

# Parallel Derandomization for Coloring

Sam Coy\*

University of Warwick

Artur Czumaj†

University of Warwick

Peter Davies-Peck‡

Durham University

Gopinath Mishra§

National University of Singapore

## Abstract

Graph coloring problems are among the most fundamental problems in parallel and distributed computing, and have been studied extensively in both settings. In this context, designing efficient *deterministic* algorithms for these problems has been found particularly challenging.

In this work we consider this challenge, and design a novel framework for derandomizing algorithms for coloring-type problems in the *Massively Parallel Computation* (MPC) model with sublinear space. We give an application of this framework by showing that a recent  $(\text{degree} + 1)$ -list coloring algorithm by Halldorsson et al. (STOC'22) in the LOCAL model of distributed computation can be translated to the MPC model and efficiently derandomized. Our algorithm runs in  $O(\log \log \log n)$  rounds, which matches the complexity of the state of the art algorithm for the  $(\Delta + 1)$ -coloring problem.

---

\*E-mail: S.Coy@warwick.ac.uk. Department of Computer Science, University of Warwick, UK. Research supported in part by the Centre for Discrete Mathematics and its Applications (DIMAP), by an EPSRC studentship, and by the Simons Foundation Award No 663281 granted to the Institute of Mathematics of the Polish Academy of Sciences for the years 2021–2023.

†E-mail: A.Czumaj@warwick.ac.uk. Department of Computer Science and Centre for Discrete Mathematics and its Applications (DIMAP), University of Warwick, UK. Research supported in part by the Centre for Discrete Mathematics and its Applications (DIMAP), by EPSRC award EP/V01305X/1, by a Weizmann-UK Making Connections Grant, by an IBM Award, and by the Simons Foundation Award No. 663281 granted to the Institute of Mathematics of the Polish Academy of Sciences for the years 2021–2023.

‡E-mail: Peter.W.Davies@durham.ac.uk. Department of Computer Science, Durham University, UK.

§E-mail: Gopinath@nus.edu.sg. Department of Computer Science, National University of Singapore. The work was done when Gopinath Mishra was a Post Doctoral Fellow at the University of Warwick. The research was supported in part by the Centre for Discrete Mathematics and its Applications (DIMAP), by EPSRC award EP/V01305X/1, and by the Simons Foundation Award No. 663281 granted to the Institute of Mathematics of the Polish Academy of Sciences for the years 2021–2023.

# 1 Introduction

The *Massively Parallel Computation* (MPC) model, introduced over a decade ago by Karloff et al. [KSV10], is a contemporary standard theoretical model for parallel algorithms. The model has evolved through the successful emulation of parallel and distributed frameworks, including but not limited to MapReduce [DG08], Hadoop [Whi15], Dryad [IBY<sup>+</sup>07], and Spark [ZCF<sup>+</sup>10]. Drawing inspiration from classical models of parallel computation (e.g., PRAM) and distributed models (e.g., *CongestedClique*), MPC incorporates numerous shared attributes with these established paradigms. We focus on the *sublinear local space* MPC regime, where machines have local space  $s = O(n^\phi)$  for any arbitrary fixed constant  $\phi \in (0, 1)$ , where  $n$  is the number of nodes in the graph. Note that  $\phi$  is a parameter of the model. This model has attracted a lot of attention recently, see, e.g., [ANOY14, ASS<sup>+</sup>18, BKM20, BKS17, BBD<sup>+</sup>19, BHH19, CFG<sup>+</sup>19, CC22, CDP21b, CDP21c, CDP21d, CLM<sup>+</sup>18, G GK<sup>+</sup>18, GKU19, GU19, GSZ11, LMOS20]. Recent works have provided algorithms for fundamental graph problems such as connectivity, approximate matching, maximal matching, maximal independent set, and  $(\Delta + 1)$  coloring.

While the main focus on MPC algorithms has been typically concentrating on the design of randomized algorithms, an area which has seen a recent surge in interest in the MPC model is the design of *deterministic algorithms* which are as efficient as their randomized counterparts. In related models of distributed computation this task is sometimes impossible: there are lower bounds in the LOCAL model which separate the randomized complexity from the deterministic complexity of several fundamental algorithmic problems (see, e.g., [CKP19]). However, a combination of the global co-ordination and of the computational power of individual machines within the MPC model allows for a broader selection of tools for derandomizing MPC algorithms. For example, a particularly effective tool in this context is the *method of conditional expectations*, which has been recently used to derandomize several algorithms for fundamental problems in the MPC model [BKM20, CC22, CDP21c, FGG22].

