

Faster Deterministic Distributed MIS and Approximate Matching

Conference Paper**Author(s):**

Ghaffari, Mohsen; [Grunau, Christoph](#) 

Publication date:

2023-06-02

Permanent link:

<https://doi.org/10.3929/ethz-b-000620429>

Rights / license:

[Creative Commons Attribution-NoDerivatives 4.0 International](#)

Originally published in:

<https://doi.org/10.1145/3564246.3585243>

Funding acknowledgement:

853109 - Distributed and Massively Parallel Graph Algorithms (EC)



Faster Deterministic Distributed MIS and Approximate Matching

Mohsen Ghaffari
MIT
Cambridge, MA, USA
ghaffari@mit.edu

Christoph Grunau
ETH Zurich
Zurich, Switzerland
cgrunau@inf.ethz.ch

ABSTRACT

We present an $\tilde{O}(\log^2 n)$ round deterministic distributed algorithm for the maximal independent set problem. By known reductions, this round complexity extends also to maximal matching, $\Delta + 1$ vertex coloring, and $2\Delta - 1$ edge coloring. These four problems are among the most central problems in distributed graph algorithms and have been studied extensively for the past four decades. This improved round complexity comes closer to the $\tilde{\Omega}(\log n)$ lower bound of maximal independent set and maximal matching [Balliu et al. FOCS '19]. The previous best known deterministic complexity for all of these problems was $\Theta(\log^3 n)$. Via the shattering technique, the improvement permeates also to the corresponding randomized complexities, e.g., the new randomized complexity of $\Delta + 1$ vertex coloring is now $\tilde{O}(\log^2 \log n)$ rounds.

Our approach is a novel combination of the previously known (and seemingly orthogonal) two methods for developing fast deterministic algorithms for these problems, namely *global derandomization via network decomposition* (see e.g., [Rozhon, Ghaffari STOC'20; Ghaffari, Grunau, Rozhon SODA'21; Ghaffari et al. SODA'23]) and *local rounding of fractional solutions* (see e.g., [Fischer DISC'17; Harris FOCS'19; Fischer, Ghaffari, Kuhn FOCS'17; Ghaffari, Kuhn FOCS'21; Faour et al. SODA'23]). We consider a relaxation of the classic network decomposition concept, where instead of requiring the clusters in the same block to be non-adjacent, we allow each node to have a small number of neighboring clusters. We also show a deterministic algorithm that computes this relaxed decomposition faster than standard decompositions. We then use this relaxed decomposition to significantly improve the integrality of certain fractional solutions, before handing them to the local rounding procedure that now has to do fewer rounding steps.

CCS CONCEPTS

• **Theory of computation** → **Distributed algorithms.**

KEYWORDS

distributed algorithms, maximal independent set, matching, derandomization, rounding

*CG received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 853109).



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

STOC '23, June 20–23, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9913-5/23/06.

<https://doi.org/10.1145/3564246.3585243>

ACM Reference Format:

Mohsen Ghaffari and Christoph Grunau. 2023. Faster Deterministic Distributed MIS and Approximate Matching. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC '23), June 20–23, 2023, Orlando, FL, USA*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3564246.3585243>

1 INTRODUCTION

In this paper, we present a faster deterministic distributed algorithm for the Maximal Independent Set (MIS) problem, which is one of the most central problems in distributed graph algorithms and has been studied extensively for the past four decades. This improvement has implications for several other problems. Our main novelty is a technique that combines the two previously known general approaches, which seemed unrelated and incompatible hitherto. We are hopeful that this new technique finds applications in a wider range of problems. We next set up the context, and then discuss our results and approach.

Distributed Model: We work with the standard synchronous message-passing model of distributed computing, often referred to as the LOCAL model, due to Linial [24]. The network is abstracted as an n -node undirected graph $G = (V, E)$, where each node represents one processor and a link between two nodes indicates that those two processors can communicate directly. Each processor has a unique b -bit identifier, where we typically assume $b = O(\log n)$. Initially, nodes do not know the topology of the network G , except for potentially knowing some global parameters such as a polynomial upper bound on n . Computations and communications take place in synchronous rounds. Per round, after doing arbitrary computations on the data that it holds, each process/node can send one message to each of its neighbors. In the LOCAL model, the message sizes are not bounded. The model variant where message sizes are limited to $O(\log n)$ -bits is known as the CONGEST model[31]. The messages sent in a round are delivered before the end of that round. At the end of the computation, each node should know its own output, e.g., in the MIS problem, each node should know whether it is in the computed maximal independent set or not.

1.1 State of the Art

Randomized algorithms, and the pursuit of deterministic algorithms. In the 1980s, Luby[27] and Alon, Babai, and Itai[1] presented a simple and elegant randomized distributed algorithm that computes an MIS in $O(\log n)$ rounds, with high probability¹. Due to known reductions, this MIS algorithm led to $O(\log n)$ round randomized algorithms for many other key graph problems, including maximal matching, $\Delta + 1$ vertex coloring, and $(2\Delta - 1)$ edge coloring. These problems are often listed as the four fundamental

¹As standard, the phrase *with high probability* indicates that an event happens with probability at least $1 - 1/n$.

symmetry-breaking problems in distributed graph algorithms and have a wide range of applications. The $O(\log n)$ -round randomized algorithm naturally led the researchers to seek a deterministic distributed algorithm with the same round complexity. In his celebrated work [24, 25], Linial asked “*can it [MIS] always be found in polylogarithmic time [deterministically]?*” He even added that “*getting a deterministic polylog-time algorithm for MIS seems hard.*” Since then, this became known as Linial’s MIS question and turned into one of the research foci in distributed graph algorithms.

The two approaches of deterministic algorithms. Linial’s MIS problem remained open for nearly three decades. During this time, two different general approaches were developed and pursued. The first approach relies on *global computations/derandomization via network decompositions*. The second approach is based on *local rounding of certain fractional solutions*. Over the past couple of years, both approaches came to fruition and led to two completely independent polylogarithmic time deterministic distributed algorithms for MIS [13, 32]. The round complexities are still $\tilde{O}(\log^3 n)$, which is somewhat far from the randomized $O(\log n)$ complexity [1, 27]. We next discuss each approach separately.

(I) *Global computation/derandomization via network decomposition.* The first approach is based on the concept of network decomposition, which was introduced by Awerbuch, Goldberg, Luby, and Plotkin [4] as the key tool in developing deterministic distributed algorithms for MIS and other symmetry-breaking problems. A (c, d) network decomposition is a partition of the vertex set into c disjoint parts, each known as a *block*, such that in the subgraph induced by the nodes in each block, each connected component has a diameter of at most d . Said differently, each block consists of non-adjacent clusters, each of diameter² at most d . Given a (c, d) network decomposition, it is easy to compute an MIS in $O(cd)$ rounds: The rounds are organized in c iterations, each consisting of $O(d)$ rounds. In iteration i , we add to the output independent set an MIS of the nodes of block i , which can be computed easily in $O(d)$ rounds as each cluster in the block has a diameter of at most d . We then remove any node in any other block that has a neighbor in this independent set and then move to the next iteration. Any n -node graph has a (c, d) network decomposition for $c = d = O(\log n)$ [2]. Awerbuch et al. [4] gave a deterministic distributed algorithm that computed a (c, d) decomposition in t rounds for $c = d = t = 2^{O(\sqrt{\log n \log \log n})}$. These bounds were improved to $2^{O(\sqrt{\log n})}$ by Panconesi and Srinivasan [30]. By a technique of Awerbuch, Berger, Cowen, and Peleg [3], one can transform any t -round (c, d) decomposition algorithm to a $(t + cd) \cdot \text{poly}(\log n)$ -round $(\log n, \log n)$ network decomposition algorithm in the LOCAL model.

The $2^{O(\sqrt{\log n})}$ bounds of Panconesi and Srinivasan remained the state of the art for over 25 years, until getting improved dramatically to $\text{poly}(\log n)$ by a new decomposition algorithm of Rozhon and Ghaffari [32]. That gave the first polylogarithmic-time deterministic distributed algorithm for MIS, and thus the first positive resolution to Linial’s MIS question. The algorithm was optimized in a follow-up work of Ghaffari, Grunau, and Rozhon [18], which

brought down MIS’s deterministic round complexity to $O(\log^5 n)$. See also [9, 11] which obtain strong-diameter guarantees with small messages. In a very recent work, Ghaffari, Grunau, Haeupler, Ilchi, and Rozhon [17] presented a completely different and faster method for computing network decompositions. Their algorithm computes a (c, d) strong-diameter network decomposition for $c = O(\log n)$ and $d = O(\log n \cdot \log \log \log n)$ in $\tilde{O}(\log^3 n)$ rounds, using $O(\log n)$ -bit messages. This is the state of art decomposition in essentially all regards, and it provides an $\tilde{O}(\log^3 n)$ round deterministic algorithm for MIS.

(II) *Local rounding of fractional solutions.* The second approach is based on obtaining fractional solutions to certain relaxations of the problem and then, locally and gradually rounding these solutions into integral solutions. Unlike the network decomposition approach, which was obviously applicable to all the symmetry-breaking problems from the start, the applicability of the rounding approach appeared limited at first and gradually increased. It started first with only the maximal matching problem, then extended to $2\Delta - 1$ edge coloring, then to $\Delta + 1$ vertex coloring, and finally to the hardest of all, the MIS problem.

The starting point is the work of Hanckowiak, Karonski, and Panconesi [22] who gave the first poly log n -time deterministic distributed algorithm for the maximal matching problem. Fischer [14] rephrased their approach in a *rounding* language and used this to improve the maximal matching complexity to $O(\log^2 \Delta \cdot \log n)$.

Fischer’s rounding was very specific to matching in graphs. Fischer, Ghaffari, and Kuhn [15] developed a different rounding method for matchings that extended to low-rank hypergraphs. By a reduction that they provided from $(2\Delta - 1)$ -edge coloring in graphs to maximal matching in hypergraphs of rank 3, this led to a poly log n -time deterministic algorithm for $(2\Delta - 1)$ -edge coloring, hence putting the second problem in the poly($\log n$) regime. Harris [23] improved the complexity to $\tilde{O}(\log^2 \Delta \cdot \log n)$.

The above local rounding approaches appeared limited to matching in graphs or hypergraphs, until a work of Ghaffari and Kuhn [20]. They developed an efficient rounding method for $\Delta + 1$ coloring (this was shortly after the first polylogarithmic-time network decomposition result [32]). This was by reexamining the analysis of the natural randomized coloring algorithm, seeing it as a fractional/probabilistic assignment of colors to the nodes which has a small bound on the expected number of monochromatic edges, and then gradually rounding the fractional assignments while approximately maintaining the upper bound on the expected number of monochromatic edges.

Finally, in a very recent work, Faour, Ghaffari, Grunau, Kuhn, and Rozhon [13] found a significant generalization of the above local roundings. They presented a unified method that can provide local rounding for any problem whose randomized solution analysis relies on only pairwise independence. This led to algorithms with round complexity $O(\log^2 \Delta \log n)$ for MIS and thus also for the other three problems mentioned above, and in a unified way.

1.2 Our Results

Unifying the two approaches and faster deterministic MIS. As discussed above, prior to the present paper, the best known deterministic complexity of MIS (and indeed any of the other three key

²If the diameter is measured in the induced subgraph, this is known as a strong-diameter network decomposition. If the distance is measured in the original graph, this is known as a weak-diameter network decomposition.

symmetry breaking problems) remained $\Omega(\log^3 n)$ [13, 17]. This is somewhat far from the randomized $O(\log n)$ round complexity [1, 27]. Furthermore, the two approaches seemed unrelated and incompatible. In this paper, we present a method to combine the two general approaches, which allows us to achieve a deterministic round complexity of $\tilde{O}(\log^2 n)$ for MIS.

THEOREM 1.1. *There is a deterministic distributed algorithm, in the LOCAL model, that computes a maximal independent set (MIS) in $\tilde{O}(\log^2 n)$ rounds. This also implies that there are LOCAL-model deterministic distributed algorithms with $\tilde{O}(\log^2 n)$ rounds complexity also for maximal matching, $(deg + 1)$ -list vertex coloring, and $(2deg - 1)$ -list edge coloring³.*

This comes closer to the $\tilde{\Omega}(\log n)$ lower bound, due to Balliu et al. [5], which holds for deterministic distributed algorithms that compute maximal independent set or maximal matching.

Faster approximate maximum matching. Our method also leads to a faster deterministic algorithm for computing a constant approximation of the maximum matching, in $\tilde{O}(\log^{4/3} n)$ rounds. By adding to this the approximation booster of Fischer, Mitrovic, and Uitto [16], we can improve the approximation to $(1+\epsilon)$ for any desirably small constant $\epsilon > 0$ with no asymptotic round complexity loss. The previous best round complexity for constant approximation of the maximum matching was $O(\log^2 \Delta)$ [13, 14].

