + t(n)) \cdot \text{poly}(\log n) = \text{poly}(\log n)$ rounds. $\square$

## 3.3 Other Implications (Deterministic & Randomized)

Here, we mention some of the other implications. This list is not exhaustive; these are just some of the prominent instances that came to our mind. A more thorough job is needed to re-examine all the related literature and list all the consequences. Moreover, in the interest of brevity and due to the large number of the implications, here we just provide a brief and sometimes informal explanation of each problem; the precise setup can be found in the references that we mention.

### 3.3.1 Lovász Local Lemma and the Sublogarithmic Complexity Lanscape

The Lovász Local Lemma has turned out to have a fundamental role in several distributed problems, and perhaps most remarkably, in the complexity of the locally checkable problems that have sublogarithmic complexity. We next review the LLL problem and outline the new result.

**Lovász Local Lemma**: Consider a probabilistic setting of events defined on a set of random variables. There is one node for each *bad* event, and $p$ denotes the maximum probability among these bad events. Moreover, each two bad events that share a variable are connected via an edge, and we use $d$ to denote the maximum degree of this graph. The Lovász Local Lemma proves that if $epd \le 1$, then there is an assignment to the variables that avoids all the bad events. In the distributed version of this problem, the question is to efficiently compute such as assignment that avoids all the bad events, where the LOCAL-model graph is the same as the dependency graph among the events. See [CPS17, CP17, FG17, GHK18].

**Improved deterministic LLL**: By running the $O(\log^2 n)$-round randomized distributed LLL algorithm of Moser and Tardos [MT10] through the derandomization method of Theorem 1.1, we

get a poly($\log n$) round deterministic distributed algorithm for Lovász Local Lemma:

**Corollary 3.8.** *There is a deterministic distributed algorithm that solves the Lovász Local Lemma problem in* poly($\log n$) *rounds, so long as the maximum probability among the bad events p and the maximum dependency degree among them d satisfy $epd \leq 1 - \delta$, for any constant $\delta > 0$ or even a slightly sub-constant $\delta > 1/$ poly($\log n$).*

**Improved randomized LLL**: By plugging this deterministic Lovász Local Lemma algorithm into the frameworks of [FG17, GKM17], we get a randomized LLL algorithm with complexity poly($\log \log n$) in constant-degree graphs.

**Corollary 3.9.** *There is a randomized distributed algorithm that solves the Lovász Local Lemma problem in $O(d^2)+$ poly($\log \log n$) *rounds, so long as the maximum probability among the bad events p and the maximum dependency degree among them d satisfy $Cpd^8 \leq 1$, for some constant $C > 1$.*

This poly($\log \log n$) round complexity for constant-degree graphs almost settles a conjecture of Chang and Pettie [CP17]; their conjecture postulates the existence of an $O(\log \log n)$ time algorithm.

**Complexity of LCLs in the sublogarithmic landscape**: Due to a beautiful result of Chang and Pettie [CP17], this improved LLL has a remarkable complexity-theoretic consequence:

**Corollary 3.10.** *Any locally-checkable problem that admits an $o(\log n)$ round randomized distributed algorithm in constant-degree graphs also admits a* poly($\log \log n$) *round randomized algorithm.*

That is, for any problem whose solution can be checked deterministically in $O(1)$ rounds, in bounded degree graphs, the randomized complexity is either $\Omega(\log n)$ and above, or poly($\log \log n$) and below. As soon as we can prove some LCL problem to admit an $o(\log n)$-round algorithm, we immediately get a poly($\log \log n$) round algorithm.

### 3.3.2 Packing/Covering Integer Linear Programs

Covering and packing integer Linear Programs are LPs in the standard form where all the coefficients are non-negative; the former is a minimization problem and the latter is a maximization problem. A wide range of optimization problem can be formulated in this manner.

A general result of Ghaffari, Kuhn, and Maus [GKM17, Section 7] shows that for any covering or packing integer linear program, there is a poly($\log n/\varepsilon$) round randomized algorithm in the LOCAL model for computing a $1 + \varepsilon$ (integral) approximation. The concrete distributed formulation of these LPs is that we have a bipartite graph where each node on the left shows one of the variables and each node on the right shows one of the constraints, and a constrain node is connected to the variable nodes that it includes. Cf. [GKM17] for details. We note that one can imagine a number of other natural formulations of the optimization problem as a graph, but in the LOCAL model, these usually can simulate each other with a constant round complexity overhead.

By plugging our network decomposition into the framework of [GKM17], we can derandomize their result and get a deterministic variant:

**Corollary 3.11.** *For any covering or packing integer linear program, there deterministic algorithm in the* LOCAL *model that computes a $1 + \varepsilon$ approximation in* poly($\log n/\varepsilon$) *rounds.*

We note that the conference version of [GKM17] describes the method explicitly only for the maximum independent set problem, but the same technique extends to other covering or packing

integer linear programs, as outlined in [GKM17]. A full description will appear in the journal version of [GKM17]. As some concrete examples, this implies poly$(\log n/\varepsilon)$-round deterministic LOCAL-model algorithms for $1+\varepsilon$ approximation of maximum independent set (as a sample packing problem) and for $1+\varepsilon$ approximation of minimum dominating set (as a sample covering problem). It should be remarked that the LOCAL model does not bound the time for local computation in one node and these two particular results take advantage of that.

### 3.3.3 Defective and Frugal Colorings

**Defective coloring**: The defective coloring problem is a variant of the standard proper coloring problem, which has turned out to be important in the study of distributed graph algorithms. In an $f$-defective coloring, we allow each node to have up to $f$ neighbors in its own color — in return for this relaxation, we hope for a smaller number of colors. Open Problem 11.7 in the book of Barenboim and Elkin asks for "*an efficient distributed algorithm for computing a $O(\Delta/p)$-defective $O(p)$-coloring*".