In this paper we develop a general derandomization framework, providing a useful tool for translating some class of randomized LOCAL algorithms to deterministic MPC in a black-box manner. Previous works have either been heavily tailored to specific algorithms [BKM20, CC22, CDP21c, FGG22] or have had severe limitations on graph degree [CDP21a]. Our approach allows generic derandomization of a large class of algorithms over a much larger degree range.

As an application of our derandomization framework, we consider the task of designing fast deterministic MPC algorithms for graph coloring — one of the most fundamental algorithmic primitives, extensively studied in various settings for several decades. Graph coloring problems have been playing a prominent role in distributed and parallel computing, not only because of their numerous applications, but also since some variants of coloring problems naturally model typical symmetry breaking problems, as frequently encountered in decentralized systems (see, e.g., [BE13] for an overview of early advances). Parallel graph coloring has been studied since the 1980s [BE13, Kar85], and nowadays  $(\Delta + 1)$ -coloring and  $(2\Delta - 1)$ -edge-coloring<sup>1</sup> are considered among the most fundamental graph problems — benchmark problems in the area (throughout the paper,  $\Delta$  refers to the maximum degree of the input graph).

In particular, we study the parallel complexity of a natural generalization of the  $(\Delta + 1)$ -coloring problem, the problem of *(degree+1)-list coloring* (D1LC). In the D1LC problem, for a given graph  $G = (V, E)$ , each node has an input palette of acceptable colors of size one more than its degree, and the objective is to find a proper coloring using these palettes. The recent increasing interest in D1LC has been largely caused by the generality and applicability of D1LC, since, for example, given a partial solution to a  $(\Delta + 1)$ -coloring problem, the remaining coloring problem on the uncolored nodes is an instance of D1LC. It also appears as a subproblem in more constrained coloring problems, e.g., as a subroutine in distributed  $\Delta$ -coloring algorithms (see, e.g., [FHM23]) and in edge-coloring algorithms (see, e.g., [Kuh20]).

---

<sup>1</sup> $(\Delta + 1)$ -coloring problem is to color a graph of maximum degree  $\Delta$  using  $\Delta + 1$  colors;  $(2\Delta - 1)$ -edge-coloring problem is to color the edges of a graph of maximum degree  $\Delta$  using  $2\Delta - 1$  colors.

## 1.1 Our contribution

In this paper, we develop a framework for generic derandomization of LOCAL algorithms in the deterministic low-space MPC model. This framework uses pseudorandom generators and the method of conditional expectations, which are known techniques, but significantly extends prior work in generality and applicability.

Our main derandomization theorem is Theorem 12. Since this theorem involves some technical definitions, we leave the statement to Section 4.3. However, informally, the result shows that any randomized LOCAL algorithm that can be decomposed into short subprocedures with some natural properties can be efficiently derandomized in MPC.

As an application of our framework, we show that D1LC can be solved efficiently deterministically in the *Massively Parallel Computation* (MPC) model with sublinear local space, matching the complexity of the state-of-the-art MPC algorithms for the simpler  $(\Delta + 1)$ -coloring and  $(\Delta + 1)$ -list coloring problems. We first show how to combine the D1LC framework for the LOCAL model due to Halldórsson et al. [HKNT22] with the techniques developed in earlier works on the MPC model, to obtain a randomized MPC algorithm for D1LC working in  $O(\log \log \log n)$  rounds, w.h.p. Then we apply our novel derandomization framework to derandomize this algorithm.

**Theorem 1 (D1LC).** *Let  $\phi \in (0, 1)$  be an arbitrary constant. There exists a deterministic algorithm that, for every  $n$ -node graph  $G = (V, E)$ , solves the D1LC problem using  $O(\log \log \log n)$  rounds, in the sublinear local space MPC model with local space  $s = O(n^\phi)$  and global space  $O(m + n^{1+\phi})$ .*