THEOREM 1.2. *There is a deterministic distributed algorithm, in the LOCAL model, that computes a $\Theta(1)$ -approximate maximum matching in $\tilde{O}(\log^{4/3} n)$ rounds. This can be boosted to an algorithm that computes a $(1 + \epsilon)$ approximation of maximum matching in $\tilde{O}(\log^{4/3} n) \cdot \text{poly}(1/\epsilon)$ rounds.*

Due to space limitations, the proof of [Theorem 1.2](#) is deferred to the full version of this paper.

Corollaries for randomized coloring. Our faster deterministic algorithm can be plugged in into the shattering framework of randomized algorithms [7, 10] and improves the randomized complexity for $\Delta + 1$ vertex coloring and $2\Delta - 1$ edge coloring.

Corollary 1.3 ([Theorem 1.1](#)+[\[10\]](#)). *There is a randomized distributed algorithm, in the LOCAL model, that computes a $\Delta + 1$ vertex coloring in $\tilde{O}((\log \log n)^2)$ rounds. The same holds also for the $2\Delta - 1$ edge coloring problem.*

1.3 Other Related Work

We discussed in [Section 1.1](#) deterministic constructions for network decomposition and the generic method of using network decomposition for symmetry-breaking problems, e.g., MIS. Let us add here three side comments and mention the related work.

(1) Better randomized constructions for network decomposition have been known. In particular, the work of Linial and Saks [26] presented an $O(\log^2 n)$ round randomized algorithm for computing decompositions with $O(\log n)$ colors and $O(\log n)$ weak diameter.

³In the $(deg + 1)$ -list vertex coloring, each node v having a prescribed list L_v of colors, of size $|L_v| \geq deg(v) + 1$ from which it should choose its color. In the $(2deg - 1)$ -list edge coloring, each edge $e = \{v, u\}$ has a prescribed list L_e of colors, of size $|L_e| \geq deg(v) + deg(u) - 1$ from which it should choose its color. Both problems reduce to MIS by a reduction of Luby [25, 27].

Elkin and Neiman [12] imported a parallel algorithm of Miller, Peng and Xu [29] into the distributed setting and obtained an $O(\log^2 n)$ round randomized algorithm for computing decompositions with $O(\log n)$ colors and $O(\log n)$ strong diameter.

(2) The deterministic MIS method described in [Section 1.1](#) for using network decompositions in computing symmetry-breaking problems such as MIS would require large messages, as it gathers the topology of each cluster in a center. For MIS, one can work with $O(\log n)$ -bit messages, by using a derandomization method of Censor-Hillel, Parter, and Schwartzman [8], and that gives a deterministic algorithm with $O(\log n)$ -bit messages and round complexity $O(cd) \cdot \text{poly}(\log n)$. Similar approaches have been presented for other problems, see e.g., Bamberger et al. [6] for results on $\Delta + 1$ coloring.

(3) Ghaffari, Harris, and Kuhn [19] and Ghaffari, Kuhn and Maus [21] showed that one can get a general derandomization method for the LOCAL model using network decompositions. This method transforms any $\text{poly}(\log n)$ -round randomized algorithm for any problem whose solution can be checked in $\text{poly}(\log n)$ rounds into a deterministic algorithm with round complexity $O(cd + t) \cdot \text{poly}(\log n)$, assuming we have a deterministic (c, d) network decomposition algorithm with round complexity t .

2 OUR APPROACH

To discuss our approach and put it in the context of prior work, let us use the problem of maximal matching. This is a special case of the maximal independent set problem (on line graphs), but it will suffice for explaining most of the key ideas in our overall approach.

2.1 The Previous Methods and Their Barriers

As mentioned before, the previous two approaches for deterministic algorithms both get stuck at a round complexity of $\tilde{O}(\log^3 n)$. Let us briefly overview the approaches and their barriers here. This discussion will help us explain how we obtain a faster algorithm.

Local rounding of fractional solutions. For maximal matching, Fischer [14] and Faour et al. [13] provide two different local rounding algorithms, both achieving an $O(\log^2 \Delta \log n)$ complexity. Both round complexities are stuck at essentially the same barrier, so we discuss only Fischer's approach. It is easy to deterministically obtain a constant-approximate fractional matching in $O(\log \Delta)$ rounds⁴ [14]. The real challenge for deterministic algorithms is in rounding this fractional matching into an integral matching. Fischer [14] gradually improves the integrality of the fractional solution from $1/\Delta$ to 1, in $\log \Delta$ gradual rounding steps, each time doubling the minimum edge value. Each doubling step in the algorithm takes $\Theta(1/\epsilon) = \Theta(\log \Delta)$ rounds, as we need to ensure that there is a loss of at most $\epsilon = 1/(2 \log \Delta)$ fraction in the matching size, so that the total loss over all the $\log \Delta$ iterations is at most a constant. One can see that, for an arbitrary fractional matching, this $\Theta(1/\epsilon)$ complexity is the best possible. Considering the $\log \Delta$ gradual rounding steps, this yields a constant-approximate integral matching in $O(\log^2 \Delta)$ rounds. We emphasize that this runtime

⁴Start with a fractional assignment of $x_e = 1/\Delta$ to all edges. Then, we have $\log \Delta$ iterations. Per iteration, for each edge e such that each endpoint v of it has $\sum_{e' \ni v} x_{e'} \leq 1/2$, set $x_e \leftarrow 2x_e$.

is quadratic in the number of rounding steps, as we need to ensure that the relative loss per step is less than the inverse of the number of steps. Via $O(\log n)$ repetitions of constant-approximate maximum matching, each time on the subgraph induced by vertices that remained unmatched, one gets a maximal matching with complexity $O(\log^2 \Delta \cdot \log n)$.

Global computation via network decomposition. The state-of-the-art network decomposition algorithm [17] computes a decomposition with $O(\log n)$ vertex-disjoint blocks, where the subgraph induced by each block consists of connected components of diameter $\tilde{O}(\log n)$, in $\tilde{O}(\log^3 n)$ rounds. Given this, it is easy to compute maximal matching or MIS in $\tilde{O}(\log^2 n)$ rounds. But the construction time is the bottleneck in these applications.

Let us revisit this decomposition-based approach for the constant-approximate maximum matching problem. This discussion will help us later to explain our approach. As discussed above, $O(\log n)$ repetitions of constant-approximate maximum matching give a maximal matching. Consider a constant-approximate *fractional* maximum matching. As mentioned above, we know how to compute this easily in $O(\log \Delta)$ rounds. The task is to turn this into an integral matching of the same size, up to a constant factor. The network decomposition algorithm [17] computes the blocks one by one, each in $\tilde{O}(\log^2 n)$ rounds, such that each block clusters a constant fraction of the remaining nodes into non-adjacent clusters, each of diameter $\tilde{O}(\log n)$. We can adjust the block construction in such a way that, instead of a constant fraction of nodes getting clustered, we ensure that at least a constant fraction of the weight of the fractional matching is on the edges inside the clusters. It is easy to see that randomized rounding can turn (each cluster's) fractional matching into an integral matching of the same size up to a constant factor. So, to deterministically do the rounding, we can simply gather each cluster's information (and constraints) into the cluster center and compute the integral matching. This takes $\tilde{O}(\log n)$ rounds, as the cluster diameter is $\tilde{O}(\log n)$. It is worth noting that the current method for building each block of the decomposition requires $\tilde{O}(\log^2 n)$ rounds, as it has $\tilde{O}(\log n)$ steps of randomness fixing, each involving coordination along $\tilde{O}(\log n)$ -hop distances.

2.2 Low-diameter Low-degree Clusterings

In our approach, we work with a relaxed variant of the standard network decompositions. In particular, we work with clusterings that have a lower diameter than the standard $\Theta(\log n)$ bound. This opens the road for faster deterministic constructions. However, it comes at the disadvantage that we do not have only $\log n$ blocks, each consisting of non-adjacent clusters. We instead maintain a small bound on the maximum number of clusters that are adjacent to each particular node. To state the formal decomposition, let us first present the formal definitions.

Definition 2.1. [Cluster, Clustering, Partition, Cluster Degree of a Vertex, Diameter of a Cluster] A subset $C \subseteq V$ is called a cluster. The (strong)-diameter of a cluster is defined as $\max_{u,v \in C} d_{G[C]}(u,v)$. A clustering \mathcal{C} is a set of disjoint clusters. We refer to \mathcal{C} as a partition if $\bigcup_{C \in \mathcal{C}} C = V$. The diameter of a clustering \mathcal{C} is defined as the maximum diameter of all its clusters. For a node $u \in V$ and a clustering \mathcal{C} , we define the degree of u with respect to \mathcal{C}

as $\deg_{\mathcal{C}}(u) = |\{C \in \mathcal{C} : d(C,u) \leq 1\}|$. We sometimes refer to $\max_{u \in V} \deg_{\mathcal{C}}(u)$ as the degree of this clustering.

We show an algorithm that computes $O(\alpha)$ -diameter ($2^{\tilde{O}(\log n)/\alpha}$)-degree clusterings in $\tilde{O}(\alpha^2 \log n)$ rounds of the LOCAL model.

THEOREM 2.1. Let $\alpha \leq \log n$ be an arbitrary value. There exists a deterministic distributed algorithm that in $\tilde{O}(\alpha^2 \log n)$ rounds of the LOCAL model computes a partition \mathcal{C} with

- (1) diameter $O(\alpha)$ and
- (2) $\max_{u \in V} \deg_{\mathcal{C}}(u) = 2^{\tilde{O}(\log n)/\alpha}$, i.e., each node has neighbors in at most $2^{\tilde{O}(\log n)/\alpha}$ clusters.

We note that the existence of such a clustering/partition, and even efficient randomized constructions for it, follow from known randomized methods [29]. We provide an efficient deterministic distributed construction here, as we outline in Section 2.4. Furthermore, we show that this relaxed notion is still useful for applications in the maximal independent set and maximal matching problems, as we next outline in Section 2.3.

As a side comment, we note that for some applications (e.g., approximate-matching), it suffices to have a weaker version of Theorem 2.1 where only a constant fraction of nodes (or a constant fraction of the weight of the nodes, according to some given weights) is clustered. For that version, we have a faster algorithm that runs in $\tilde{O}(\alpha \log n)$ rounds of the LOCAL model. The precise statement is as follows.

THEOREM 2.2. Let $\alpha \leq \log n$ be an arbitrary value and assume that each node $u \in V$ is equipped with a value $x_u \in [\frac{1}{n}, 1]$. There exists a deterministic distributed algorithm that in $\tilde{O}(\alpha \log n)$ rounds of the LOCAL model computes a partition \mathcal{C} with

- (1) diameter $O(\alpha)$ and
- (2) let $V^{good} = \{u \in V : \deg_{\mathcal{C}}(u) = 2^{\tilde{O}(\log n)/\alpha}\}$, then we have $\sum_{u \in V^{good}} x_u \geq 0.9 \sum_{u \in V} x_u$.

2.3 Interplay Between Global Computations via Clustering and Local Rounding

Let us now revisit the maximal matching problem. Here, we provide an intuition of how we solve this problem in $\tilde{O}(\log^2 n)$ rounds, by combining the two approaches. First, we invoke Theorem 2.1 for $\alpha = \sqrt{\log n}$ and compute a partition \mathcal{C} into clusters of diameter $O(\sqrt{\log n})$ such that each node has $d = 2^{\tilde{O}(\sqrt{\log n})}$ neighboring clusters, in $O(\log^2 n)$ rounds. This vertex partition also induces an edge partition $E = \sqcup_{C \in \mathcal{C}} E_C$ where E_C contains all edges whose endpoint with the larger identifier is in C . Now, we use this clustering to gradually compute the maximal matching in $O(\log n)$ iterations. In each iteration, we first compute a constant-approximate fractional matching in the graph induced by the unmatched vertices in $O(\log \Delta) = O(\log n)$ rounds [14]. Then, the key part of the iteration is to turn this into a constant-approximate integral matching in $\tilde{O}(\log n)$ additional rounds. For that, we use the clustering to quickly perform a significant rounding of the fractional matching. In particular, the value assigned to a given edge $e \in E_C$ after the rounding is fully determined by the fractional values assigned to the edges in E_C before the rounding. Thus, this initial rounding can

be computed in just $O(\sqrt{\log n})$ rounds. Afterward, we will resort to local rounding which now needs only $\tilde{O}(\sqrt{\log n})$ doubling steps and thus can be performed in $\tilde{O}(\log n)$ rounds. Let us elaborate.

First, we divide all the fractional matching values by a 4 factor, so that now each node v has $\sum_{e' \ni v} x_{e'} \leq 1/4$. Then, for the rounding, for each node v and cluster C , we allow the summation of the edges in E_C incident to v to increase from the original value by a multiplicative factor of 2 and, on top of that, an additive increase of $1/(2d) = 1/2^{\tilde{O}(\sqrt{\log n})}$. Hence, the multiplicative deviations bring the summation to at most $1/2$ and the additive deviations, over all the d neighboring clusters, raise the total sum to at most 1. Thus, this would still ensure that we have a valid fractional matching. How does each cluster compute such a rounded fractional solution that satisfies this constraint for all possible neighboring nodes v ?