We note that an iterative-improvement algorithm of Lovász [Lov66]—which starts with an arbitrary coloring and changes node colors one by one, so long as that improves the node's defect— ensures the existence of such a defective coloring in all graphs. Kuhn [Kuh09] showed that a $\Delta/p$-defective $O(p^2)$ coloring can be computed in $O(\log^* n)$ rounds. Chung, Pettie, and Su [CPS17] gave a randomized algorithm that in $O(\log n)$ rounds computes an $O(\Delta/p)$-defective $O(p)$ coloring. By running their randomized algorithm through our derandomization result (Theorem 1.1), we get an efficient deterministic variant which settles Open Problem 11.7:

**Corollary 3.12.** *There is a deterministic distributed algorithm in the* LOCAL *model that, for any $p$, computes an $O(\Delta/p)$-defective $O(p)$ coloring in* poly$(\log n)$ *rounds.*

**Frugal coloring**: A $k$-frugal coloring is a coloring where each color appears at most $k$ times in the neighborhood of each node (independent of the color of that node itself, which is what makes this definition different from defective coloring). We are not aware of any deterministic distributed algorithm for frugal coloring (with good parameters), but there are some efficient randomized algorithms: Chung, Pettie, and Su [CPS17] show a randomized algorithm that computes an $O(\log^2 \Delta/\log\log \Delta)$-frugal $\Delta+1$ coloring in $O(\log n)$ rounds of the LOCAL model, and a $\beta$-frugal $O(\Delta^{1+1/\beta})$-coloring in $O(\log n \log^2 \Delta)$ rounds of the LOCAL model. By derandomizing these algorithms, we get

**Corollary 3.13.** *There are deterministic distributed algorithm that in* poly$(\log n)$ *rounds of the* LOCAL *model compute*

*(I) a $O(\log^2 \Delta/\log\log \Delta)$-frugal $\Delta+1$ coloring, and*

*(II) $\beta$-frugal $O(\Delta^{1+1/\beta})$-coloring.*

### 3.3.4 Forest Decomposition and Low Out-degree Orientation

Consider a graph with arboricity at most $a$, that is, a graph where edges can be decomposed into $a$ forests. Due to a result of Barenboim and Elkin [BE10], there is a deterministic distributed algorithm that decomposes any graph of arboricity $a$ into $2a$ forests, in $O(\log n)$ rounds. In Open Problem 11.10 of their book [BE13], Barenboim and Elkin ask for an "*efficient distributed algorithm for computing a decomposition of graph with arboricity $a$ into less than $2a$ forests*". A result of [GS17] provides a randomized poly$(\log n)$ round algorithm that decomposes the graph into $(1 + o(1))a$ forests, when $a = \Omega(\log n)$, and into $(1 + o(1))a$ pseudo-forests when $a = o(\log n)$.

Recall that a pseudo-forest is an undirected graph where each connected component has at most one cycle. In both cases, the decomposition provides an orientation of the edges where each node has out-degree at most $(1 + o(1))a$. To the best of our knowledge, in all distributed applications of the aforementioned forest decomposition, a decomposition into pseudo-forests (or alternatively, just the orientation with the bounded out-degree) would also suffice. Plugging this randomized algorithm into our derandomization result (Theorem 1.1), we get an algorithm that almost settles Open Problem 11.10 of [BE13]:

**Corollary 3.14.** *There is a deterministic* $\mathrm{poly}(\log n)$ *round algorithm in the* LOCAL *model that, for any graph with arboricity at most* $a$*, computes an orientation with maximum outdegree at most* $(1+o(1))a$*. Moreover, the algorithm decomposes the graph into* $(1+o(1))a$ *forests, if* $a = \Omega(\log n)$*, and into* $(1 + o(1))a$ *pseudo-forests if* $a = o(\log n)$*.*

### 3.3.5 Derandomizations in the CONGEST Model: Neighborhood Cover, Spanners, and Dominating Set

We have already mentioned that our network decomposition algorithm extends to the CONGEST model, and even has the nice property that each edge is in $\mathrm{poly}(\log n)$ many Steiner trees. We used these to derive our CONGEST model efficient deterministic MIS algorithm, in Theorem 3.3. But there is one more generality of our network decomposition, which opens the road for other applications: the algorithm readily extents to powers $G^k$ of the graph $G$, where we connect any two nodes within distance $G$. As stated in Theorem 2.12, in $\mathrm{poly}(\log n)$ rounds of the CONGEST model, we can compute a decomposition into clusters, each with a Steiner tree of depth $\mathrm{poly}(\log n)$, colored with $\mathrm{poly}(\log n)$ colors so that any two clusters wihin distance $k$ have different colors. Moreover, each edge is used in $\mathrm{poly}(\log n)$ Steiner trees. This can be directly plugged into some of the recent work on derandomization in the CONGEST model, for particular graph problems, to improve the related round complexities. We overview these next.

**Sparse neighbohood covers**: One prominent corollary is that we get an efficient deterministic algorithm in the CONGEST model for the *sparse neighborhood cover* problem — one of the central and versatile algorithmic tools in the study of locality-sensitive distributed graph algorithms [Pel00, AGLP89]. This corollary follows from using our improved network decomposition in the method provided by Ghaffari and Kuhn [GK18].

**Corollary 3.15.** *There is a deterministic distributed algorithm that for any radius* $r \geq 1$*, computes an* $\mathrm{poly}(\log n)$*-sparse neighborhood cover of the* $r$*-neighborhoods of the graph, with clusters of radius* $r\,\mathrm{poly}(\log n))$*, in* $r\,\mathrm{poly}(\log n)$ *rounds of the* CONGEST *model. In other words, this gives a clustering of the graph into overlapping clusters of radius* $r\,\mathrm{poly}(\log n)$ *such that for each node, its* $r$*-hop neighborhood is entirely contained in at least one of the clusters and moreover, each node is in at most* $\mathrm{poly}(\log n)$ *clusters.*

We note that the above neighborhood cover also settles a question of Elkin [Elk06], giving a deterministic variant of his minimum spanning tree algorithm with the same round complexity up to logarithmic factors.