Observe that the bound in Theorem 1 matches the state-of-the-art bound for the complexity of the simpler  $(\Delta + 1)$ -coloring problem in the (sublinear local space) MPC model (see [CFG<sup>+</sup>19] for the randomized bound and [CDP21c] for the deterministic bound). Furthermore, the recently developed framework connecting the complexity of LOCAL and sublinear local space MPC algorithms (see [CDP21a, GKU19]), provides some evidence that our upper bound cannot be asymptotically improved, unless the complexity of the  $(\Delta + 1)$ -coloring problem is  $(\log \log n)^{o(1)}$  in the LOCAL model. This is because [CDP21a, GKU19] show that for a class of component stable algorithms and conditioned on the so-called 1-vs-2-cycles conjecture, no sublinear local space MPC algorithm can run faster than the logarithm of the complexity of LOCAL algorithms. (Still, even conditioned on the 1-vs-2-cycles conjecture, it might be conceivable that a non-component stable randomized MPC sublinear local space algorithm can solve  $(\Delta + 1)$ -coloring in  $o(\log \log \log n)$  rounds, and further, we do not have any good enough LOCAL lower bounds for coloring, and so maybe an  $(\log \log n)^{o(1)}$ -rounds LOCAL algorithm is possible.) Finally, notice that Roughgarden et al. [RVW18] showed that proving any super-constant lower bound in the sublinear local space MPC for any problem in P would separate  $\text{NC}^1$  from P, making any *unconditional* super-constant (strongly sublinear local space) MPC lower bound unlikely.

## 1.2 Technical overview

A natural approach to design efficient MPC algorithms is to simulate a LOCAL algorithm using the so-called graph exponentiation approach (see, e.g., [ASS<sup>+</sup>18, Gha16, GU19, LW10]): in the first round each node learns its 2-hop neighborhood, then its 4-hop neighborhood, and so on. In some cases we can even obtain deterministic algorithms in low-space MPC which are exponentially faster than the corresponding *randomized* complexity in LOCAL [CDP21a]. However, doing so is often not easy, and faces two main challenges: the first is how to efficiently derandomize algorithms, and the second is that large neighborhoods may not fit onto machines for high-degree instances, which renders challenging many common subroutines in LOCAL algorithms (e.g., in the coloring setting, computing an almost-clique decomposition).

We address the first problem, of derandomization, by developing a general derandomization framework for transferring some class of randomized LOCAL algorithms to deterministic low-space MPC ones. Our framework relies on the use of pseudorandom generators (PRGs) combined with the method of conditional expectations. The approach of using PRGs in derandomization has been used in the past (see, e.g., [CDP21a, CDP21c]), but faces a major difficulty in general application. This difficulty is that the stringent constraint of the local space requirement of sublinear local space MPC causes PRGs to fail on a non-trivial proportion of nodes, even if the base randomized procedure succeeds with high probability. Furthermore, attempts to defer

or change the outputs of these failed nodes can cause a chain reaction of nodes failing to meet the success requirements of the randomized procedure. To derandomize  $(\Delta + 1)$ -coloring, [CDP21c] addressed this issue by painstakingly analyzing the effects of PRG failures throughout the course of the (highly complex) base randomized algorithm, but this was highly tailored to the specific algorithm and not easily generalizable. [CDP21a] did provide a form of general framework, but this required collecting the neighborhoods needed for entire LOCAL algorithms onto machines, and therefore only worked for low values of  $\Delta$ . These low- $\Delta$  examples sufficed for [CDP21a] since the aim was to show that component-stable lower bounds could be surpassed in some cases, but did not aid derandomization of general algorithms.

We overcome these difficulties by providing a suitable framework that formalizes (see Definition 5) a collection of properties that allow us to fully derandomize randomized algorithm using the PRGs. The main technical difference from [CDP21a] is that rather than derandomizing full algorithms in one go, we instead decompose them into short subprocedures with certain useful properties. We then have the space necessary to derandomize these subprocedures on individual machines, allowing us to tackle much higher-degree instances.

The framework and its analysis summarized in Theorem 12 form the main technical contribution of our work. We hope that our framework, and in particular, the key lemma (Theorem 12), will prove useful as a powerful black-box derandomization technique in MPC.

We apply this framework to D1LC as our main application. However, for instances with very high degree, we must still reduce the degree before we can fit even 1-hop neighborhoods onto machines. To do so, we employ a *deterministic recursive sparsification* approach similar to [CDP21c, CDP21d], where we repeatedly partition an instance of D1LC with maximum degree  $\Delta$  into  $n^\delta$  D1LC instances, each with maximum degree  $\Delta/n^\delta$ . Here  $\delta$  with  $0 \leq \delta \leq \phi$  is to be fixed later, and  $\phi$  is our local space parameter, i.e.,  $s = O(n^\phi)$ . All but one of these instances are *valid* D1LC instances and so can be solved using this recursive sparsification if the degree is still too high, and the final instance can only be solved when it is determined which colors are unused in the other instances. In this way, we can reduce the maximum degree of the D1LC instances which we have to solve to an arbitrarily small polynomial in  $n$ , which is small enough to apply the derandomization framework.