For each cluster C , such a solution exists by a simple probabilistic rounding argument. Let us provide an informal explanation. Suppose we keep the fractional values that exceed $1/\Theta(d \log n)$ intact. For any edge that has $x_e \leq 1/\Theta(d \log n)$, let us round it probabilistically: with probability $p = x_e \cdot \Theta(d \log n)$, we set it to $1/\Theta(d \log n)$, and with the rest of the probability, we set it to zero. One can see by a standard Chernoff bound that, w.h.p., for the edges in E_C incident on node v , the new summation of fractional values is within a 2 factor of the old summation, modulo an additive error of at most $1/(2d)$. With this rounding, now the smallest fractional value is at least $1/\Theta(d \log n) = 1/2^{\tilde{O}(\sqrt{\log n})}$. Furthermore, with a reasonably high probability, the size of the fractional matching is preserved up to a constant factor inside the cluster⁵. To deterministically find the solution proven to exist by this probabilistic argument, it suffices for the cluster C to gather the current fractional values of the edges in E_C . Then, it can find such a fractional solution in a centralized fashion. Since the diameter of the cluster is $O(\sqrt{\log n})$, this can be done easily in $O(\log n)$ rounds of the LOCAL model.

Once each cluster does this rounding (and note that all can be done in parallel), we have computed a fractional matching with integrality $1/2^{\tilde{O}(\sqrt{\log n})}$ whose size is at least a constant fraction of the fractional matching with which we started. Hence, it suffices to apply only $\tilde{O}(\sqrt{\log n})$ steps of local rounding of [14] on this fractional matching, and that takes $\tilde{O}(\sqrt{\log n})^2 = \tilde{O}(\log n)$ rounds. Therefore, we now have a method to compute a constant-approximate integral matching in $\tilde{O}(\log n)$ rounds. This is after the initial $\tilde{O}(\log^2 n)$ time spent for computing the clustering, but that clustering is computed only once. With $O(\log n)$ repetitions of this constant-approximate integral matching in $\tilde{O}(\log n)$ rounds, each time removing the matched nodes from the future matching computations, we get a maximal matching in $\tilde{O}(\log^2 n)$ rounds.

2.4 Construction of the Clustering, and Hitting Set with Pipelining

Our low-diameter clustering algorithms are similar in spirit to Ghaffari et al. [17]. We start by briefly overviewing their approach.

Review of the clustering algorithm of [17]. Recently, Ghaffari et al. [17] obtained a network decomposition with $O(\log n)$ colors and diameter $O(\log n \log \log \log n)$. One key step towards that result

⁵We defer the precise discussions of the details to Section 5.

is an algorithm that, in $\tilde{O}(\log^2 n)$ rounds, computes a partition of diameter $O(\log n)$ such that a large constant fraction of the vertices have a clustering degree of $O(\log \log n)$. Their algorithm to compute such a partition can be seen as a derandomization of the randomized algorithm of Miller, Peng and Xu [29]. The randomized algorithm of Miller, Peng and Xu computes a partition of diameter $O(\log n)$ by first assigning each vertex a (random) delay between 0 and $O(\log n)$. Then, each node u gets clustered to the node v minimizing $del(v) + d(v, u)$ where $del(v)$ is the delay assigned to node v . The delays assigned to all of the vertices can be computed by repeatedly, for $\log n$ repetitions, subsampling all the remaining active nodes with probability $1/2$. Here, all the nodes would be active at the beginning, and in each repetition we keep each previously active node with probability $1/2$. With high probability, no active node remains after $O(\log n)$ subsampling steps. The deterministic algorithm of Ghaffari et al. derandomizes each of the $O(\log n)$ subsampling steps. To do so, they phrase each subsampling step as an instance of a certain hitting set variant. They show that the randomized algorithm produces a partition such that the expected cluster degree of each node is $O(\log \log n)$. The hitting set viewpoint in each subsampling step lends itself to pairwise analyses and opens the road for efficient derandomization, allowing one to “sample” the vertices of each step in $\tilde{O}(\log n)$ rounds in a deterministic manner.

Our Clustering Algorithms. For our weaker clustering result, namely Theorem 2.2 which clusters only a constant fraction of the vertices (or their weights), we follow a similar general approach. One can see that, for any $\alpha \leq \log n$, by a simple parameter adjustment in [29], their randomized construction produces a partition with diameter $O(\alpha)$ and cluster degree $2^{O(\log n/\alpha)}$ – essentially, it suffices to have α subsampling steps, each with probability $2^{-O(\log n/\alpha)}$. We can then follow a similar derandomization approach as [17] to turn this into a deterministic algorithm. Indeed, because of the parameter regime, we can work here with a even slightly simpler hitting set analyses, as we describe in Section 4.1. This leads to clustering a constant fraction of the nodes (or a constant fraction of their weights) in $\tilde{O}(\alpha \log n)$ rounds. Intuitively, in contrast to the $\tilde{O}(\log^2 n)$ complexity of the clustering of [17], the complexity here is $\tilde{O}(\alpha \log n)$ because the distance in the hitting set problem (to coordinate between the nodes to be hit and the active nodes) is reduced from $O(\log n)$ to $O(\alpha)$.

For the stronger clustering algorithm that we have, namely Theorem 2.1, we want to cluster all vertices. Doing this by repetitions of Theorem 2.2, which clusters only a constant fraction of nodes (or their weights), would incur a factor of $O(\log n)$ loss in the round complexity. That would bring the complexity of our clustering to $\tilde{O}(\alpha \log^2 n)$. Theorem 2.1 however achieves a complexity of $\tilde{O}(\alpha^2 \log n)$. For that, we present a novel pipelining idea in the hitting set framework, as we outline next.

Efficient Hitting Set Derandomization with Pipelining. Recently, Faour et al. [13] gave a local rounding method that essentially can derandomize pairwise analyses. Consider some random process where each node gets sampled with probability p , pairwise independently. Then, assuming the pairwise analyses looks at only pairs of nodes that are nearby, their method allows one to derandomize the sampling in roughly $O(\log^2(1/p))$ rounds of the LOCAL model.

Let us provide a brief intuition: their algorithm views the probabilistic solutions as a p -integral fractional assignment, and gradually rounds it to an integral solution in $s = O(\log(1/p))$ doubling steps. The approach views the objectives analyzed via pairwise analyses as its target function, which should be approximately preserved. Each step uses a certain defective coloring and then decides the rounding of the nodes of each color simultaneously. This ensures that, with the exception of the small loss in the target function over edges that are colored monochromatically, the rest of the target function is preserved. The coloring is chosen such that the loss in each step is roughly a $1/(2s)$ fraction of the overall target function, so that the total loss over all $s = O(\log(1/p))$ steps is still below a constant fraction. This aspect of the parameterization is similar to the basic rounding of Fischer [14] for matchings, as discussed before.

Now, let us illustrate our pipelining idea by considering a simplified setup. Let $G = (U \sqcup V, E)$ be a bipartite graph. We assume that each node $u \in U$ has degree Δ and we define $n = |U|$. The goal is to compute a small subset $V' \subseteq V$ such that each node in U is neighboring with at least one node in V' . Such a set V' is referred to as a hitting set.

Consider the probabilistic approach where each node in V joins V' independently with probability $p = O(\log(n)/\Delta)$. The expected size of V' is $O\left(\frac{\log(n)|V|}{\Delta}\right)$. Moreover, each node $u \in U$ is hit with probability $1 - (1-p)^\Delta \geq 1 - \frac{1}{n^{10}}$. Unfortunately, this analysis completely breaks down if we only assume pairwise independence.

A hitting set with the same guarantees can be computed using only pairwise independence as follows: First, let each node in V join V' with probability $\frac{1}{10\Delta}$, pairwise independently. A simple pairwise analysis shows that each node has one neighbor in V' with probability at least $\frac{1}{100}$. Therefore, repeating this subsampling step $O(\log n)$ times, and adding every node to V' that was sampled at least once, results in each node being hit with high probability. Moreover, the expected size of V' is $O\left(\frac{\log(n)|V|}{\Delta}\right)$.

The method of Faour et al. [13] allows to directly derandomize the pairwise analysis. In particular, it allows to compute in $O(\log^2 \Delta)$ rounds a subset $V' \subseteq V$ of size $O(|V|/\Delta)$ such that at least $\frac{n}{100}$ vertices of U have a neighbor in V' . Hence, in $O(\log n \log^2 \Delta)$ rounds, one can deterministically compute a hitting set of size $O(\log(n)|V|/\Delta)$. This is not efficient enough for us to achieve [Theorem 2.2](#), especially for large $\Delta \gg \log n$.

We give a method that, in just $\tilde{O}(\log n)$ rounds, deterministically computes a slightly larger hitting set of size $O(|V|/\Delta^{1/\log \log n})$. Let $k = O(\log \Delta / \log \log n)$ be the largest integer such that $(100 \log n)^k \leq \Delta$. For $j \in [0, k]$, let $\deg_j = (100 \log n)^{k-j}$. The algorithm computes a sequence of sets $V := V_0 \supseteq V_1 \supseteq \dots \supseteq V_k$ such that each node in U has at least \deg_j neighbors in V_j . In particular, each node in U has at least $\deg_k \geq 1$ neighbors in V_k and therefore V_k is a hitting set. The algorithm also ensures that $|V_j| \leq |V|/2^j$ and therefore $|V_k| = O(|V|/\Delta^{1/\log \log n})$, as promised. In the randomized world, one could obtain V_{j+1} from V_j by subsampling each vertex with probability $1/2$. Then, given that each node in U neighbors at least \deg_j nodes in V_j , one can show that with high probability each node in U neighbors at least \deg_{j+1} nodes in V_{j+1} . One can also derandomize this with round complexity $\tilde{O}(\log n)$, which then would

result in a round complexity of $\tilde{O}(\log n \log \Delta)$. We give a method that pipelines the computation of the sets, in the sense that it starts the computation of V_{j+1} before we have finished computing V_j . The actual ingredients of this pipelining are more involved and are thus deferred to the technical section. We hope that similar pipelining ideas might find applications in other similar contexts.

3 PRELIMINARIES

3.1 Definitions and Basic Inequalities

Notation and Basic Definitions. We use n to denote the number of nodes and N to denote a polynomial upper bound on the number of nodes. For two integers a, b where $a \leq b$, we define $[a, b] = \{a, a+1, \dots, b\}$. For an event \mathcal{E} , we define the indicator variable $I(\mathcal{E})$ to be equal to 1 if \mathcal{E} happens and 0 otherwise.

Given a graph $G = (V, E)$, a fractional matching is an assignment of a value $x_e \in [0, 1]$ to each $e \in E$ such that for every vertex $v \in V$, we have $\sum_{e \ni v} x_e \leq 1$. The size of this fractional matching is $\sum_e x_e$. We call the fractional matching $1/Q$ -integral if each edge e satisfies either $x_e = 0$ or $x_e \geq 1/Q$. We use this $1/Q$ -integrality term more generally also for other fractional assignments.

We next state some concentration inequalities that we will use throughout the paper.

THEOREM 3.1 (CHERNOFF BOUND). *Let $X := \sum_{i \in [1, n]} X_i$, where $X_i, i \in [1, n]$ are independently distributed and $0 \leq X_i \leq 1$. Then, for a given $\delta > 0$, we have*

$$\Pr[|X - \mathbb{E}[X]| \geq \delta \mathbb{E}[X]] \leq 2e^{-\frac{\min(\delta, \delta^2) \mathbb{E}[X]}{3}}.$$

Corollary 3.2 (Chernoff Bound variant). *Let $Y := \sum_{i \in [1, n]} Y_i$, where $Y_i, i \in [1, n]$ are independently distributed and $0 \leq Y_i \leq b$ for some $b > 0$. Then, for a given $t \geq 0.5E[Y]$, it holds that*

$$\Pr[|Y - \mathbb{E}[Y]| \geq t] \leq 2e^{-\frac{t}{6b}}.$$

3.2 The Local Rounding Framework

We use the local rounding framework of Faour et al. [13]. Their rounding framework works via computing a particular weighted defective coloring of the vertices, which allows the vertices of the same color to round their values simultaneously, with a limited loss in some objective functions that can be written as a summation of functions, each of which depends on only two nearby nodes. Next, we provide a related definition and then state their black-box local rounding lemma.

Definition 3.1. [Pairwise Utility and Cost Functions] *Let $G = (V_G, E_G)$ be a graph. For any label assignment $\vec{x} : V_G \rightarrow \Sigma$, a pairwise utility function is defined as $\sum_{u \in V_G} \mathbf{u}(u, \vec{x}) + \sum_{e \in E_G} \mathbf{u}(e, \vec{x})$, where for a vertex u , the function $\mathbf{u}(u, \vec{x})$ is an arbitrary function that depends only on the label of u , and for each edge $e = \{u, v\}$, the function $\mathbf{u}(e, \vec{x})$ is an arbitrary function that depends only on the labels of v and u . These functions can be different for different vertices u and also for different edges e . A pairwise cost function is defined similarly.*

For a probabilistic/fractional assignment of labels to vertices V_G , where vertex v assumes each label in Σ with a given probability, the utility and costs are defined as the expected values of the utility and cost functions, if we randomly draw integral labels for the vertices

from their corresponding distributions (and independently, though of course each term in the summation depends only on the labels of two vertices and thus pairwise independence suffices).