**Spanner**: Another example is the first efficient deterministic distributed algorithm, in the CONGEST model, for constructing spanners with almost optimal parameters. This follows from plugging our network decomposition into the algorithms of [GK18]:

**Corollary 3.16.** *There is a deterministic distributed algorithm that in* $\mathrm{poly}(\log n)$ *rounds of the* CONGEST *model, computes a spanner with stretch* $2k - 1$ *and size* $O(kn^{1+1/k} \log n)$*.*

**Dominating set and set cover**: As another example, by putting together our CONGEST-model network decomposition with the work of Deurer et al. [DKM19], we get the first efficient deterministic CONGEST model approximation of minimum dominating set. Moreover, as outlined in [DKM19], this can also be extended to an approximation of set cover. These lead to the following corollary:

**Corollary 3.17.** *There are* $\mathrm{poly}(\log n)$*-round deterministic distributed algorithms in the* CONGEST *model that compute: (I) a* $(1 + o(1)) \log \Delta$ *approximation of minimum dominating set, where* $\Delta$ *denotes the maximum degree, and (II) a* $(1 + o(1)) \log \Delta$ *approximation of the minimum set cover problem, where* $\Delta$ *denotes the maximum set size.*

# Acknowledgment

# References

[ABCP96]  Baruch Awerbuch, Bonnie Berger, Lenore Cowen, and David Peleg. Fast network decompositions and covers. *J. of Parallel and Distributed Computing*, 39(2):105–114, 1996.

[ABI86]  Noga Alon, Laszlo Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7(4):567–583, 1986.

[AGLP89]  Baruch Awerbuch, Andrew V. Goldberg, Michael Luby, and Serge A. Plotkin. Network decomposition and locality in distributed computation. In *Proc. 30th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 364–369, 1989.

[BBH+19]  Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. Lower bounds for maximal matchings and maximal independent sets. In *Proc. Foundations of Computer Science (FOCS)*, 2019.

[BE10]  Leonid Barenboim and Michael Elkin. Deterministic distributed vertex coloring in polylogarithmic time. In *Proc. 29th Symp. on Principles of Distributed Computing (PODC)*, pages 410–419, 2010.

[BE11]  Leonid Barenboim and Michael Elkin. Deterministic distributed vertex coloring in polylogarithmic time. *Journal of the ACM (JACM)*, 58(5):23, 2011.

[BE13]  Leonid Barenboim and Michael Elkin. *Distributed Graph Coloring: Fundamentals and Recent Developments*. Morgan & Claypool Publishers, 2013.

[BEPS16]  Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. The locality of distributed symmetry breaking. *Journal of the ACM*, 63:20:1–20:45, 2016.

[CFG+19]  Yi-Jun Chang, Manuela Fischer, Mohsen Ghaffari, Jara Uitto, and Yufan Zheng. The complexity of $(\Delta + 1)$ coloring in congested clique, massively parallel computation, and centralized local computation. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 471–480. ACM, 2019.

[CHPS17]  Keren Censor-Hillel, Merav Parter, and Gregory Schwartzman. Derandomizing local distributed algorithms under bandwidth restrictions. In *31st International Symposium on Distributed Computing (DISC 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[CLP18]  Yi-Jun Chang, Wenzheng Li, and Seth. Pettie. An optimal distributed $(\Delta + 1)$-coloring algorithm? In *Proc. 50th ACM Symp. on Theory of Computing (STOC)*, 2018.

[CP17]  Yi-Jun Chang and Seth Pettie. A time hierarchy theorem for the LOCAL model. In *Proc. 58th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 156–167, 2017.

[CPS17]  Kai-Min Chung, Seth Pettie, and Hsin-Hao Su. Distributed algorithms for the lovász local lemma and graph coloring. *Distributed Computing*, 30(4):261–280, 2017.

[DKM19]  Janosch Deurer, Fabian Kuhn, and Yannic Maus. Deterministic distributed dominating set approximation in the congest model. In *Proc. Principles of Distributed Computing (PODC)*, pages 94–103, 2019.

[Elk06]  Michael Elkin. A faster distributed protocol for constructing a minimum spanning tree. *Journal of Computer and System Sciences*, 72(8):1282–1308, 2006.

[FG17]  Manuela Fischer and Mohsen Ghaffari. Sublogarithmic distributed algorithms for Lovász local lemma, and the complexity hierarchy. In *Proc. 31st Symp. on Distributed Computing (DISC)*, pages 18:1–18:16, 2017.

[FGK17]  Manuela Fischer, Mohsen Ghaffari, and Fabian Kuhn. Deterministic distributed edge-coloring via hypergraph maximal matching. In *Proc. 58th IEEE Symp. on Foundations of Computer Science (FOCS)*, 2017.

[Gha16]  Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In *Proc. 27th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 270–277, 2016.

[GHK18]  Mohsen Ghaffari, David Harris, and Fabian Kuhn. On derandomizing local distributed algorithms. In *Proc. Foundations of Computer Science (FOCS)*, pages 662–673, 2018.

[GK18]  Mohsen Ghaffari and Fabian Kuhn. Derandomizing distributed algorithms with small messages: Spanners and dominating set. In *32nd International Symposium on Distributed Computing (DISC 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[GKM17]  Mohsen Ghaffari, Fabian Kuhn, and Yannic Maus. On the complexity of local distributed graph problems. In *Proc. 49th ACM Symp. on Theory of Computing (STOC)*, pages 784–797, 2017.

[GKU19]  Mohsen Ghaffari, Fabian Kuhn, and Jara Uitto. Conditional hardness results for massively parallel computation from distributed lower bounds. In *Proc. Foundations of Computer Science (FOCS)*, 2019.

[GS17]     Mohsen Ghaffari and Hsin-Hao Su.  Distributed degree splitting, edge coloring, and orientations. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2505–2523. Society for Industrial and Applied Mathematics, 2017.

[KK19]     Dariusz Kowalski and Piotr Krysta. Deterministic coloring algorithms in the local model. *arXiv preprint arXiv:1907.12857*, 31 July 2019.

[KMW16]    Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local computation: Lower and upper bounds. *Journal of the ACM*, 63(2), 2016.

[KSV10]    Howard Karloff, Siddharth Suri, and Sergei Vassilvitskii. A model of computation for mapreduce. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 938–948. SIAM, 2010.

[Kuh09]    Fabian Kuhn. Weak graph colorings: distributed algorithms and applications. In *Proceedings of the twenty-first annual symposium on Parallelism in algorithms and architectures*, pages 138–144. ACM, 2009.

[Lin87]    Nathan Linial.  Distributive graph algorithms – global solutions from local data.  In *Proc. 28th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 331–335, 1987.

[Lin92]    Nati Linial.  Locality in distributed graph algorithms.  *SIAM Journal on Computing*, 21(1):193–201, 1992.

[Lov66]    László Lovász.  On decomposition of graphs.  *Studia Sci. Math. Hungar*, 1(273):238, 1966.