### 1.3 Related work

#### Parallel derandomization.

Our work relies on some sparsification and derandomization techniques developed for parallel and distributed coloring algorithms. Derandomization techniques such as the method of conditional expectations are long-studied and well-understood in classical sequential computation, and were first employed in linear-space MPC (actually, in the mostly equivalent *CongestedClique* model) by [CHPS20]. This was followed by a work extending the techniques to low-space MPC [CDP21b]. These works both applied the method of conditional expectations to families of bounded-independence hash functions, a technique for reducing the length of random seeds required for randomized algorithms. However, they are only applicable to algorithms that do not heavily exploit the independence of random choices.

Many algorithms (including all known sublogarithmic LOCAL coloring algorithms) do not appear to have this property, and their analyses effectively use  $\Delta$ -wise independence or higher, which is too high to efficiently apply families of bounded-independence hash functions. So, some subsequent works on parallel derandomization instead employ pseudorandom generators (PRGs) to reduce the seed space. PRGs are again a well-studied topic in classical computing (see, e.g., [Vad12] for an introduction), and do not have the restriction on independence. The drawback is that PRGs achieving optimal parameters are only known existentially, and computing them requires exponential-time computation — but this is generally permitted in MPC, and even if not, the PRG can be pre-computed sequentially and hard-coded onto MPC machines. [CDP21c] used PRG-based derandomization to give an  $O(\log \log \log n)$ -round deterministic low-space MPC algorithm for  $\Delta + 1$ -coloring, and [CDP21a] used it for general derandomization, in order to demonstrate that the method of conditional expectations could be exploited to surpass the component-stable lower bounds of [GKU19]. However, as mentioned, the derandomization results therein were only for low-degree instances.

Implementations of the method of conditional expectations have also recently been used for derandomization in related distributed models, see e.g., [BKM20, GHK18, Kuh20].

### Parallel and distributed coloring.

Our application to D1LC continues a long line of research studying the parallel and distributed complexity of graph coloring problems. For the references to earlier work on distributed coloring algorithms we refer to the monograph by Barenboim and Elkin [BE13] (see also the influential papers by Linial [Lin87, Lin92]). We will discuss here only more recent advances (and final results) for the four most relevant coloring problems,  $(\Delta + 1)$ -coloring,  $(\Delta + 1)$ -list-coloring, D1LC, and  $\Delta$ -coloring. After extensive research in the field of distributed computing concerning the  $(\Delta + 1)$ -coloring problem, we understand its complexity for the LOCAL, CongestedClique, and also for the MPC model. For CongestedClique (and also for MPC with linear memory,  $s = O(n)$ ), we now know how to solve  $(\Delta + 1)$ -coloring in a constant number of rounds, see [CLP18, CDP21d]. For the LOCAL model, after a very long line of research, the current state of the art upper bound for randomized algorithm is  $\tilde{O}(\log^2 \log n)$  [CLP18, RG20].<sup>2 3</sup> The best known deterministic algorithm is of  $\tilde{O}(\log^2 n)$  rounds [GG23] in the LOCAL model.

For the sublinear local space MPC, it is known that the  $(\Delta + 1)$ -coloring algorithm due to Chang et al. [CFG<sup>+</sup>19] can be combined with the network decomposition result of [RG20] to obtain a randomized  $O(\log \log \log n)$ -round MPC algorithm, which is currently the state-of-the-art for the problem; this result was derandomized by Czumaj et al. [CDP21c]. Furthermore, all algorithms mentioned above for  $(\Delta + 1)$ -coloring can be extended to solve also  $(\Delta + 1)$ -list-coloring.