Lemma 3.3. [cf. Lemma 2.5 of [13]] Let $G = (V_G, E_G)$ be a multi-graph, which is equipped with utility and cost functions \mathbf{u} and \mathbf{c} and with a fractional label assignment λ such that for every label $\alpha \in \Sigma$ and every $v \in V_G$, $\lambda_\alpha(v) \geq \lambda_{\min}$ for some given value $\lambda_{\min} \in (0, 1]$. Further assume that G is equipped with a proper ζ -vertex coloring. If $\mathbf{u}(\lambda) - \mathbf{c}(\lambda) \geq 0.1\mathbf{u}(\lambda)$ and if each node knows the utility and cost functions of its incident edges, there is a deterministic $O(\log^2(\frac{1}{\lambda_{\min}}) + \log(\frac{1}{\lambda_{\min}}) \cdot \log^* \zeta)$ -round distributed message passing algorithm on G , in the LOCAL model, that computes an integral label assignment ℓ for which

$$\mathbf{u}(\ell) - \mathbf{c}(\ell) \geq 0.9(\mathbf{u}(\lambda) - \mathbf{c}(\lambda)).$$

3.3 The Hitting Set Subroutines

We also make use of a deterministic distributed algorithm for computing a certain hitting set. The formal statement is provided below.

Lemma 3.4. There exists a deterministic distributed algorithm in the LOCAL model such that, for every $\Delta \in \mathbb{N}$, $b \in \mathbb{N}$, $k \in [1, \Delta]$, $p = \Omega(1/k)$ and $\text{Norm} \geq 0$, it provides the following guarantees: The input is a bipartite graph $H = (U_H \sqcup V_H, E_H)$ with $\deg_H(u) = \Delta$ for every $u \in U_H$. Initially, each node is equipped with a unique b -bit identifier and each node u is assigned a weight $w_u \geq 0$. Each node also knows at the beginning to which side of the bipartition it belongs. The algorithm computes a subset $V^{sub} \subseteq V_H$ satisfying

$$\begin{aligned} & \sum_{u \in U_H: |N_H(u) \cap V^{sub}| \leq 0.5 \lfloor \Delta/k \rfloor} w_u + \text{Norm} \cdot |V^{sub}| \\ & \leq 4 \left(e^{-pk} \sum_{u \in U_H} w_u + \text{Norm} \cdot p \cdot |V_H| \right) \end{aligned}$$

The algorithm runs in $O(kp(\log^2(k) + \log(k) \log^* b))$ rounds.

Notice that a simple probabilistic scheme where we include each node of V_H in V^{sub} independently with probability p achieves the desired properties in expectation (and indeed with even stronger guarantees). In particular, let us even group the neighbors of each $u \in U_H$ into $\lfloor \Delta/k \rfloor$ groups each of size k —we will rely on independence only within each group and this simplifies the task for derandomization. The probability that a group does not have a node in V^{sub} is at most $(1-p)^k \leq e^{-pk}$. If $u \in U_H: |N_H(u) \cap V^{sub}| \leq 0.5 \lfloor \Delta/k \rfloor$, then half of its groups are not hit. However, we expect only e^{-pk} fraction of groups not to be hit (and the fraction can be weighted, by taking the w_u weights into account). So, roughly speaking, the weighted fraction of nodes u such that $u \in U_H: |N_H(u) \cap V^{sub}| \leq 0.5 \lfloor \Delta/k \rfloor$ is at most $2e^{-pk}$. The lemma statement has another 2 factor of slack beyond this bound, which simplifies the argument and suffices for the applications. We present a method to deterministically compute such a set V^{sub} . This is by an adaptation of the method of Faour et al. [13], and statements somewhat similar to this lemma were implicit in [13, 17]. The proof of Lemma 3.4 appears in the full version of this paper.

4 CLUSTERING

This section is devoted to proving our two clustering results, as restated below. Please see Definition 2.1 for the related definitions.

THEOREM 2.1. Let $\alpha \leq \log n$ be an arbitrary value. There exists a deterministic distributed algorithm that in $\tilde{O}(\alpha^2 \log n)$ rounds of the LOCAL model computes a partition \mathcal{C} with

- (1) diameter $O(\alpha)$ and
- (2) $\max_{u \in V} \deg_{\mathcal{C}}(u) = 2^{\tilde{O}(\log n)/\alpha}$, i.e., each node has neighbors in at most $2^{\tilde{O}(\log n)/\alpha}$ clusters.

THEOREM 2.2. Let $\alpha \leq \log n$ be an arbitrary value and assume that each node $u \in V$ is equipped with a value $x_u \in [\frac{1}{n}, 1]$. There exists a deterministic distributed algorithm that in $\tilde{O}(\alpha \log n)$ rounds of the LOCAL model computes a partition \mathcal{C} with

- (1) diameter $O(\alpha)$ and
- (2) let $V^{good} = \{u \in V: \deg_{\mathcal{C}}(u) = 2^{\tilde{O}(\log n)/\alpha}\}$, then we have $\sum_{u \in V^{good}} x_u \geq 0.9 \sum_{u \in V} x_u$.

Notice the first statement clusters all nodes and takes $\tilde{O}(\alpha^2 \log n)$ rounds, while the second statement clusters only a 0.9 fraction of the vertices, in a weighted sense, but runs faster in $\tilde{O}(\alpha \log n)$ rounds. Both algorithms can be seen as a derandomization of the randomized clustering algorithm of Miller, Peng and Xu [29]. The high-level approach in all of the three algorithms is the same and can be found in Algorithm 1. The idea is to first compute a delay $del(v) \in \{0, 1, \dots, 50\alpha\}$ for every $v \in V$. Then, these delays define a partition as follows: First, each node starts to broadcast a token at time $del(v)$. Then, in the following $O(\alpha)$ rounds, a node forwards all the token it has received in the previous rounds to its neighbors. For a given node u , let c_u be the node whose token reaches u first, breaking ties by choosing the node with the smallest identifier. Then, u joins the cluster corresponding to the node c_u . Formally speaking, we define $c_u := \arg \min_{v \in V} (del(v) + d(v, u), ID(v))$ where pairs are compared according to the lexicographic order and for a node v we define the cluster associated with v as $C_v := \{u \in V: c_u = v\}$. Finally, the output clustering consists of all non-empty clusters C_v for $v \in V$. To compute the delay $del(v)$ for every node $v \in V$, all three algorithms compute a sequence of sets $V_0^{active} := V \supseteq V_1^{active} \supseteq \dots \supseteq V_{10\alpha}^{active} = \emptyset$. Then, i_v is defined as the largest index i with $v \in V_i^{active}$ and the delay of v is defined as $del(v) := 50\alpha - 5i_v$. Our two clustering algorithms and the randomized clustering of Miller, Peng, and Xu only differ in the way V_{i+1}^{active} is computed given the set V_i^{active} ; the randomized clustering algorithm simply obtains V_{i+1}^{active} by including each vertex in V_i^{active} independently with probability $p_{rand} = 2^{-\log(n)/\alpha}$. Note that this implies $V_{10\alpha}^{active} = \emptyset$ with high probability. Furthermore, since $del(v) \in [0, 50\alpha]$, the diameter of the output partition is $O(\alpha)$.

Claim 4.1. Algorithm 1 produces a partition with diameter $O(\alpha)$.

The next lemma will be used later on to upper bound the cluster degree of a vertex. Before stating it, we introduce one more definition.

Definition 4.1 ($S_i(u)$). For every $u \in V$ and $i \in [0, 10\alpha - 1]$, we define

$$S_i(u) = V_i^{active} \cap B_G(u, d(u, \min(V_i^{active}) + 2, 100\alpha)).$$

Algorithm 1 Generic Partitioning AlgorithmInput: Parameter $\alpha \in \mathbb{N}$, Algorithm \mathcal{A} Output: A partition C of the vertex set

```

1: procedure PARTITIONING
2:                                      $\triangleright$  Delay Computation
3:    $V_0^{active} \leftarrow V$ 
4:   for  $i \leftarrow 0, 1, \dots, 10\alpha - 1$  do
5:      $V_{i+1}^{active} \leftarrow \mathcal{A}(V_i^{active}, i, \alpha)$     $\triangleright V_{i+1}^{active} \subseteq V_i^{active}$ 
6:   Assert :  $V_{10\alpha}^{active} = \emptyset$ 
7:   For every  $v \in V$ , let  $i_v$  be the largest index  $i$  with  $v \in V_i^{active}$ 
8:   For every  $v \in V$ , define  $del(v) \leftarrow 50\alpha - 5i_v$ 
9:                                      $\triangleright$  Partitioning given the delays
10:   $\forall u \in V$ :  $c_u \leftarrow \arg \min_{v \in V} (del(v) + d(v, u), ID(v))$ 
11:   $\forall v \in V$ :  $C_v \leftarrow \{u \in V : c_u = v\}$ 
12:  return  $C = \{C_v : v \in V, C_v \neq \emptyset\}$ 

```

Informally speaking, the set $S_i(u)$ contains those active nodes in V_i^{active} which are almost as close to u as the closest active node in V_i^{active} , unless the closest node in V_i^{active} is far away. The last part ensures that each node u can compute the set $S_i(u)$ in $O(\alpha)$ rounds, which will be important later on.

Lemma 4.2. *Let $u \in V$ be arbitrary, C the partition returned by Algorithm 1 and*

$$R_u := \max_{i \in \{0, 1, \dots, 10\alpha - 1\} : S_i(u) \cap V_{i+1}^{active} = \emptyset} |S_i(u)|.$$

Then, it holds that $\deg_C(u) = O(\alpha R_u)$.

Randomized Intuition. The proof of Lemma 4.2 is deferred to the full version. Here, we discuss some intuition by applying the lemma to the randomized clustering algorithm. In particular, a simple hitting set argument on top of Lemma 4.2 gives that $\deg_C(u) = O(\log n / p_{rand})$, w.h.p. Recall that the randomized algorithm obtains V_{i+1}^{active} by including each vertex in V_i^{active} with probability $p_{rand} = 2^{-\log(n)/\alpha}$. Consider an arbitrary $i \in \{0, 1, \dots, 10\alpha - 1\}$ and assume that $|S_i(u)| \geq 10 \log(n) / p_{rand}$. With high probability, at least one node in $S_i(u)$ will be sampled and added to V_{i+1}^{active} . Hence, with high probability $R_u \leq 10 \log(n) / p_{rand}$, in which case Lemma 4.2 gives $\deg_C(u) = O(\alpha \cdot \log(n) 2^{\log(n)/\alpha})$. We note that this analysis is not tight. Using the memoryless property of the exponential distribution, one can show that $\deg_C(u) = O(2^{\log(n)/\alpha})$ with positive constant probability and $\deg_C(u) = O(\log(n) 2^{\log(n)/\alpha})$ with high probability. The latter bound improves our bound by a factor of α , i.e., for $\alpha = O(\log n)$ our proof gives that with high probability $\deg_C(u) = O(\log^2 n)$ while a tighter analysis shows that $\deg_C(u) = O(\log n)$. However, this simple hitting set analysis is sufficient to obtain our desired result.

4.1 Clustering with Constant Fraction Good

This section is dedicated to proving Theorem 2.2. As a reminder, the goal is to compute a partition with diameter $O(\alpha)$ such that at least 90 percent of the vertices (in a weighted sense) have a cluster degree of $2^{\tilde{O}(\log n)/\alpha}$.

The algorithm follows the outline of Algorithm 1. The only missing part is to specify how to compute V_{i+1}^{active} given V_i^{active} for every $i \in [0, 10\alpha - 1]$.

In phase i , the goal is to compute V_{i+1}^{active} in such a way that almost all nodes u with $|S_i(u) \cap V_i^{active}| = 2^{\tilde{\Omega}(\log n)/\alpha}$ satisfy $S_i(u) \cap V_{i+1}^{active} \neq \emptyset$. We refer to a node u that does not satisfy this condition as being bad in phase i . For a node u , if there does not exist a phase i in which u is bad, Lemma 4.2 implies that the clustering degree of u is $2^{\tilde{O}(\log n)/\alpha}$. Below, we give the formal definition of being bad in phase i .

Definition 4.2 (Bad in phase i , U_i^{bad}). *For every $i \in [0, 10\alpha - 1]$, we refer to a node as bad in phase i if $|S_i(u)| \geq \lceil 1000 \log \log(N) \rceil^{\log(N)/\alpha}$ and $S_i(u) \cap V_{i+1}^{active} = \emptyset$. We denote by U_i^{bad} the set consisting of all bad nodes in phase i .*

Recall, in the LOCAL model nodes don't have access to the precise number of vertices in the graph, but are only given a polynomial upper bound, which we denote by N . Also note that $\lceil 1000 \log \log(N) \rceil^{\log(N)/\alpha} = 2^{\tilde{O}(\log n)/\alpha}$.