[LS93]     Nati Linial and Michael Saks.  Low diameter graph decompositions.  *Combinatorica*, 13(4):441–454, 1993.

[Lub86]    Michael Luby.  A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15:1036–1053, 1986.

[MT10]     Robin Moser and Gabor Tardos.  A constructive proof of the general Lovász local lemma. *Journal of the ACM*, 57(2):11, 2010.

[NS95]     Moni Naor and Larry Stockmeyer.  What can be computed locally?  *SIAM Journal on Computing*, 24(6):1259–1277, 1995.

[Pel00]    David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.

[PS92]     Alessandro Panconesi and Aravind Srinivasan.  Improved distributed algorithms for coloring and network decomposition problems. In *Proc. 24th ACM Symp. on Theory of Computing (STOC)*, pages 581–592, 1992.

# A    Comparison of previous work with Theorem 1.2

Recall that before our work, the state of the art deterministic algorithm for network decomposition was the $2^{O(\sqrt{\log n})}$-round algorithm of Panconesi and Srinivasan [PS92]. This result itself

was a refinement of the approach of Awerbuch et al. [AGLP89], who obtained an $2^{O(\sqrt{\log n \log \log n})}$-round algorithm for network decomposition. In this section, we first briefly recall this approach, in Appendix A.2. Then, in Appendix A.3, we describe a new algorithm that also achieves this $2^{O(\sqrt{\log n \log \log n})}$round complexity. In some sense, the approaches of [PS92, AGLP89] and the algorithm that we describe here can be viewed as two fundamentally different methods, both of which seem to get stuck at the $2^{O(\sqrt{\log n \log \log n})}$round complexity. Finally, in Appendix A.4 then discuss how these compare with the approach of Theorem 2.3 that achieves a $poly(\log n)$ round complexity.

## A.1 Recap on Ruling sets

Both the approach of Awerbuch et al. [AGLP89] and of the new algorithm that we present here would achieve round-complexity $2^{O(\sqrt{\log n})}$, assuming we have an efficient algorithm for maximal independent set (MIS). However, the only known way of computing MIS is via network decomposition (cf. Theorem 3.1). Hence, instead, we will use a *ruling set* procedure. As we recall next, ruling set is a certain weakening of MIS, and it can be computed deterministically in $O(\log n)$ rounds [AGLP89] (for certain parameters).

**Definition A.1.** *Given a graph $G$, an $(\alpha, \beta)$-ruling set is a subset $S \subseteq G$ such that*

- *for each two different vertices $u, v \in S$ their distance in $G$ is at least $\alpha$,*
- *for each vertex $u \in G$ there is a vertex $v \in S$ such that the distance of $u$ and $v$ in $G$ is at most $\beta$.*

To give an example, a maximal independent set is simply a $(2, 1)$-ruling set. For the sake of completeness, we now show how to construct a $(2, O(\log n))$-ruling set, as presented by Awerbuch et al. [AGLP89]. In other words we find a set $S$ such that no two of its vertices are neighbouring and for each vertex $v$ there is a vertex $u \in S$ within its $O(\log n)$ distance.

**Proposition A.2.** *There is a deterministic $O(\log n)$-round complexity algorithm that computes a $(2, O(\log n))$-ruling set of any graph $G = (V, E)$ in the LOCAL model, given that each vertex holds a unique $O(\log n)$-bit identifier.*

*Proof.* We start with a set $S_0 = V$ of all vertices and gradually prune this set in $O(\log n)$ steps, producing a chain $V = S_0 \supseteq S_1 \supseteq \cdots \supseteq S_{O(\log(n))} =: S$, where the final set $S$ in the chain is the desired $(2, O(\log n))$-ruling set.

This is done as follows: in the $i$-th step, each vertex $u \in S_{i-1}$ that has 1 at the $i$-th position of its identifier checks whether it has a neighbour with identifier containing 0 at its $i$-th bit. If this is the case, $u$ decides to leave $S_{i-1}$.

This process clearly finishes with a set $S_{O(\log n)}$ without any neighbouring vertices. To see that for each vertex $v \in G$ we can find $u \in S$ at distance $O(\log n)$, consider the previous process again and for each vertex $w$ that is removed from the ruling set at some iteration $i$, draw an oriented edge towards its neighbour that made $w$ disappear from $S_{i-1}$. These edges form disjoint oriented trees of depth $O(\log n)$ and with vertices of $S$ being their roots. This implies that $S$ is indeed $(2, O(\log n))$-ruling set. $\square$

## A.2 Recap on the Approach of Awerbuch et al. [AGLP89]

Next, we sketch the $2^{O(\sqrt{\log n \log \log n})}$-round algorithm for network decomposition from Awerbuch et al. [AGLP89]. We note that the algorithm of Panconesi and Srinivasan [PS92] is a improvement and refinement of this approach; while the algorithms is different, the overall approach outline

is the same. In particular, both of these algorithms seem to fall short of reaching a poly($\log n$) complexity, for the same reasons.

**Proposition A.3.** *There is a deterministic $2^{O(\sqrt{\log n \log\log n})}$-round algorithm in the* LOCAL *model that computes a network decomposition of the underlying graph $G$ into $O(\log n)$ color classes, each one with clusters of diameter at most $O(\log n)$.*

*Proof Sketch.* It suffices to construct a decomposition with $2^{O(\sqrt{\log n \log\log n})}$ color classes and clusters of weak-diameter $2^{O(\sqrt{\log n \log\log n})}$[7], since then we can reduce the diameter and the number of color classes by known reduction [ABCP96]. See also Theorem 2.14 for such a transformation.

We show an iterative process with a parameter $d$ that will have $\log_d n$ iterations. The $i$-th iteration has round complexity $O(\log n)^{i-1} \cdot O(d^2 + \log^* n)$. This process will construct network decomposition with $d \log_d n$ colors and with the diameter of each cluster being $O(\log n)^{\log_d n}$. We optimize $d$ by setting $d = 2^{O(\sqrt{\log n \log\log n})}$ which gives the advertised parameters.

Now we describe one phase of the process that works with a graph $G_i = (V_i, E_i)$. Initially, we set $G_0 := G$. We shall ensure that one round of communication in $G_i$ can be simulated in $O(\log n)^{i-1}$ rounds of communication in $G$.