For the D1LC problem, which is a generalization of  $(\Delta + 1)$ -coloring and  $(\Delta + 1)$ -list-coloring, there have not been many comparable bounds until the very recent work of Halldórsson et al. [HKNT22]. In particular, D1LC admits a randomized  $\tilde{O}(\log^2 \log n)$ -round distributed algorithm in the LOCAL model [HKNT22, GG23] matching the state-of-the-art complexity for the  $(\Delta + 1)$ -coloring problem [CLP18, GG23]. In [HKNT22], the authors significantly extended the approaches for  $(\Delta + 1)$ -coloring (in particular, to allow efficient management of nodes of various degrees). As a byproduct, the framework of Halldórsson et al. [HKNT22] can be incorporated into a constant-round MPC algorithm assuming the local MPC space is slightly *superlinear*, i.e.,  $O(n \log^4 n)$  [HKNT22, Corollary 2]. We make extensive use of the framework laid out by Halldórsson et al. [HKNT22] in their algorithm for LOCAL in the design of our D1LC algorithm. A similar approach has been applied recently for the CONGEST model in [HNT22, GK21], solving D1LC in  $\tilde{O}(\log^3 \log n)$  CONGEST rounds, w.h.p. Very recently, a deterministic constant-round algorithm for D1LC has been obtained for the CongestedClique model [CCDM23], settling the complexity of D1LC for CongestedClique.

## 2 Preliminaries

### 2.1 Notations, the problem, and the model

For  $k \in \mathbb{N}$ ,  $[k]$  denotes the set  $\{1, \dots, k\}$ . For  $a, b \in \mathbb{N}$ ,  $[a, b]$  denotes the set of integers in  $\{a, a + 1, \dots, b\}$ . We consider a graph  $G = (V, E)$  with  $V$  as the node set and  $E$  as the edge set with  $|V| = n$  and  $|E| = m$ . The set of neighbors of a node  $v$  is denoted by  $N(v)$  and the degree of a node  $v$  is denoted by  $d(v)$ . For a node  $v$ ,  $\Psi(v)$  denotes the list of colors in the color palette of node  $v$  and  $p(v)$  denotes the size of  $\Psi(v)$ . The maximum degree of any node in  $G$  is denoted by  $\Delta$ . As we go on coloring the nodes of the graph  $G$ , the graph will change and the color palettes of the nodes will also change. Often, we denote the current (rather than the input) graph by  $G$ . For all graphs we consider, we have  $p(v) \geq d(v) + 1$ . For a subset  $X \subseteq V$ ,  $G[X]$  denotes the subgraph induced by  $X$  and  $m(X)$  denotes the number of edges in  $G[X]$ .

---

<sup>2</sup> $\tilde{O}(f)$  hides a polynomial term in  $\log f$ .

<sup>3</sup>Through out this paper,  $\log^k x$  denotes  $(\log x)^k$ .

### Degree+1 list coloring (D1LC).

The *degree+1 list coloring (D1LC) problem* is for a given graph  $G = (V, E)$  and given color palettes  $\Psi(u)$  assigned to each node  $u \in V$ , such that  $|\Psi(u)| \geq d(u) + 1$ , the objective to find a proper coloring of nodes in  $G$  such that each node is assigned to a color from its color palette and no edge in  $G$  is monochromatic.

### Massively Parallel Computation model.

We consider the *Massively Parallel Computation (MPC)* model, which is a parallel system with some number of machines, each of them having some local space  $s$ . At the beginning of computation, each machine receives some part of the input, with the constraint that it must fit within its local space. In our case, for the D1LC problem, the input is a set of  $n$  nodes,  $m$  edges, and  $n$  color palettes of total size  $O(n + m)$ . Hence we will require that the number of machines is  $\Omega(\frac{n+m}{s})$ , for otherwise the input would not fit the system. The computation on an MPC proceeds in synchronous rounds. In each round, each machine processes its local data and performs an arbitrary local computation on its data without communicating with other machines. At the end of each round, machines can exchange messages. Each message is sent only to a single machine specified by the machine that is sending the message. All messages sent and received by each machine in each round have to fit into the machine's local space. Hence, the total number of messages sent by any machine and received by any machine is bounded by  $s$ , and the total amount of communication across the whole MPC is bounded by  $s$  times the number of machines. At the beginning of the next round, each machine can process all messages received in the previous round. When the algorithm terminates, machines collectively output the solution.