Starting with V_i^{active} , we compute V_{i+1}^{active} in $\log(N)/\alpha$ steps. In step j , we start with a set $V_{i,j}^{active}$ and compute a subset $V_{i,j+1}^{active} \subseteq V_{i,j}^{active}$. Initially, we define $V_{i,0}^{active} = V_i^{active}$ and after the last step, we set $V_{i+1}^{active} = V_{i, \log(N)/\alpha}^{active}$. Note that we can assume without loss of generality that N is a sufficiently large constant and that α divides $\log(N)$.

Before we explain in more detail what happens in each step, we first give two more definitions. The first definition associates a degree threshold for each step.

Definition 4.3 (\deg_j). *For every $j \in [0, \log(N)/\alpha]$, we define $\deg_j = \lceil 1000 \log \log(N) \rceil^{\log(N)/\alpha - j}$.*

If a node u is bad in phase i , then $|S_i(u) \cap V_i^{active}| \geq \deg_0$ and $|S_i(u) \cap V_{i+1}^{active}| < \deg_{\log(N)/\alpha}$.

Next, we define what it means for a node to be bad during a step.

Definition 4.4 (Bad in step j of phase i , $U_{i,j}^{bad}$). *For every $i \in [0, 10\alpha - 1]$ and $j \in [0, \log(N)/\alpha - 1]$ we say that a node is bad in step j of phase i if $|S_i(u) \cap V_{i,j}^{active}| \geq \deg_j$ and $|S_i(u) \cap V_{i,j+1}^{active}| < \deg_{j+1}$. We denote by $V_{i,j}^{bad}$ the set of bad nodes during step j of phase i .*

If a node u is bad in phase i , then $|S_i(u) \cap V_{i,0}^{active}| \geq \deg_0$ and $|S_i(u) \cap V_{i, \log(N)/\alpha}^{active}| < \deg_{\log(N)/\alpha}$. Hence, there exists some $j \in [0, \log(N)/\alpha - 1]$ such that $|S_i(u) \cap V_{i,j}^{active}| \geq \deg_j$ and $|S_i(u) \cap V_{i,j+1}^{active}| < \deg_{j+1}$ and therefore u is bad in at least one step of phase i .

Fix some $i \in [0, 10\alpha - 1]$ and $j \in [0, \log(N)/\alpha - 1]$. We compute $V_{i,j+1}^{active}$ given $V_{i,j}^{active}$ by first defining a bipartite graph H , then invoking the hitting set subroutine of Lemma 3.4 to deduce that there exists an efficient LOCAL algorithm on the bipartite graph which computes a subset of $V_{i,j}^{active}$ with the desired properties, and finally use the fact that we can simulate the algorithm on the bipartite graph on the original graph with an $O(\alpha)$ -multiplicative overhead.

Consider the bipartite graph $H = (U_H \sqcup V_H, E_H)$ with $U_H = \{u \in V : |S_i(u) \cap V_{i,j}^{active}| \geq \deg_j\}$, $V_H = V_{i,j}^{active}$ and where we connect each node $u \in U_H$ to \deg_j nodes in $S_i(u) \cap V_{i,j}^{active}$.

Using the fact that $S_i(u) \subseteq B_G(u, 1000\alpha)$ for every $u \in V$, it follows that each round in H can be simulated in $O(\alpha)$ rounds in the original graph. Moreover, as H contains at most two copies of each node in G , we can assign each node a unique b -bit identifier with $b = O(\log N)$.

One of the parameters of [Lemma 3.4](#) is a normalization constant, which we define below.

Definition 4.5 (Normalization Constant). *For every $i \in [0, 10\alpha - 1]$ and $j \in [0, \log(N)/\alpha]$, we define $Norm_{i,j} = \frac{2^{i \cdot \log(N)/\alpha + j}}{N^2}$.*

We now invoke [Lemma 3.4](#) with parameters $\Delta = \deg_j$, $k = \lceil 100 \log \log(N) \rceil$ and $p = \frac{1}{16}$ to conclude that there exists a LOCAL algorithm running in $O(kp(\log^2(k) + \log(k) \log^* b)) = \tilde{O}(\log \log n)$ rounds on H , and hence can be simulated in $\tilde{O}(\log \log n)\alpha$ rounds in G , which computes a set $V^{sub} \subseteq V_H$ satisfying

$$\begin{aligned} & \sum_{u \in U_H : |N_H(u) \cap V^{sub}| \leq 0.5 \lceil \Delta/k \rceil} x_u + Norm_{i,j} \cdot |V^{sub}| \\ & \leq 4 \left(e^{-pk} \sum_{u \in U_H} x_u + Norm_{i,j} \cdot p \cdot |V_H| \right). \end{aligned}$$

We then define $V_{i,j}^{active} = V^{sub}$. Consider some node u that is bad in step j of phase i . By definition, $|S_i(u) \cap V_{i,j}^{active}| \geq \deg_j$ and $|S_i(u) \cap V_{i,j+1}^{active}| < \deg_{j+1}$. Therefore, $u \in U_H$ and

$$\begin{aligned} |N_H(u) \cap V^{sub}| & \leq |S_i(u) \cap V_{i,j+1}^{active}| < \deg_{j+1} \\ & = \frac{\deg_j}{\lceil 1000 \log \log(N) \rceil} \leq 0.5 \lceil \Delta/k \rceil \end{aligned}$$

Since $e^{-pk} = e^{-(100/16) \log \log N} \leq \frac{1}{800 \log(N)}$, we can conclude that

$$\begin{aligned} & \sum_{u \in U_{i,j}^{bad}} x_u + Norm_{i,j} \cdot |V_{i,j+1}^{active}| \\ & \leq \frac{1}{200 \log(N)} \sum_{u \in V} x_u + \frac{Norm_{i,j}}{4} |V_{i,j+1}^{active}|. \end{aligned}$$

From this, we can deduce the following two claims:

Claim 4.3 (No active nodes in the end). *For every $i \in [0, 10\alpha - 1]$ and $j \in [0, \log(N)/\alpha]$, we have*

$$|V_{i,j}^{active}| \leq \frac{\sum_{u \in V} x_u}{50 \log(N) Norm_{i,j}},$$

which in particular also implies that $V_{10\alpha}^{active} = \emptyset$.

PROOF. First, note that

$$|V_{0,0}^{active}| \leq n \leq \frac{1}{50 \log(N)(1/N^2)} \leq \frac{\sum_{u \in V} x_u}{50 \log(N) Norm_{0,0}}.$$

Also, recall that for every $i \in [1, 10\alpha - 1]$, $V_{i+1,0}^{active} := V_{i+1}^{active} := V_{i, \log(N)/\alpha}^{active}$. Together with $Norm_{i+1,0} = Norm_{i, \log(N)/\alpha}$, we get that $|V_{i, \log(N)/\alpha}^{active}| \leq \frac{\sum_{u \in V} x_u}{50 \log(N) Norm_{i, \log(N)/\alpha}}$ trivially implies that $|V_{i+1,0}^{active}| \leq \frac{\sum_{u \in V} x_u}{50 \log(N) Norm_{i+1,0}}$.

Now, assume that $|V_{i,j}^{active}| \leq \frac{\sum_{u \in V} x_u}{50 \log(N) Norm_{i,j}}$ for a given $i \in [0, 10\alpha - 1]$ and $j \in [0, \log(N)/\alpha - 1]$. This implies that

$$\begin{aligned} |V_{i,j+1}^{active}| & \leq \frac{\sum_{u \in V} x_u}{200 \log(N) Norm_{i,j}} + \frac{1}{4} |V_{i,j}^{active}| \\ & \leq \frac{\sum_{u \in V} x_u}{100 \log(N) Norm_{i,j}} = \frac{\sum_{u \in V} x_u}{50 \log(N) Norm_{i,j+1}} \sum_{u \in V} x_u. \end{aligned}$$

Thus, $|V_{i,j}^{active}| \leq \frac{\sum_{u \in V} x_u}{50 \log(N) Norm_{i,j}}$, as one can formalize by an induction. In particular,

$$|V_{10\alpha-1, \log(N)/\alpha}^{active}| \leq \frac{\sum_{u \in V} x_u}{50 \log(N) Norm_{10\alpha-1, \log(N)/\alpha}} < 1 \quad \square$$

and as $V_{10\alpha}^{active} := V_{10\alpha-1, \log(N)/\alpha}^{active}$, we get that $V_{10\alpha}^{active} = \emptyset$.

Claim 4.4 (Total number of bad nodes in step j of phase i). *For every $i \in [0, 10\alpha - 1]$ and $j \in [0, \log(N)/\alpha - 1]$, we have*

$$\sum_{u \in U_{i,j}^{bad}} x_u \leq \frac{1}{100 \log(N)} \sum_{u \in V} x_u.$$

PROOF. We have

$$\sum_{u \in U_{i,j}^{bad}} x_u \leq \frac{1}{200 \log(N)} \sum_{u \in V} x_u + \frac{Norm_{i,j}}{4} |V_{i,j}^{active}| \leq \frac{\sum_{u \in V} x_u}{100 \log(N)}. \quad \square$$

PROOF OF THEOREM 2.2. According to [Claim 4.1](#), the diameter of the output partition is $O(\alpha)$.

[Lemma 4.2](#) implies that $V \setminus (\cup_{i=0}^{10\alpha-1} U_i^{bad}) \subseteq \{u \in V : \deg_G(u) = 2\tilde{O}(\log n)/\alpha\}$ and [Claim 4.4](#) implies that $\sum_{u \in \cup_{i=0}^{10\alpha-1} U_{i,j}^{bad}} x_u \leq 10\alpha \cdot (\log(N)/\alpha) \frac{\sum_{u \in V} x_u}{100 \log(N)} = \frac{\sum_{u \in V} x_u}{10}$. Together with $U_i^{bad} \subseteq \cup_{j=0}^{\log(N)/\alpha} U_{i,j}^{bad}$ for every $j \in [0, \log(N)/\alpha - 1]$, we therefore get $\sum_{u \in V : \deg_G(u) = 2\tilde{O}(\log n)/\alpha} x_u \geq 0.9 \sum_{u \in V} x_u$, as needed.

We have argued that each step takes $\tilde{O}(\log \log n \cdot \alpha)$ rounds. There are $O(\alpha) \cdot O(\log n/\alpha)$ steps in total, and thus the total runtime to compute the delays is $\tilde{O}(\alpha \log n)$. Given the delays, the output partition can be computed in $O(\alpha)$ additional rounds. Hence, the algorithm runs in $\tilde{O}(\alpha \log n)$ rounds. \square

4.2 Clustering with All Nodes Good

This section is dedicated to proving [Theorem 2.1](#). As a reminder, the goal is to compute a partition with diameter $O(\alpha)$ such that all vertices have a cluster degree of $2\tilde{O}(\log n)/\alpha$.

As in the previous section, the algorithm follows the outline of [Algorithm 1](#). The only missing part is to specify how to compute V_{i+1}^{active} given V_i^{active} for every $i \in [0, 10\alpha - 1]$.

In phase i , the goal is to compute V_{i+1}^{active} in such a way that all nodes u with $|S_i(u) \cap V_i^{active}| = 2^{\Omega(\log n)/\alpha}$, which we refer to as being important in phase i , satisfy $S_i(u) \cap V_{i+1}^{active} \neq \emptyset$. By using [Lemma 4.2](#), this is sufficient to ensure that all vertices have a cluster degree of $2^{\tilde{O}(\log n)/\alpha}$. We start with the formal definition of being important in phase i . Recall that we use N to denote a polynomial upper bound on the number of nodes, which is known to all the nodes.

Definition 4.6 (Important in phase i , $U_i^{important}$). *For every $i \in [0, 10\alpha - 1]$, we refer to a node as being important in phase i if $|S_i(u)| \geq (2000 \log N)^{\log(N)/\alpha}$.*

As before, phase i consists of $\log(N)/\alpha$ steps and we compute a sequence of sets $V_{i,0}^{active} := V_i^{active} \supseteq V_{i,1}^{active} \supseteq V_{i,2}^{active} \supseteq \dots \supseteq V_{i,\log(N)/\alpha}^{active} := V_{i+1}^{active}$.

Compared to [Section 4.1](#), one crucial difference is that the algorithm starts computing $V_{i,j+1}^{active}$ before it has finished computing $V_{i,j}^{active}$.

As before, we associate a degree threshold with each step. We remark that the degree threshold is slightly larger compared to the one in [Section 4.1](#).

Definition 4.7 (\deg_j). *For every $j \in [0, \log(N)/\alpha]$, we define $\deg_j = (2000 \log N)^{\log(N)/\alpha-j}$.*

Note that u is important in phase i if $|S_i(u)| \geq \deg_0$ and that $S_i(u) \cap V_{i+1}^{active} \neq \emptyset$ is equivalent to $|S_i(u) \cap V_{i+1}^{active}| = |S_i(u) \cap V_{i,\log(N)/\alpha}^{active}| \geq \deg_{\log(N)/\alpha}$. In general, each node u which is important in phase i will satisfy $|S_i(u) \cap V_{i,j}^{active}| \geq \deg_j$ for every $j \in [0, \log(N)/\alpha]$.