We split the vertex set $V_i$ into the set $H_i$ of vertices of degree at least $d$ and $L_i$ of vertices of degree smaller than $d$. To deal with low degree vertices, we construct their $d^2$-coloring by the algorithm of Linial [Lin87] in time $O(\log n)^{i-1} \cdot O(d^2 + \log^* n)$ – here the first term is the number of rounds in $G$ we need to simulate one round in $G_{i-1}$ and the second term is the complexity of Linial's algoirthm. Each constructed color class will constitute one class of the final network decomposition.

Next, the nodes locally construct a graph $G_i^2$ where two vertices are connected if their distance in $G_i$ is at most 2. Now we use Proposition A.2 to find a $(2, O(\log n))$ ruling set in $G_i^2$ restricted to vertices of $H_i$. The ruling set procedure also gives us a decomposition of $H_i$ into oriented trees of depth $O(\log n)$, hence we get a decomposition of vertices of $H_i$ into clusters of diameter $O(\log n)$ in $G_i$. These clusters will be vertices of $G_{i+1}$ and there is an edge in $E(G_{i+1})$ for each pair of clusters adjacent in $G_i$.

In the next round, the additional communication overhead of the algorithm on $G^{i+1}$ is only $O(\log n)$ times higher than for $G^i$. More precisely, each vertex of $G^{i+1}$ corresponds to a set of vertices of $G$ that have weak-diameter at most $O(\log n)^{i+1}$. To finish the argument we need to bound the number of iterations. This follows from the fact that each vertex in the ruling set found in the $i$-th iteration has degree at least $d$ in $G_i$. Moreover, the set of its neighbours is disjoint with the neighbours of all other vertices in the ruling set due to its construction in graph $G_i^2$. Hence, the number of vertices in $G_{i+1}$ is at least $(d+1)$ times smaller than the number of vertices of $G_i$. Thus the number of iterations is bounded by $\log_d n$. $\square$

## A.3   A New Network Decomposition with Complexity $2^{O(\sqrt{\log n \log\log n})}$

We now propose a new algorithm, which appears to be fundamentally different and it also achieves the same $2^{O(\sqrt{\log n \log\log n})}$ round complexity. In the next subsection, we will contrast these two approaches with the poly($\log n$)-time algorithm described in Theorem 2.3.

Similarly to the case of Theorem 1.2, we only show how to find first $O(\log n)$ color classes of the resulting decomposition that ensure that the overall number of uncolored vertices drop by a factor of $(1 - 2^{-O(\sqrt{\log n \log\log n})})$. The result then follows after repeated application of Proposition A.4 and,

---

[7]I.e., any two vertices of a cluster have $2^{O(\sqrt{\log n \log\log n})}$ distance in $G$.

finally, after reduction of the number of colour classes and their diameter using the transformations of [ABCP96] (See also Theorem 2.14 for such a transformation).

The idea of the new algorithm is to simulate the sequential process of building network decomposition from the proof of Theorem 2.14. In particular, each vertex $u$ will consider a ball $B(u, r(u))$ around itself of radius $r(u)$. This is similar to the sequential ball carving process in the proof of Theorem 2.14. The difference is that, here, instead of additively growing the radius of the ball by one in each step, we will consider radii that are powers of a small value $t$ (later fixed to be $t = O(\log n)$). We discuss this choice after the proof of Proposition A.4. The balls of different vertices will possibly intersect considerably. The idea is then to output only a subfamily $I := \{B(u_1, r(u_1)), B(u_2, r(u_2)), \ldots, B(u_k, r(u_k))\}$ composed of balls that do not intersect nor touch, pairwise. In the actual proof we follow this idea, though the resulting output is little bit more complicated.

The main challenge here is to ensure that the family $I$ covers a substantial part of the whole graph. To this end, we need a stronger condition, which is effectively like thinking of a "wider boundary" for each ball. Concretely, instead of requiring a lower bound on the ratio $|B(u, r(u))|/|B(u, r(u) + 1)|$ — as we had in the sequential algorithm of Theorem 2.14 and our main distributed algorithm from Theorem 1.2 — we require a lower bound on the ratio $|B(u, r(u))|/|B(u, O(\log n) \cdot r(u)|$. This allows us to find such a good set $I$ with the help of ruling sets. The above condition leads to an optimization problem giving the same round complexity as in [AGLP89].

**Proposition A.4.** *There is a deterministic* LOCAL*-model algorithm that, for any graph $G = (V, E)$ with $n$ nodes, in $2^{O(\sqrt{\log n \log \log n})}$ rounds, computes a $S \subseteq V$ of vertices and a coloring of this set with $O(\log n)$ colors such that (1) each each connected component of any color class of $S$ has weak-diameter $2^{O(\sqrt{\log n \log \log n})}$, and (2) $|S| \geq |V|/2^{O(\sqrt{\log n \log \log n})}$.*

*Proof.* First fix $t = O(\log n)$ to be such that Proposition A.2 can give us $(2, t/4)$-ruling sets in $O(\log n)$ rounds. Moreover, set $\varepsilon = 2^{-\sqrt{\log n/\log \log n}}$. Each vertex $u$ gradually grows a sequence of balls $B(u, t^i)$ for $i \geq 0$. It stops growing the ball around it at the first point when $|B(u, t^i)| \geq \varepsilon \cdot |B(u, t^{i+1})|$ and then sets $r(u) := t^i$. Note that the maximum $i$ the vertex needs to consider is bounded by

$$\log_{1/\varepsilon} n = \frac{\log n}{\log 1/\varepsilon} = \sqrt{\log n \log \log n}.$$

This is because, in each step, the vertex either finishes growing or the volume of its ball grows by multiplicative factor of at least $1/\varepsilon$. Therefore, the overall radius the vertex needs to consider and, thus also the round complexity of this step, is bounded by

$$t^{\log(1/\varepsilon)} = 2^{O(\log \log n) \cdot \sqrt{\log n/\log \log n}} = 2^{O(\sqrt{\log n \log \log n})}.$$

Next, we define $\log_{1/\varepsilon} n + 1$ graphs $G_0, G_1, \ldots, G_{\log_{1/\varepsilon} n}$, where each vertex $u$ is in the vertex set of $G_i$ if and only if $r(u) = t^i$. Two vertices $u, v \in G_i$ are connected in $G_i$ if their distance in $G$ is at most $3 \cdot t^i$. This is to ensure that if there is not an edge between $u$ and $v$ in $G_i$, the corresponding balls $B(u, r(u))$ and $B(v, r(v))$ neither overlap, nor are adjacent.