Observe that if a single machine can store the entire input, then any problem (like, e.g., D1LC) can be solved in a single round, since no communication is required. In order for our algorithms to be as scalable as possible, normally one wants to consider graph problems in the *sublinear local space regime*, where local space  $s = n^\phi$  for any given constant  $\phi \in (0, 1)$ . (There has been some research considering also the case when  $s = \Theta(n)$ , or even when  $s = n^{1+\phi}$  (in which case one wants to study the case that  $s \ll m$ ) but we will not consider such setting in the current paper.) We will require that the number of machines is not significantly more than required, specifically that it is  $\tilde{O}(n + \frac{m}{s})$  (note that the optimal amount would be  $\tilde{O}(\frac{n+m}{s})$ , but our algorithm requires the ability to assign a machine to each node). A major challenge in the design of MPC algorithms in the sublinear local space regime is that the local space of each machine is (possibly) not sufficient to store all the edges incident to a single node. This constraint naturally requires an MPC algorithm to rely on extensive communication between machines, and most of the techniques known are based on some graph sparsification. It is important to note here that even in the sublinear local space regime, the MPC model is known [GSZ11] to be stronger than the PRAM model, e.g., it is known that sorting<sup>4</sup> (and in fact, many related tasks, like prefix sum computation) can be performed in a constant number of rounds, even deterministically, see [GSZ11]. Observe that with this tool, we can gather nodes' neighborhoods to contiguous blocks of machines, and learn their degrees, in  $O(1)$  rounds, and that we can assume, without loss of generality, that the input can be distributed arbitrarily on the first  $\Theta(\frac{n+m}{s})$  MPC machines.

## 2.2 An overview of LOCAL D1LC algorithm of Halldórsson et al. [HKNT22]

In this section, we give a detailed overview of the work of [HKNT22]. Among several other contributions, [HKNT22] give an algorithm which, in the LOCAL model of distributed computing, given an input instance of  $(\deg + 1)$ -list coloring for a graph  $G$ , colors all nodes of  $G$  having degrees at least  $\log^7 n$  in  $O(\log^* n)$  rounds. We give an overview of their algorithm here, because our algorithm in MPC to color all nodes in the degree range  $[\log^7 n, n^{7\delta}]$  uses steps of the algorithm in [HKNT22], where  $\delta$  is set to be sufficiently smaller than  $\phi$  (recall that  $\phi$  is our local space parameter, i.e.,  $s = O(n^\phi)$ ).

The LOCAL algorithm for D1LC due to Halldórsson et al. [HKNT22] handles the input graph in “ranges” of degree. It begins by coloring vertices with degrees in the range  $[\log^7 n, n]$ , followed by vertices with degrees in the range  $[\log^7 \log n, \log^7 n]$  vertices in the range  $[\log^7 \log \log n, \log^7 \log n]$ , and so on, giving  $O(\log^* n)$  ranges

---

<sup>4</sup>Here we consider sorting of  $N$  objects on an MPC with local space  $N^\gamma$  and on  $N^{1-\gamma}$  machines, for any constant  $\gamma > 0$ .

overall. The coloring algorithm runs in  $O(\log^* n)$  rounds for each range. However, for ranges after the first the algorithm does not color all nodes with high probability, leaving a set of “bad” nodes. Some standard “shattering” arguments are used (showing that these bad vertices form small components), and then these bad vertices can be colored using a deterministic D1LC algorithm. The post-shattering coloring step on the second degree range ( $[\log^7 n, \log^7 n]$ ) is the bottleneck in the algorithm, and gives the overall complexity of  $\tilde{O}(\log^2 \log n)$  LOCAL rounds.<sup>5</sup>

The algorithm of [HKNT22] (for the degree range  $[\log^7 n, n]$ ) requires several careful arguments which rely on nodes being able to calculate certain parameters, in particular the following:

**Definition 2** (Parameters from [HKNT22]). *The following node-parameters are used in the coloring algorithm of [HKNT22]:*

- Slack  $s(v) = p(v) - d(v)$ : *The difference between the size of the palette of a node and its degree.*
- Sparsity  $\zeta_v = \frac{1}{d(v)} \cdot \left[ \binom{d(v)}{2} - m(N(v)) \right]$ <sup>6</sup>: *The proportion of non-edges between neighbors of  $v$ .*
- Disparity  $\bar{\eta}_{u,v} = |\Psi(u) \setminus \Psi(v)| / |\Psi(u)|$ : *The (proportional) difference between the palettes of  $u$  and  $v$ .*
- Discrepancy  $\bar{\eta}_v = \sum_{u \in N(v)} \bar{\eta}_{u,v}$ : *The sum of disparities between  $v$  and its neighbors.*
- Unevenness  $\eta_v = \sum_{u \in N(v)} \frac{\max(0, d(u) - d(v))}{d(u) + 1}$ : *A quantity that relates to how many of the neighbors of  $v$  have much higher degree than  $v$ .*
- Slackability  $\bar{\sigma}_v = \bar{\eta}_v + \zeta_v$ , Strong Slackability  $\sigma_v = \eta_v + \zeta_v$ : *How easy it is to create slack for a node.*