For each $i \in [0, 10\alpha - 1]$ and $j \in [1, \log(N)/\alpha]$, we compute $V_{i,j}^{active}$ by computing a sequence of sets $V_{i,j,0}^{active} := \emptyset \subseteq V_{i,j,1}^{active} \subseteq \dots \subseteq V_{i,j,4\log(N)}^{active} := V_{i,j}^{active}$. To simplify notation, we also define $V_{i,0,0}^{active} = V_{i,0,1}^{active} = \dots = V_{i,0,4\log(N)}^{active} = V_{i,0}^{active} := V_i^{active}$ for every $i \in [0, 10\alpha - 1]$.

For every $j \in [0, \log(N)/\alpha]$ and $\ell \in [0, 4\log(N)]$, we define a partition $U_i^{important} = U_{i,j,\ell}^{good} \sqcup U_{i,j,\ell}^{bad}$.

Definition 4.8 ($U_{i,j,\ell}^{good}$, $U_{i,j,\ell}^{bad}$). *For every $i \in [0, 10\alpha - 1]$, $j \in [0, \log(N)/\alpha]$, $\ell \in [0, 4\log(N)]$, we define*

$$U_{i,j,\ell}^{good} = \{u \in U_i^{important} : |S_i(u) \cap V_{i,j,\ell}^{active}| \geq \deg_j\}$$

and

$$U_{i,j,\ell}^{bad} = U_i^{important} \setminus U_{i,j,\ell}^{good}.$$

Computing $V_{i,j,\ell}^{active}$ given $V_{i,j-1,\ell}^{active}$ and $V_{i,j,\ell-1}^{active}$. Now, fix some $i \in [0, 10\alpha - 1]$, $j \in [1, \log(N)/\alpha]$ and $\ell \in [1, 4\log(N)]$. We compute $V_{i,j,\ell}^{active}$ given $V_{i,j-1,\ell}^{active}$ and $V_{i,j,\ell-1}^{active}$ by first defining a bipartite graph H , then invoking the hitting set subroutine of [Lemma 3.4](#) to deduce that there exists an efficient LOCAL algorithm on the bipartite graph which computes a subset of $V_{i,j-1,\ell}^{active}$ with the desired properties, and finally use the fact that we can simulate the

algorithm on the bipartite graph on the original graph with an $O(\alpha)$ -multiplicative overhead. Consider the bipartite graph $H = (U_H \sqcup V_H, E_H)$ with $U_H = U_{i,j-1,\ell}^{good} \cap U_{i,j,\ell-1}^{bad}$, $V_H = V_{i,j-1,\ell}^{active}$ and where we connect each node $u \in U_H$ to \deg_{j-1} nodes in $S_i(u) \cap V_{i,j-1,\ell}^{active}$. Similar as in [Section 4.1](#), each round in H can be simulated in $O(\alpha)$ rounds in the original graph and we can assume that each node is assigned a unique b -bit identifier with $b = O(\log N)$. We also define $Norm_{i,j} = 2^{i \cdot \log(N)/\alpha + j}$ for every $i \in [0, 10\alpha - 1]$ and $j \in [0, \log(N)/\alpha]$.

We now invoke [Lemma 3.4](#) with parameters $\Delta = \deg_{j-1}$, $k = 500 \log(N)$, $p = \frac{1}{64 \log N}$ and $Norm = Norm_{i,j} 2^{j-\ell}$ to conclude that there exists a LOCAL algorithm running in $O(\log^2 \log n)$ rounds on H , and hence can be simulated in $O(\alpha \log^2 \log n)$ rounds in G , which computes a set $V_{i,j,\ell}^{sub} \subseteq V_H$ satisfying

$$\begin{aligned} |\{u \in U_H : |N_H(u) \cap V_{i,j,\ell}^{sub}| \leq 0.5 \lfloor \Delta/k \rfloor\}| + 2^{j-\ell} Norm_{i,j} \cdot |V_{i,j,\ell}^{sub}| \\ \leq 4 \left(e^{-pk} |U_H| + 2^{j-\ell} Norm_{i,j} \cdot p \cdot |V_H| \right). \end{aligned}$$

We then define $V_{i,j,\ell}^{active} = V_{i,j,\ell-1}^{active} \cup V_{i,j,\ell}^{sub}$.

Consider some $u \in U_{i,j-1,\ell}^{good} \cap U_{i,j,\ell}^{bad}$. As $u \in U_{i,j,\ell}^{bad}$ implies $u \in U_{i,j,\ell-1}^{bad}$, we have $u \in U_{i,j-1,\ell}^{good} \cap U_{i,j,\ell-1}^{bad} := U_H$. Moreover,

$$\begin{aligned} |N_H(u) \cap V_{i,j,\ell}^{sub}| &\leq |S_i(u) \cap V_{i,j,\ell}^{active}| < \deg_j \\ &= \frac{\deg_{j-1}}{2000 \log(N)} \leq 0.5 \lfloor \frac{\Delta}{k} \rfloor. \end{aligned}$$

Hence, using that $e^{-pk} \leq \frac{1}{32}$, we get

$$\begin{aligned} 2^{\ell-j} |U_{i,j-1,\ell}^{good} \cap U_{i,j,\ell}^{bad}| + Norm_{i,j} \cdot |V_{i,j,\ell}^{sub}| \\ \leq \frac{2^{\ell-j}}{8} |U_{i,j,\ell-1}^{bad}| + \frac{Norm_{i,j}}{16 \log(N)} |V_{i,j-1,\ell}^{active}|. \end{aligned}$$

The claim below is the key claim in the analysis.

Claim 4.5. *For every $i \in [0, 10\alpha - 1]$, $j \in [0, \log(N)/\alpha]$ and $\ell \in [0, 4\log(N)]$, it holds that*

- (1) $|V_{i,j}^{active}| \leq \frac{2 \log(N) \cdot n}{Norm_{i,j}}$ and
- (2) $|U_{i,j,\ell}^{bad}| \leq n \cdot 2^{j-\ell}$.

In particular, $V_{i,0}^{active} = \emptyset$ and $\{u \in U_i^{important} : S_i(u) \cap V_{i+1}^{active} = \emptyset\} = \emptyset$ for every $i \in [0, 10\alpha - 1]$.

PROOF. First, consider some fixed $i \in [0, 10\alpha - 1]$ and $j \in [1, \log(N)/\alpha]$ and assume $|V_{i,j-1}^{active}| \leq \frac{2 \log(N) \cdot n}{Norm_{i,j-1}}$ and $|U_{i,j-1,\ell}^{bad}| \leq n \cdot 2^{(j-1)-\ell}$ for every $\ell \in [0, 4\log(N)]$. In particular, for every $\ell \in [0, 4\log(N) - 1]$, $V_{i,j-1,\ell}^{active} \subseteq V_{i,j-1}^{active}$ directly gives

$$\frac{Norm_{i,j}}{16 \log(N)} |V_{i,j-1,\ell}^{active}| \leq \frac{Norm_{i,j}}{16 \log(N)} \frac{2 \log(N) \cdot n}{Norm_{i,j-1}} = \frac{n}{4}.$$

We first show by induction on ℓ that this implies $|U_{i,j,\ell}^{bad}| \leq n \cdot 2^{j-\ell}$ for every $\ell \in [0, 4\log(N)]$ and afterwards we show that $|V_{i,j}^{active}| \leq \frac{2 \log(N) \cdot n}{Norm_{i,j}}$.

The base case $\ell = 0$ trivially holds as $|U_{i,j,0}^{bad}| \leq n \leq n \cdot 2^{j-0}$. Now, fix some $\ell \in [1, 4\log(N)]$ and assume that $|U_{i,j,\ell-1}^{bad}| \leq 2^{j-(\ell-1)}$.

First, note that $U_{i,j,\ell}^{bad} \subseteq V_i^{important} = U_{i,j-1,\ell}^{good} \sqcup U_{i,j-1,\ell}^{bad}$ together with the initial assumption $|U_{i,j-1,\ell}^{bad}| \leq 2^{(j-1)-\ell}$ implies

$$|U_{i,j,\ell}^{bad}| \leq |U_{i,j-1,\ell}^{bad}| + |U_{i,j-1,\ell}^{good} \cap U_{i,j,\ell}^{bad}| \leq \frac{n \cdot 2^{j-\ell}}{2} + |U_{i,j-1,\ell}^{good} \cap U_{i,j,\ell}^{bad}|.$$

Moreover, using induction, we get

$$\begin{aligned} |U_{i,j-1,\ell}^{good} \cap U_{i,j,\ell}^{bad}| &\leq \frac{1}{8} |U_{i,j-1,\ell}^{bad}| + 2^{j-\ell} \frac{Norm_{i,j}}{16 \log(N)} |V_{i,j-1,\ell}^{active}| \\ &\leq \frac{n 2^{j-(\ell-1)}}{8} + \frac{n \cdot 2^{j-\ell}}{4} \leq \frac{n \cdot 2^{j-\ell}}{2} \end{aligned}$$

and therefore $|U_{i,j,\ell}^{bad}| \leq n \cdot 2^{j-\ell}$, finishing the induction. Now, consider some $\ell \in [1, 4 \log(N)]$. Using $|U_{i,j,\ell-1}^{bad}| \leq n \cdot 2^{j-(\ell-1)}$, we get

$$\begin{aligned} |V_{i,j,\ell}^{sub}| &\leq \frac{1}{Norm_{i,j}} \left(\frac{2^{\ell-j}}{8} |U_{i,j,\ell-1}^{bad}| + \frac{Norm_{i,j}}{16 \log(N)} |V_{i,j-1,\ell}^{active}| \right) \\ &\leq \frac{(n/4) + (n/4)}{Norm_{i,j}} = \frac{n}{2 Norm_{i,j}}. \end{aligned}$$

As $V_{i,j}^{active} = \bigcup_{\ell=1}^{4 \log(N)} V_{i,j,\ell}^{sub}$, we conclude $|V_{i,j}^{active}| \leq \frac{2 \log(N) \cdot n}{Norm_{i,j}}$.

Consider some fixed $i \in [0, 10\alpha - 1]$ and assume that $|V_{i,0}^{active}| \leq \frac{2 \log(N) \cdot n}{Norm_{i,0}}$. We use induction on j to conclude that $|V_{i,j}^{active}| \leq \frac{2 \log(N) \cdot n}{Norm_{i,j}}$ and $|U_{i,j,\ell}^{bad}| \leq n \cdot 2^{j-\ell}$ for every $j \in [0, \log(N)/\alpha]$ and $\ell \in [0, 4 \log(N)]$. For the base case $j = 0$, note that $|V_{i,0}^{active}| \leq \frac{2 \log(N) \cdot n}{Norm_{i,0}}$ is just what we assumed and as $V_{i,0,\ell}^{active} := V_{i,0}^{active}$ for every $\ell \in [0, 4 \log(N)]$, we have $|U_{i,0,\ell}^{bad}| = 0 \leq n \cdot 2^{0-\ell}$. We already did the induction step going from $j-1$ to j above and therefore we are done.

Finally, we prove [Claim 4.5](#) by induction on i . For the base case $i = 0$, note that $|V_{0,0}^{active}| \leq n \leq \frac{2 \log(N) \cdot n}{Norm_{0,0}}$. By using the previous induction on j , this is enough to show that the bounds of [Claim 4.5](#) hold for $i = 0$. Now, consider some $i \in [1, 10\alpha - 1]$ and assume that the bounds hold for $i-1$. In particular,

$$|V_{i,0}^{active}| = |V_{i-1, \log(N)/\alpha}^{active}| \leq \frac{2 \log(N) \cdot n}{Norm_{i-1, \log(N)/\alpha}} = \frac{2 \log(N) \cdot n}{Norm_{i,0}}.$$

Now, we can again use the previous induction on j to finish the induction step.

To conclude [Claim 4.5](#), note that $|V_{10\alpha-1, \log(N)/\alpha}^{active}| \leq \frac{2 \log(N) \cdot n}{Norm_{0,0}} < 1$ and as $V_{10\alpha}^{active} := V_{10\alpha-1, \log(N)/\alpha}^{active}$, we get that $V_{10\alpha}^{active} = \emptyset$. Consider some $i \in [0, 10\alpha - 1]$. Note that $|U_{i, \log(N)/\alpha, 4 \log(N)}^{bad}| \leq n \cdot 2^{\log(N)/\alpha - 4 \log(N)} < 1$ and therefore $U_{i, \log(N)/\alpha, 4 \log(N)}^{bad} = \emptyset$.

In particular, there does not exist a node $u \in U_i^{important}$ with $|S_i(u) \cap V_{i, \log(N)/\alpha, 4 \log(N)}^{active}| < \deg_{\log(N)/\alpha} = 1$. As $V_{i+1}^{active} := V_{i, \log(N)/\alpha}^{active} := V_{i, \log(N)/\alpha, 4 \log(N)}^{active}$, we therefore get that there exists no node $u \in U_i^{important}$ with $|S_i(u) \cap V_{i+1}^{active}| = \emptyset$, as desired. \square

PROOF OF THEOREM 2.1. According to [Claim 4.1](#), the diameter of the output partition is $O(\alpha)$. Moreover, combining [Claim 4.5](#) with [Lemma 4.2](#) shows that the clustering of each node is $2^{\tilde{O}(\log n)/\alpha}$. It remains to argue about the round complexity of the algorithm.