Now for each $G_i$ we run the $(2, t/4)$-ruling set algorithm of Proposition A.2 on $G_i$ in $O(\log n) \cdot 2^{O(\sqrt{\log n \log \log n})} = 2^{O(\sqrt{\log n \log \log n})}$ rounds. We call the output ruling set $R_i$ and define the set $S$ of vertices we color as

$$S = \bigcup_{0 \leq i \leq \log_{1/\varepsilon} n} \bigcup_{u \in R_i} B(u, r(u)).$$

26

Moreover, each vertex $v \in S$ is colored by the smallest index $i$ such that $v$ is in some ball $B(u, r(u))$ for a vertex $u \in R_i$.

Next, we argue that this output satisfies all the desired conditions. First, notice that the number of color classes we output is $\log_{1/\varepsilon} n = \sqrt{\log n \log\log n} = O(\log n)$. For each color class we output a subset of balls with diameter $2^{O(\sqrt{\log n \log\log n})}$ that were neither overlapping, nor adjacent due to the construction of the ruling set and hence their subsets have still bounded weak-diameter and they are neither overlapping, nor adjacent.

It remains to be seen that the number of vertices of $S$ constitutes a substantial fraction of the vertices of $G$. The crucial observation is that due to the property of our ruling sets, each $v \in G_i$ has distance in $G$ at most $t/4 \cdot 3t^i < t^{i+1}$ to some $u \in R_i$. Hence, while $S = \bigcup_{0 \leq i \leq \log_{1/\varepsilon} n} \bigcup_{u \in R_i} B(u, r(u))$ is a union of $O(\log n)$ disjoint sets, the union $\bigcup_{0 \leq i \leq \log_{1/\varepsilon} n} \bigcup_{u \in R_i} B(u, t \cdot r(u))$ already covers the whole $G$. Finally, recall that each ball $B(u, r(u))$ already substitutes a substantial fraction of the bigger ball $B(u, t \cdot r(u))$. This enables us to finish off with the following simple double counting argument.

$$
\begin{aligned}
|S| &= \left| \bigcup_{0 \leq i \leq \log_{1/\varepsilon} n} \bigcup_{u \in R_i} B(u, r(u)) \right| \\[2mm]
&\geq \frac{1}{\log_{1/\varepsilon} n} \sum_{0 \leq i \leq \log_{1/\varepsilon} n} \left| \bigcup_{u \in R_i} B(u, r(u)) \right| \qquad \text{each vertex is overcounted at most } \log_{1/\varepsilon} n \text{ times} \\[2mm]
&= \frac{1}{\log_{1/\varepsilon} n} \sum_{0 \leq i \leq \log_{1/\varepsilon} n} \sum_{u \in R_i} |B(u, r(u))| \qquad \text{balls with centers from the same ruling set are disjoint} \\[2mm]
&\geq \frac{\varepsilon}{\log_{1/\varepsilon} n} \sum_{0 \leq i \leq \log_{1/\varepsilon} n} \sum_{u \in R_i} |B(u, t \cdot r(u))| \quad |B(u, r(u))| \geq \varepsilon |B(u, t \cdot r(u))| \\[2mm]
&\geq \frac{\varepsilon}{\log_{1/\varepsilon} n} \left| \bigcup_{0 \leq i \leq \log_{1/\varepsilon} n} \bigcup_{u \in R_i} B(u, t \cdot r(u)) \right| \\[2mm]
&\geq \frac{\varepsilon}{\log_{1/\varepsilon} n} \cdot n \qquad \text{this union covers the whole graph} \\[2mm]
&= n/2^{O(\sqrt{\log n \log\log n})}.
\end{aligned}
$$

$\square$

Note that choosing the radius as a power of $t$, instead of additive increases similar to the sequential ball carving process, is crucial in the above proof. This is because it is necessary to ensure that not only the boundary of each ball $B(u, r(u))$ is small as in the sequential ball carving process, but also the volume of the ball $B(u, t \cdot r(u))$ is not much bigger than the volume of $B(u, r(u))$, as the factor $|B(u, r(u))|/|B(u, t \cdot r(u))|$ is also the fraction of all vertices that we output in the end output as the set $S$.

## A.4 Contrasting the $\text{poly}(\log n)$-round and $2^{O(\sqrt{\log n \log\log n})}$-round algorithms

In this subsection, we discuss the similarities in the above two approaches and describe why they both fall short of reaching a $\text{poly}(\log n)$ round complexity. In particular, we discuss one particular

graph on which both of these approaches get stuck at a $2^{O(\sqrt{\log n})}$ complexity. We note that the discussions in this subsection are not written formally, and they instead try to provide an intuitive explanation of the shortcoming of the two approaches, which is circumvented by the poly($\log n$)-round algorithm of Theorem 2.3.

First, let us note the similarities between the algorithm from [AGLP89] and Proposition A.4. Both algorithms use a ruling set computation just as a replacement for a maximal independent set, which allows efficient computation. It is straightforward to observe that, if we could solve the maximal independent set problem in poly($\log n$) rounds, then the complexity of both of these approaches would improve from $2^{O(\sqrt{\log n \log \log n})}$ to $2^{O(\sqrt{\log n})}$. This is roughly because, in that case, per iteration, we would lose an $O(1)$ factor instead of an $O(\log n) = O(2^{\log \log n})$ factor. For the approach of [AGLP89], the algorithm of [PS92] cleverly circumvents this issue via additional insights (which recursively build the desired maximal independent set, at each point, using the network decomposition that is constructed up to that point).

Moreover, both algorithms cannot bound the increase of the radii of growing clusters in an additive manner. Instead, they both the increase in the radii by a multiplicative factor of $O(\log n)$ (or $O(1)$ should they have access to an efficient algorithm for maximal independent set). They have to achieve as much as possible for one multiplicative increase and here their strategies differ: The approach of Awerbuch et al. [AGLP89] uses iterations of *contracting clusters*, while the algorithm of Proposition A.4 uses a sped up *ball carving*. More concretely, the algorithm of [AGLP89] decreases the overall number of vertices by a factor of roughly $2^{O(\sqrt{\log n})}$, via contractions, the algorithm from Proposition A.4 multiplies volume of a given ball by the same factor, via growing the ball's radius.