The coloring algorithm [HKNT22] (for a single degree range) is outlined in Algorithm 1. As described, the algorithm first computes *almost-clique decomposition* (ACD). After computing ACD, each vertex of the graph will be classified as either as a sparse vertex or a uneven vertex or a dense vertex. Then the we call COLORSPARSE to color the vertices that are either sparse or uneven, and then we call COLORDENSE to color the dense vertices. In Section 2.2.1, we discuss how to compute ACD and classify the vertices into sparse and dense vertices. In Section 2.2.2 and Section 2.2.3, we discuss how to color the sparse and dense vertices respectively. There are additional subroutines, called by COLORSPARSE and COLORDENSE, that we will discuss from Section 2.2.1–Section 2.2.3.

---

**Algorithm 1:** COLORMIDDLE: Colors nodes with degrees in  $[\log^7 n, n^{7\delta}]$  (extracted from [HKNT22])

---

- 1 Compute an almost-clique decomposition (ACD).
  - 2 COLORSPARSE (Algorithm 5).
  - 3 COLORDENSE (Algorithm 7).
- 

### 2.2.1 Computing an almost-clique decomposition

First, as already mentioned, the authors of [HKNT22] compute an *almost-clique decomposition* of the input graph. The almost-clique decomposition is as follows:

**Definition 3** ((deg+1)-ACD [AA20, HKNT22]). *Let  $G = (V, E)$  be a graph and  $\varepsilon_{ac}, \varepsilon_{sp} \in (0, 1)$  be constants. A partition of  $V$  into  $V_{sparse} \sqcup V_{uneven} \sqcup V_{dense}$ , such that  $V_{dense}$  can be partitioned into  $C_1, \dots, C_t$  for some  $t$ , is said to be an almost-clique decomposition for  $G$  if*

- (i) *Every  $v \in V_{sparse}$  is  $\varepsilon_{sp}d(v)$ -sparse;*
- (ii) *Every  $v \in V_{uneven}$  is  $\varepsilon_{sp}d(v)$ -uneven;*

---

<sup>5</sup>The actual bound obtained by [HKNT22] is  $O(\log^3 \log n)$ , and when this is combined with the result in [GG23], one gets the round complexity to be  $\tilde{O}(\log^2 \log n)$ .

<sup>6</sup> $m(N(v))$  denotes the number of edges in the subgraph induced by  $N(v)$ .

- (iii) For every  $i \in [t]$  and  $v \in C_i$ ,  $d(v) \leq (1 + \varepsilon_{ac}) |C_i|$ ;
- (iv) For every  $i \in [t]$  and  $v \in C_i$ ,  $|C_i| \leq (1 + \varepsilon_{ac}) |N(v) \cap C_i|$ .

Intuitively:

- $V_{\text{sparse}}$ : A set of (sufficiently) *sparse* nodes. A node is  $\varepsilon$ -sparse if there are  $\varepsilon \cdot d(v)$  many non-edges between neighbors of  $v$ .
- $V_{\text{uneven}}$ : A set of *uneven* nodes. A node is uneven if many of its neighbors have a much higher degree than  $v$ . Formally, a node  $v$  is  $\eta$ -uneven if  $\eta_v \geq \sum_{u \in N(v)} \frac{\max(0, d(u) - d(v))}{d(u) + 1}$ .
- $V_{\text{dense}} = C_1 \sqcup \dots \sqcup C_t$ : Each  $C_i$  is an *almost-clique*. Intuitively, almost-cliques are very well connected parts of the graph. Formally, for some  $\varepsilon_{sc}$  and for all almost-cliques  $C_i$  and all  $v \in C_i$ ,  $d(v) \leq (1 + \varepsilon_{ac}) |C_i|$ , and  $|C_i| \leq (1 + \varepsilon_{ac}) |N(v) \cap C_i|$ .