Fix some $i \in [0, 10\alpha - 1]$. Note that given V_i^{active} , we can compute $V_{i,j,0}^{active}$ for every $j \in [0, \log(N)/\alpha]$ and $V_{i,0,\ell}^{active}$ for every $\ell \in [1, 4 \log(N)]$ without further communication. Moreover, we have seen above that given $V_{i,j-1,\ell}^{active}$ and $V_{i,j,\ell-1}^{active}$, we can compute $V_{i,j,\ell}^{active}$ in $O(\alpha \log^2 \log n)$ rounds, for every $j \in [1, \log(N)/\alpha]$ and $\ell \in [1, 4 \log(N)]$. Hence, a simple induction shows that given V_i^{active} , we can compute $V_{i,j,\ell}^{active}$ in $(j+\ell)O(\alpha \log^2 \log n)$ rounds for every $j \in [0, \log(N)/\alpha]$ and $\ell \in [0, 4 \log(N)]$ and therefore we can compute V_{i+1}^{active} in $(\log(N)/\alpha + 4 \log(N)) \cdot O(\alpha \log^2 \log n) = \tilde{O}(\alpha \log n)$ rounds. Hence, the overall algorithm runs in $\tilde{O}(\alpha^2 \log n)$ rounds, as desired. \square

5 MIS

In this section, we use the low-diameter clustering results obtained in the previous section to obtain a deterministic distributed algorithm that computes an MIS in $\tilde{O}(\log^2 n)$ rounds of the LOCAL model.

THEOREM 1.1. *There is a deterministic distributed algorithm, in the LOCAL model, that computes a maximal independent set (MIS) in $\tilde{O}(\log^2 n)$ rounds. This also implies that there are LOCAL-model deterministic distributed algorithms with $\tilde{O}(\log^2 n)$ rounds complexity also for maximal matching, $(deg+1)$ -list vertex coloring, and $(2deg-1)$ -list edge coloring⁶.*

The rest of the section is devoted to the proof of [Theorem 1.1](#). The missing proofs are deferred to the full version. We first recall Luby's classic algorithm[27], which in each iteration chooses an independent set of nodes such that, when we add them to the output and remove them from the graph along with their neighbors, in expectation a constant fraction of the edges of the graph are removed.

Luby's Randomized MIS Algorithm. The starting point is to recall Luby's classic randomized algorithm from [28]. Each iteration of it works as follows. We mark each node u with probability $1/(10 \deg(u))$. Then, for each edge $\{u, v\}$, let us orient the edge as $u \rightarrow v$ if and only if $\deg(u) < \deg(v)$ or $\deg(u) = \deg(v)$ and $ID(u) < ID(v)$. For each marked node u , we add u to the independent set if and only if v does not have a marked out-neighbor. Finally, as a clean-up step at the end of this iteration, we remove all nodes that have been added to the independent set along with their neighbors. We then proceed to the next iteration.

Derandomizing Luby via Rounding. It is well-known that in each iteration of Luby's algorithm a constant fraction of the edges of the remaining graph gets removed, in expectation. Hence, the process terminates in $O(\log n)$ iterations with probability $1 - 1/\text{poly}(n)$. We explain how to derandomize each iteration of the algorithm in $\tilde{O}(\log n)$ rounds, such that we still remove a constant fraction of the edges per iteration. For the rest of this proof, we focus on an arbitrary iteration, and we assume that $H = (V_H, E_H)$ is the graph induced by the remaining vertices at the beginning of this iteration.

⁶In the $(deg+1)$ -list vertex coloring, each node v having a prescribed list L_v of colors, of size $|L_v| \geq deg(v) + 1$ from which it should choose its color. In the $(2deg-1)$ -list edge coloring, each edge $e = \{u, v\}$ has a prescribed list L_e of colors, of size $|L_e| \geq deg(v) + deg(u) - 1$ from which it should choose its color. Both problems reduce to MIS by a reduction of Luby [25, 27].

Let $\vec{x} \in \{0, 1\}^{V_H}$ be the indicator vector of whether different nodes are marked, that is, we have $x_v = 1$ if v is marked and $x_v = 0$ otherwise. Let $R_v(\vec{x})$ be the indicator variable of the event that v gets removed, for the marking vector \vec{x} . Let $Z(\vec{x})$ be the corresponding number of removed edges. Luby's algorithm determines the markings \vec{x} randomly. Our task is to derandomize this and select the marked nodes in a deterministic way such that when we remove nodes added to the independent set (those marked nodes that do not have a marked out-neighbor) and their neighbors, along with all the edges incident on these nodes, at least a constant fraction of edges E get removed.

Below, we give a pairwise analysis which shows that the expected number of removed edges is $\Omega(|E_H|)$. The rounding framework of Faour et al. [13] would then allow us to select the marked nodes in a deterministic way while retaining the guarantee that $\Omega(|E_H|)$ edges are removed. The rounding procedure runs in $O(\log^2(1/p_{min})) = O(\log^2 \Delta)$ rounds where p_{min} is the smallest marking probability. In fact, this is how Faour et al. [13] obtained an MIS algorithm running in $O(\log n \log^2 \Delta)$ rounds. We derandomize the marking process in $\tilde{O}(\log n)$ rounds. To do so, we first perform an intra-cluster rounding step and only then apply the rounding framework of Faour et al. [13]. In the intra-cluster rounding step, we compute for each node a marking probability which is either 0 or $1/2^{\tilde{O}(\sqrt{\log n})}$ and such that the expected number of removed edges is still $\Omega(|E_H|)$. Then, we apply the rounding framework of Faour et al. to the same pairwise analysis as before, but this time with the new marking probabilities. As all non-zero marking probabilities are $1/2^{\tilde{O}(\sqrt{\log n})}$, the rounding procedure runs in just $(\log(2^{\tilde{O}(\sqrt{\log n})}))^2 = \tilde{O}(\log n)$ rounds, as desired.

Good and bad nodes and prevalence of edges incident on good nodes. We call any node $v \in V_H$ *good* if and only if it has at least $\deg(v)/3$ incoming edges. A node v that is not good is called *bad*. It can be proven [28] that

$$\sum_{\text{good vertex } v} \deg(v) \geq |E_H|/2. \quad (1)$$

Lower bounding removed edges. We can lower bound the number of removed edges as

$$Z(\vec{x}) \geq \sum_{\text{good vertex } v} \deg(v) \cdot R_v(\vec{x})/2.$$

The 2 factor in the denominator is because for an edge, both endpoints might be good nodes. Since $\sum_{\text{good vertex } v} \deg(v) \geq |E_H|/2$, to prove that $\mathbb{E}[Z(\vec{x})] = \Omega(|E_H|)$, it suffices to show that each good vertex v has $\Pr[R_v(\vec{x})] = \Omega(1)$. This fact can be proven via elementary probability calculations. Next, we discuss how to prove it using only pairwise independence in the analysis.

Pessimistic estimator of removed edges via pairwise-independent analysis. Let us use $IN(u)$ and $OUT(u)$ to denote in-neighbors and out-neighbors of a vertex u . Consider a good node v and consider all its incoming neighbors u , i.e., neighbors u such that $(\deg(u), ID(u)) < (\deg(v), ID(v))$. Since v is good, it has at least $\deg(v)/3$ such neighbors. Hence, we have

$$\sum_{\text{incoming neighbor } u} \frac{1}{\deg(u)} \geq 1/3.$$

Choose a subset $IN^*(v) \subseteq IN(v)$ of incoming neighbors such that

$$\sum_{u \in IN^*(v)} \frac{1}{\deg(u)} \in [1/3, 4/3]. \quad (2)$$

Notice that such a subset $IN^*(v)$ exists since the summation over all incoming neighbors is at least $1/3$ and each neighbor contributes at most 1 to the summation. On the other hand, notice that for any node u , we have

$$\sum_{w \in OUT(u)} \frac{1}{\deg(w)} \leq 1. \quad (3)$$

This is because $|OUT(u)| \leq \deg(u)$ and for each $w \in OUT(u)$, we have $(\deg(w), ID(w)) > (\deg(u), ID(u))$.

A sufficient event $\mathcal{E}(v, u)$ that causes v to be removed is if some $u \in IN^*(v)$ is marked and no other node in $IN^*(v) \cup OUT(u)$ is marked. By union bound, this event's indicator is lower bounded by

$$x_u - \sum_{u' \in IN^*(v), u' \neq u} x_u \cdot x_{u'} - \sum_{w \in OUT(u)} x_u \cdot x_w.$$

Furthermore, the events $\mathcal{E}(v, u_1), \mathcal{E}(v, u_2), \dots, \mathcal{E}(v, u_{|IN^*(v)|})$ are mutually disjoint for different $u_1, u_2, \dots, u_{|IN^*(v)|} \in IN^*(v)$. Hence, we can sum over these events for different $u \in IN^*(v)$ and conclude that

$$\begin{aligned} R_v(\vec{x}) &\geq \sum_{u \in IN^*(v)} \left(x_u - \sum_{u' \in IN^*(v), u' \neq u} x_u \cdot x_{u'} - \sum_{w \in OUT(u)} x_u \cdot x_w \right) \\ &= \sum_{u \in IN^*(v)} x_u - \sum_{u, u' \in IN^*(v)} x_u \cdot x_{u'} \\ &\quad - \sum_{u \in IN^*(v)} \sum_{w \in OUT(u)} x_u \cdot x_w \end{aligned}$$

Therefore, our overall pessimistic estimator for the number of removed edges gives that

$$\begin{aligned} Z(\vec{x}) &\geq \sum_{\text{good vertex } v} (\deg(v)/2) \cdot R_v(\vec{x}) \\ &\geq \sum_{\text{good vertex } v} (\deg(v)/2) \cdot \left(\sum_{u \in IN^*(v)} x_u - \sum_{u, u' \in IN^*(v)} x_u \cdot x_{u'} \right. \\ &\quad \left. - \sum_{u \in IN^*(v)} \sum_{w \in OUT(u)} x_u \cdot x_w \right). \end{aligned}$$

Intra-Cluster Rounding. For the intra-cluster rounding step, we assume that we are given a partition C_G of the input graph G with diameter $O(\sqrt{\log n})$ and such that $\deg_{C_G}(u) \leq f(n)$ for every $u \in V$ and for some function $f(n) = 2^{\tilde{O}(\sqrt{\log n})}$.

Indeed, invoking [Theorem 2.1](#) with $\alpha = \sqrt{\log(N)}$, we can compute such a partition C_G at the very beginning in $\tilde{O}(\log^2 n)$ rounds. Also, we denote by $C_H = \{C \cap V_H : C \in C_G\}$ the partition of H that

we obtain from the partition C_G by removing from each cluster the vertices not contained in V_H .

Now, consider that we relax the label assignment such that it also allows for a fractional assignment $\vec{x} \in [0, 1]^{V_H}$. Then, $Z(\vec{x})$ is a pessimistic estimator on the expected number of edges removed if we mark each vertex u fully independently (or pairwise independent) with probability x_u . In Luby's algorithm, one marks each node $u \in V_H$ with probability $x_u := \frac{1}{20 \deg(u)}$, and a simple calculation shows that for this marking probability we get $Z(\vec{x}) = \Omega(|E_H|)$. In the intra-cluster rounding step, the goal is to compute a $\frac{1}{10000f(n) \log(n)}$ -integral assignment $\vec{x}^{intra} \in [0, 1]^{V_H}$ such that we have $Z(\vec{x}^{intra}) = \Omega(|E_H|)$.

Consider some fixed cluster $C \in C_H$. For any $u \in C$, we define \vec{x}^{intra} in such a way that we can compute x_u^{intra} by only knowing the cluster C (together with the k -hop neighborhood around C for some $k = O(1)$). As the weak-diameter of C in the original input graph is $O(\sqrt{\log n})$, we can compute x_u^{intra} in $O(\sqrt{\log n})$ rounds. Similar as in the intra-cluster rounding step for constant approximate matching, we will use the probabilistic method to show that a fractional assignment with certain desirable properties exist.

To that end, we introduce one random variable X_u^{intra} for every vertex $u \in C$. If $\deg(u) \leq 1000f(n) \log(n)$, we simply set $X_u^{intra} = \frac{1}{10 \deg(u)}$. Otherwise, if $\deg(u) > 1000f(n) \log(n)$, we set $X_u^{intra} = \frac{1}{10000f(n) \log(n)}$ with probability $\frac{1000f(n) \log(n)}{\deg(u)}$ and with the remaining probability we set $X_u^{intra} = 0$, fully independently. Note that for every $u \in C$, $\mathbb{E}[X_u] = \frac{1}{10 \deg(u)}$.

The following claim follows by a simple Chernoff bound.