Despite this departure, it appears that both approaches fall short of reaching a poly($\log n$) round complexity for similar reasons. In particular, next we discuss an example graph $G$ that is "hard" for both of these approaches; that is, both approaches achieves at best a $2^{O(\sqrt{\log n})}$ round complexity on this particular graph. This is a simple high-dimensional torus-like graph $G$, as follows: The vertices of $G$ are vectors of length $\sqrt{\log n}$ with coordinates from the set $\{0, 1, \ldots, 2^{\sqrt{\log n}} - 1\}$. Two vertices/vectors are connected via an edge if and only if the difference in each coordinate is at most one modulo $2^{\sqrt{\log n}}$. We will refer to the first parameter, i.e., the length of vectors which is set to $\sqrt{\log n}$, as the *dimension*, and to the second parameter, i.e., the maximum coordinate value which is set to $2^{\sqrt{\log n}}$, as the *side length*.

The vertex-transitive graph $G$ has the following two properties. First, its diameter is $\Theta(2^{\sqrt{\log n}})$. Second, the graph has a "volume expansion" factor of $\Theta(2^{\sqrt{\log n}})$ as we double the radius. That is, for each vertex $u \in V(G)$ and for any radius $r \leq \Theta(2^{\sqrt{\log n}})$, we have $|B(u, 2r)|/|B(u, r)| = (4r + 1)^{\sqrt{\log n}}/(2r + 1)^{\sqrt{\log n}} = \Theta(2^{\sqrt{\log n}})$. In fact, $G$ optimizes the trade-off between these two parameters, diameter and "volume expansion".

The two $2^{O(\sqrt{\log n \log \log n})}$-round algorithms in fact optimize the same trade-off and we now observe that the approaches used there cannot yield $2^{o(\sqrt{\lg n})}$ running time. Let us first observe it on the example of algorithm from Proposition A.4. There, we are growing a ball around each vertex and doubling its diameter at every step. If we continue doubling the ball until it covers the whole graph, its radius reaches the diameter of $G$ and we will need $O(2^{\sqrt{\log(n)}})$ rounds. If we, instead, stop growing the ball with any radius $r$, we will have $|B(u, O(\log n) \cdot r)| \geq |B(u, 2r)| = \Theta(2^{\sqrt{\log n}}) \cdot |B(u, r)|$. Hence, we can guarantee that only $\Theta(1/2^{\sqrt{\log n}})$ fraction of vertices will be colored in one phase of the algorithm.

Similarly, if we run the algorithm of Awerbuch et al. [AGLP89], we will either decide that all vertices of $G$ have small degree and color them in $2^{\Theta(\sqrt{\log(n)})}$ rounds, or we decide that they have high degree. In the latter case, we compute a ruling set and use it to contract vertices and form a new cluster graph. However, one can see that for an arbitrary ruling set — e.g., if the ruling

set is given by vertices with all coordinates equal 0 modulo 3 and vertices outside ruling sets join the cluster of the unique adjacent vertex in the ruling set — the new cluster graph will again be a $\sqrt{\log(n)}$-dimensional torus, this time with coordinates from the set (up to rounding errors) $\{0, 1, \ldots, 2^{\sqrt{\log(n)}}/3\}$. That is, the new contracted graph will also be a torus-like graph with the same structure described above and only with a constant factor smaller side length. Even after we contract vertices for $O(\sqrt{\log(n)})$ repetitions, we will still have a torus with dimension $\sqrt{\log(n)}$ and side length $2^{\Theta(\sqrt{\log(n)})}$. But already at this point, simulating one round of communication on this contracted graph takes $2^{\Theta(\sqrt{\log(n)})}$ communication rounds in the original graph. Hence, the algorithm cannot achieve a round complexity $2^{o(\sqrt{\log(n)})}$.

Our algorithm from Theorem 2.3 circumvents this issue by growing the clusters more carefully: crucially, we get only additive increase in diameter of each cluster per step, instead of a multiplicative increase as in the examples above.

# B    CONGEST Network Decomposition for Power Graphs

Here, we present the formal proof of Theorem 2.12. This formal proof expands on the proof sketch provided in Section 2.2 and provides addition low-level details.

*Proof of Theorem 2.12.* We adopt the algorithm from Lemma 2.4 and the notation used throughout the proof; we also apply the lemma as in the proof of Theorem 2.3. The only change is that the process we run in a given step of a given phase will involve all red nodes at distance $k$ from some blue node, instead of only red nodes neighbouring to blue nodes in the original algorithm. More concretely, all the red nodes in $k$-hop distance of some blue node propose to some blue cluster. This is done as follows.

We describe a process with $k$ *iterations* that we run in a given step of a given phase. The process can be thought of as a variant of a breadth first search (BFS) algorithm run from all blue nodes at once.

In the first iteration, each blue node starts with a token with the label $Y$ of its cluster $S'(Y)$ (we dropped the index $i$ from the original notation $S'_i(Y)$ since we already fixed a phase) and it sends this token to all of its neighbours.

In the iterations 2 to $k$, the following happens. If at any point of the algorithm any node $u$ from $G$ (here we consider even dead nodes that generally include also nodes of the host graph that were already colored) receives *for the first time* some nonzero number $t$ of tokens, say a set $\{Y_1, Y_2, \ldots, Y_t\}$ with the label $Y_1$ being the smallest, it does the following.

- If $u$ is a living blue node, it does not do anything.
- If $u$ is a living red node, it adds itself as a new terminal node to the underlying Steiner tree of the cluster $S'(Y_1)$ together with an oriented edge pointing towards some node $v$ that sent a token with $Y_1$. The node then sends a token $Y_1$ to all of its neighbours.

  The node $u$ will later propose to the blue cluster with the identifier $Y_1$. To propose, the node will broadcast via the Steiner tree of the cluster $S'(Y_1)$ that it just joined.

- If $u$ is a dead node, after receiving tokens in the iteration $i$, $u$ first checks whether in the previous run it received tokens also in the iteration $i$, or later. If the latter is the case, it adds itself as a new nonterminal node to the underlying Steiner tree of the cluster $S'(Y_1)$ together with an oriented edge pointing towards some node $v$ that sent a token with $Y_1$. The node then sends a token $Y_1$ to all of its neighbours.