Claim 5.1. *For any subset $S \subseteq C$, we have*

$$\Pr \left[\sum_{u \in S} X_u^{intra} \leq \frac{1}{5} \sum_{u \in S} \frac{1}{\deg(u)} + \frac{1}{100f(n)} \right] \geq 1 - \frac{1}{n^{10}} \text{ and}$$

$$\Pr \left[\sum_{u \in S} X_u^{intra} \geq \frac{1}{20} \sum_{u \in S} \frac{1}{\deg(u)} - \frac{1}{100f(n)} \right] \geq 1 - \frac{1}{n^{10}}.$$

Claim 5.1 together with a simple union bound implies that the following holds with strictly positive probability: for every good vertex v ,

$$\frac{1}{20} \sum_{u \in IN^*(v) \cap C} \frac{1}{\deg(u)} - \frac{1}{100f(n)} \leq \sum_{u \in IN^*(v) \cap C} X_u^{intra}$$

$$\leq \frac{1}{5} \sum_{u \in IN^*(v) \cap C} \frac{1}{\deg(u)} + \frac{1}{100f(n)}$$

and for every vertex $u \in V_H$,

$$\sum_{w \in OUT(u) \cap C} X_w^{intra} \leq \frac{1}{5} \sum_{w \in OUT(u) \cap C} \frac{1}{\deg(w)} + \frac{1}{100f(n)}.$$

Therefore, each cluster C can compute in $O(\sqrt{\log n})$ rounds a value $x_u^{intra} \in \{0\} \cup \left[\frac{1}{10000f(n) \log(n)}, 1 \right]$ for every node $u \in C$ such

that for every good vertex v ,

$$\frac{1}{20} \sum_{u \in IN^*(v) \cap C} \frac{1}{\deg(u)} - \frac{1}{100f(n)} \leq \sum_{u \in IN^*(v) \cap C} x_u^{intra}$$

$$\leq \frac{1}{5} \sum_{u \in IN^*(v) \cap C} \frac{1}{\deg(u)} + \frac{1}{100f(n)}$$

and for every vertex $u \in V_H$,

$$\sum_{w \in OUT(u) \cap C} x_w^{intra} \leq \frac{1}{5} \sum_{w \in OUT(u) \cap C} \frac{1}{\deg(w)} + \frac{1}{100f(n)}.$$

We use these properties together with the fact that the degree of the clustering C_H is at most $f(n)$ to prove the claim below.

Claim 5.2. *For every good vertex v , we have*

$$1/1000 \leq \sum_{u \in IN^*(v)} x_u^{intra} \leq 1/3.$$

For every $u \in V_H$, we have $\sum_{w \in OUT(u)} x_w^{intra} \leq 1/4$.

Local Rounding. We next round the fractional solution \vec{x}^{intra} to an integral solution \vec{y} using the rounding framework of Faour et al. [13], as discussed in [Section 3.2](#).

For a given label assignment $\vec{x} \in \{0, 1\}^{V_H}$, we define the utility function as

$$\mathbf{u}(\vec{x}) = \sum_{\text{good vertex } v} (\deg(v)/2) \cdot \left(\sum_{u \in IN^*(v)} x_u \right),$$

and the cost function as

$$\mathbf{c}(\vec{x}) = \sum_{\text{good vertex } v} (\deg(v)/2) \cdot \left(\sum_{u, u' \in IN^*(v)} x_u \cdot x_{u'} \right)$$

$$+ \sum_{u \in IN^*(v)} \sum_{w \in OUT(u)} x_u \cdot x_w.$$

If the label assignment is relaxed to be a fractional assignment $\vec{x} \in [0, 1]^{V_H}$, where intuitively now x_u is the probability of u being marked, the same definitions apply for the utility and cost of this fractional assignment.

Let H^2 denote the graph where any two nodes of distance at most 2 in H are connected by an edge. Note that $\mathbf{u}(\vec{x})$ is a utility function in the graph H^2 and similarly $\mathbf{c}(\vec{x})$ is a cost function in the graph H^2 . We next argue that the utility and cost function also satisfy the key requirement of [Lemma 3.3](#):

Claim 5.3. *For the fractional label assignment $\vec{x}^{intra} \in [0, 1]^{V_H}$ computed during the intra-cluster rounding step we have $\mathbf{u}(\vec{x}^{intra}) - \mathbf{c}(\vec{x}^{intra}) \geq \mathbf{u}(\vec{x}^{intra})/3$.*

Hence, we can apply [Lemma 3.3](#) on these fractional assignments with $\lambda_{min} = \frac{1}{100f(n) \log(n)} = 1/2\tilde{O}(\sqrt{\log n})$. The algorithm runs in $O(\log^2(1/\lambda_{min}) + \log(1/\lambda_{min}) \log^* n) = \tilde{O}(\log n)$ rounds in H^2 , and hence can be simulated with no asymptotic overhead in H , and as a result we get an integral label assignment $\vec{y} \in \{0, 1\}^{V_H}$ which satisfies $\mathbf{u}(\vec{y}) - \mathbf{c}(\vec{y}) \geq 0.5(\mathbf{u}(\vec{x}^{intra}) - \mathbf{c}(\vec{x}^{intra}))$. We know that $Z(\vec{y}) = \mathbf{u}(\vec{y}) - \mathbf{c}(\vec{y}) \geq (1/2) \cdot (\mathbf{u}(\vec{x}^{intra}) - \mathbf{c}(\vec{x}^{intra}))$. Next, we argue that this implies $Z(\vec{y}) \geq |E_H|/24000$.

Claim 5.4. For the fractional label assignment $\bar{x}^{intra} \in [0, 1]^{V_H}$ computed during the intra-cluster rounding step we have $Z(\bar{x}^{intra}) = \mathbf{u}(\bar{x}^{intra}) - \mathbf{c}(\bar{x}^{intra}) \geq |E_H|/12000$. Hence, for the integral marking assignment \bar{y} we obtain from rounding \bar{x} by invoking [Lemma 3.3](#), we have $Z(\bar{y}) = \mathbf{u}(\bar{y}) - \mathbf{c}(\bar{y}) \geq (1/2) \cdot (\mathbf{u}(\bar{x}) - \mathbf{c}(\bar{x})) \geq |E_H|/24000$.

Putting Everything Together. From the rounding procedure described above, which runs in $\tilde{O}(\log n)$ rounds of the LOCAL model, we get an integral marking assignment \bar{y} with the following guarantee: if we add marked nodes u that have no marked out-neighbor to the independent set and remove them along with their neighbors, we remove at least a $1/24000$ fraction of the remaining edges. Hence, $O(\log n)$ such iterations suffice to complete the computation and have a maximal independent set, for a total round complexity of $\tilde{O}(\log^2 n)$. Note that the low-diameter partition only has to be computed once in the beginning, which takes $\tilde{O}(\log^2 n)$ rounds. Thus, we can indeed compute an MIS in $\tilde{O}(\log^2 n)$ rounds of the LOCAL model.

REFERENCES

- [1] Noga Alon, László Babai, and Alon Itai. 1986. A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem. *Journal of Algorithms* 7, 4 (1986), 567–583.
- [2] Baruch Awerbuch. 1985. Complexity of Network Synchronization. *J. ACM* 32, 4 (1985), 804–823.
- [3] Baruch Awerbuch, Bonnie Berger, Lenore Cowen, and David Peleg. 1996. Fast Network Decompositions and Covers. *J. of Parallel and Distributed Computing* 39, 2 (1996), 105–114.
- [4] B. Awerbuch, A. V. Goldberg, M. Luby, and S. A. Plotkin. 1989. Network Decomposition and Locality in Distributed Computation. In *Proc. 30th IEEE Symp. on Foundations of Computer Science (FOCS)*. 364–369.
- [5] Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikael Rabie, and Jukka Suomela. 2019. Lower bounds for maximal matchings and maximal independent sets. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*.
- [6] Philipp Bamberger, Fabian Kuhn, and Yannic Maus. 2020. Efficient deterministic distributed coloring with small bandwidth. In *ACM Symp. on Principles of Distributed Computing (PODC)*. 243–252.
- [7] Leonid Barenboim. 2012. On the Locality of Some NP-Complete Problems. In *Proc. 39th Coll. on Automata, Languages, and Programming (ICALP)*. 403–415.
- [8] Keren Censor-Hillel, Merav Parter, and Gregory Schwartzman. 2020. Derandomizing Local Distributed Algorithms under Bandwidth Restrictions. *Distributed Comput.* 33, 3-4 (2020), 349–366.
- [9] Yi-Jun Chang and Mohsen Ghaffari. 2021. Strong-diameter network decomposition. In *ACM Symposium on Principles of Distributed Computing (PODC)*. 273–281.
- [10] Yi-Jun Chang, Wenzheng Li, and Seth Pettie. 2018. An Optimal Distributed $(\Delta + 1)$ -Coloring Algorithm?. In *ACM Symp. on Theory of Computing (STOC)*.
- [11] Michael Elkin, Bernhard Haeupler, Václav Rozhoň, and Christoph Grunau. 2022. Deterministic Low-Diameter Decompositions for Weighted Graphs and Distributed and Parallel Applications. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, to appear, arXiv:2204.08254.
- [12] Michael Elkin and Ofer Neiman. 2016. Distributed Strong Diameter Network Decomposition. In *ACM Symp. on Principles of Distributed Computing (PODC)*. 211–216.
- [13] Salwa Faour, Mohsen Ghaffari, Christoph Grunau, Fabian Kuhn, and Václav Rozhoň. 2023. Local Distributed Rounding: Generalized to MIS, Matching, Set Cover, and Beyond. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, to appear, arXiv:2209.11651.
- [14] Manuela Fischer. 2020. Improved deterministic distributed matching via rounding. *Distributed Computing* 33, 3 (2020), 279–291.
- [15] Manuela Fischer, Mohsen Ghaffari, and Fabian Kuhn. 2017. Deterministic Distributed Edge-Coloring via Hypergraph Maximal Matching. In *IEEE Symp. on Foundations of Computer Science (FOCS)*.
- [16] Manuela Fischer, Slobodan Mitrović, and Jara Uitto. 2022. Deterministic $(1 + \epsilon)$ -approximate maximum matching with $\text{poly}(1/\epsilon)$ passes in the semi-streaming model and beyond. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 248–260.
- [17] Mohsen Ghaffari, Christoph Grunau, Bernhard Haeupler, Saeed Ilchi, and Václav Rozhoň. 2023. Improved Distributed Network Decomposition, Hitting Sets, and Spanners, via Derandomization. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, to appear, arXiv:2209.11669.
- [18] Mohsen Ghaffari, Christoph Grunau, and Václav Rozhoň. 2021. Improved Deterministic Network Decomposition. In *ACM-SIAM Symp. on Discrete Algorithms (SODA)*. 2904–2923.
- [19] Mohsen Ghaffari, David Harris, and Fabian Kuhn. 2018. On derandomizing local distributed algorithms. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*. 662–673.
- [20] Mohsen Ghaffari and Fabian Kuhn. 2021. Deterministic Distributed Vertex Coloring: Simpler, Faster, and without Network Decomposition. In *Proc. 62nd IEEE Symp. on Foundations of Computer Science (FOCS)*. 1009–1020.
- [21] Mohsen Ghaffari, Fabian Kuhn, and Yannic Maus. 2017. On the Complexity of Local Distributed Graph Problems. In *Proc. 49th ACM Symp. on Theory of Computing (STOC)*. 784–797.
- [22] Michał Hańćkowiak, Michał Karoński, and Alessandro Panconesi. 2001. On the Distributed Complexity of Computing Maximal Matchings. *SIAM Journal on Discrete Math.* 15, 1 (2001), 41–57.
- [23] David G. Harris. 2019. Distributed local approximation algorithms for maximum matching in graphs and hypergraphs. In *Proc. 60th IEEE Symp. on Foundations of Computer Science (FOCS)*. 700–724.
- [24] Nati Linial. 1987. Distributive graph algorithms – Global solutions from local data. In *Proc. 28th IEEE Symp. on Foundations of Computer Science (FOCS)*. 331–335.
- [25] Nati Linial. 1992. Locality in Distributed Graph Algorithms. *SIAM J. Comput.* 21, 1 (1992), 193–201.
- [26] Nati Linial and Michael Saks. 1993. Low Diameter Graph Decompositions. *Combinatorica* 13, 4 (1993), 441–454.
- [27] Michael Luby. 1986. A Simple Parallel Algorithm for the Maximal Independent Set Problem. *SIAM J. Comput.* 15 (1986), 1036–1053.
- [28] Michael Luby. 1993. Removing randomness in parallel computation without a processor penalty. *J. of Computer and System Sciences* 47, 2 (1993), 250–286.
- [29] Gary L. Miller, Richard Peng, and Shen Chen Xu. 2013. Parallel graph decompositions using random shifts. In *ACM symposium on Parallelism in algorithms and architectures (SPAA)*. 196–203.
- [30] Alessandro Panconesi and Aravind Srinivasan. 1992. Improved distributed algorithms for coloring and network decomposition problems. In *Proc. 24th ACM Symp. on Theory of Computing (STOC)*. 581–592.
- [31] David Peleg. 2000. *Distributed Computing: A Locality-Sensitive Approach*. SIAM.
- [32] Václav Rozhoň and Mohsen Ghaffari. 2020. Polylogarithmic-time deterministic network decomposition and distributed derandomization. In *Proc. 52nd ACM Symp. on Theory of Computing (STOC)*. 350–363.

Received 2022-11-07; accepted 2023-02-06