However, if, during the last BFS, $u$ received some tokens during the same iteration $i$ and chose to forward a token $Y_{\text{last}}$, it will forward the same token also this time, not considering the tokens it actually received.

We will see that it is not possible that $u$ received tokens earlier than in the $i$-th iteration in the previous run of BFS via standard argument about correctness of BFS.

If the node already received some tokens during this breadth first search algorithm, it does not receive or send any more tokens.

After the breadth first search algorithm finishes, roots of all Steiner trees collect the number of proposing red nodes and each root decides to either accept all proposing red vertices and recolor them to blue, or it makes them die and stops growing with the same decision rule as in Lemma 2.4. The Steiner trees, however, stay the same even if some of its vertices die, the red nodes that died are just labeled as nonterminals. This finishes the description of one step of current phase.

Besides these changes, the algorithm (the number of phases, steps in each phase and decisions of blue clusters whether to grow or not) stays the same.

**Analysis**: To argue about the correctness of the new version of the algorithm, we first check that some basic properties of BFS.

**Observation B.1.** *Throughout the course of the algorithm, the following holds:*

1. *The underlying Steiner trees of clusters are indeed trees oriented towards the root of the cluster throughout the course of the algorithm.*

2. *In the $i+1$-th iteration tokens are sent exactly by vertices whose distance from the set of all blue vertices is $i$.*

3. *If a node $u$ sends a token $Y$ in some iteration, it belongs to the Steiner tree of $S'(Y)$ as either terminal or nonterminal node.*

4. *If in some step $j$ a dead node $u$ added itself to the Steiner tree of cluster $S'(Y)$ and later in step $\tilde{j} > j$ it added itself to the Steiner tree of cluster $S'(\tilde{Y})$, the distance of $u$ to the set of blue vertices was strictly smaller in step $\tilde{j}$ than in step $j$.*

*Proof.* The first two bullet points follow from the standard analysis of BFS. For the third bullet point, we recall that if a dead node $u$ sends a token $Y_{\text{last}}$ in some iteration, then node $u$ is already part of the Steiner tree of $S'(Y_{\text{last}})$ by induction.

The last bullet point follows from the fact that a dead node adds itself to a Steiner tree only if it receives token in an earlier iteration than during the previous step, which by the second bullet point means that the distance of $u$ to the set of blue clusters decreased. ☐

Now we can mostly replicate the proof of Lemma 2.4. We state equivalents of observations from the proof of Lemma 2.4 and argue that they are satisfied if they differ substantially from the arguments for Lemma 2.4. To argue about $k$-hop separation of the clusters, instead of the invariants (I) and (II) from Lemma 2.4 we keep stronger invariants (I') and (II'). We keep invariant (III) the same.

(I') For each $i$-bit string $Y$, the set $S_i'(Y) \subseteq S_i'$ of all living nodes whose label ends in suffix $Y$ has no other living nodes $S_i' \setminus S_i'(Y)$ in its $k$-hop neighbourhood.

In other words, the set $S_i'(Y)$ is a union of some connected components of the subgraph $G[S_i']$ induced by living nodes $S_i'$ and in the $k$-hop neighbourhood in $G$ around $S_i'(Y)$ all nodes are

either dead or they do not belong to the set $S$ (they were colored by previous application of the algorithm).

(II') For each label $L$ and the corresponding cluster $S_i'(L)$, the related Steiner tree $T_L$ has radius at most $i \cdot k \cdot R$, where $R = O(\log^2 n)$.

(III) We have $|S_{i+1}'| \geq |S_i'|(1 - 1/2b)$.

Now we repeat the list of observations from the analysis of Lemma 2.4 and remark on them whenever they differ from their counterparts.

**Observation B.2.** *Any blue cluster stops after at most $4b \log n$ steps.*

*Proof.* The proof stays the same as in Observation 2.5. □

**Observation B.3.** *Once a blue cluster $A$ stops growing, there is no red node in its $k$-hop neighbourhood in $G$ and there never will be one in this phase.*

*Proof.* As in Lemma 2.4, consider the step in which cluster $A$ stops. In that step, each red node in its $k$-hop neighbourhood in $G$ (if there is one) either requested to join $A$ or some other blue cluster. This is because even if a dead node $u$ did not forward the token with the identifier of $A$, it forwarded a token of a blue cluster $A'$ that is of the same distance to $u$ or even closer.

Any red node $v$ that requested to join some blue cluster either adopts a blue label or dies. In either case, $v$ is not a living red node anymore (and it will never become one). From this point onward, this blue cluster $A$ never grows or shrinks. □

**Observation B.4.** *In each step, the radius of the Steiner tree of each blue cluster grows by at most $k$, while the radius of the Steiner tree of each red cluster does not grow. This implies invariant (II').*

**Observation B.5.** *The total number of red vertices in $S_i'(Y)$ that die during this phase is at most $|S_i'(Y)|/(2b)$. This implies invariant (III).*

*Proof.* The proof stays the same as in Observation 2.8. □

Now we can bound the number of Steiner trees that use each particular edge.

**Observation B.6.** *In construction of one color class of the decomposition, each edge is used in $O(\log n \cdot \min(k + \log^2 n))$ Steiner trees. Thus, overall, each edge is used in $O(\log^2 n \min(k + \log^2 n))$ Steiner trees.*

*Proof.* We run through $O(\log n)$ phases and in each phase we run $O(\log^2 n)$ steps. First, note that during each step, each edge $uv$ will be used by at most one additional Steiner tree. This is because the edge is added only in the case when either $u$ sends a token to $v$, or the other way around.

We also claim that during all the steps of a given phase, one edge $uv$ is used for a Steiner tree at most $k$ times. This follows from the last bullet point of Observation B.1. □

Finally, we bound the running time. We have $O(\log n)$ color classes, each one is constructed in $O(\log n)$ phases, where each phase has $O(\log^2 n)$ steps. For each step, we need to run breadth first search for $O(k)$ steps and broadcast information to root via Steiner trees, which takes $O(k \log^3 n \cdot \log n \min(k + \log^2 n))$ where the first term is the diameter of the underlying Steiner tree that is bounded by Observation B.4 and the second term is due to number of Steiner trees per edge that was bounded by Observation B.6. This implies running time $O(k \log^8 n \min(k + \log^2 n))$. □