After computing an ACD and partitioning the nodes into  $V_{\text{sparse}}$ ,  $V_{\text{uneven}}$  and  $V_{\text{dense}}$ , nodes in  $V_{\text{sparse}} \sqcup V_{\text{uneven}}$  are colored first by COLORMIDDLE (Algorithm 1) by subroutine COLORSPARSE and then the nodes in  $V_{\text{dense}} = C_1 \sqcup \dots \sqcup C_t$  are colored by subroutine COLORDENSE.

## 2.2.2 Coloring sparse nodes

Before discussing about coloring the sparse nodes, we first discuss about a subroutine SLACKCOLOR (Algorithm 2) which is used by Step 2 and 3 of Algorithm 1 by the subroutines COLORSPARSE and COLORDENSE, respectively.

### SlackColor

SLACKCOLOR colors all nodes which have slack linear in their degree (and the degree is large enough) in  $O(\log^* n)$  rounds. The authors of [HKNT22] point out that similar results were shown in [SW10]: in this paper we focus on simulating and derandomizing the algorithm SLACKCOLOR given in [HKNT22]. Note that SLACKCOLOR takes two parameter  $s_{\min}$  and  $\kappa$  as input where  $s_{\min}$  is an integer that lies between 1 and the minimum amount of slack of any node (that we want to color) and  $\kappa$  lies between  $1/s_{\min}$  and 1.<sup>7</sup>

---

**Algorithm 2:** SLACKCOLOR( $s_{\min}$ ), for node  $v$ , from [HKNT22]

---

- 1 Do TRYRANDOMCOLOR( $v$ ) for  $O(1)$  rounds.
  - 2 If  $s(v) < 2d(v)$  then terminate.
  - 3 Let  $\rho \leftarrow s_{\min}^{1/(1+\kappa)}$ . //  $s_{\min}$  is globally known such that  $1 < s_{\min} \leq \min_v s(v)$  and  $\kappa$  is an input parameter such that  $1/s_{\min} < \kappa \leq 1$ .
  - 4 **for**  $i$  from 0 to  $\log^* \rho$  **do**
  - 5      $x_i \leftarrow 2 \uparrow\uparrow i$  //  $2 \uparrow\uparrow i$  refers to iterated exponentiation, i.e.,  $2 \uparrow\uparrow 0 = 1$  and  $2 \uparrow\uparrow (i+1) = 2^{2 \uparrow\uparrow i}$ .
  - 6     MULTITRIAL( $x_i$ ) 2 times
  - 7     If  $d(v) > s(v) / \min(2^{x_i}, \rho^\kappa)$  then terminate.
  - 8 **end**
  - 9 **for**  $i$  from 1 to  $\lceil 1/\kappa \rceil$  **do**
  - 10      $x_i \leftarrow \rho^{i \cdot \kappa}$ .
  - 11     MULTITRIAL( $x_i$ ) 3 times.
  - 12     If  $d(v) > s(v) / \min(\rho^{(i+1) \cdot \kappa}, \rho)$  then terminate.
  - 13 **end**
  - 14 MULTITRIAL( $\rho$ ).
- 

SLACKCOLOR (Algorithm 2) requires two auxiliary subroutines TRYRANDOMCOLOR (Algorithm 3) and MULTITRIAL (Algorithm 4) which we describe below.

<sup>7</sup>The success probability of SLACKCOLOR depends on  $s_{\min}$  and  $\kappa$ . Also, [HKNT22] uses SLACKCOLOR only when  $s_{\min} = \Omega(\ell)$  and  $\kappa$  is a suitable constant, where  $\ell = \log^{2.1} \Delta$ .



**Auxiliary subroutine TryRandomColor** In TRYRANDOMCOLOR, each participating node  $v$  selects a color  $\psi_v$  from its current color palette uniformly at random.  $v$  colors itself with  $\psi_v$  if none of the neighbors of  $v$ , whose colors conflict with  $v$ , selects  $\psi_v$ .

---

**Algorithm 3:** TRYRANDOMCOLOR(node  $v$ ) from [HKNT22]

---

- 1 Pick  $\psi_v$  u.a.r. from  $\Psi(v)$ .
  - 2 Send  $\psi_v$  to each  $u \in N(v)$ , receive the set  $T = \{\psi_u : u \in N^+(v)\}$ , where  $N^+(v)$  is the set of neighbours whose colors “conflict” with  $v$ .
  - 3 If  $\psi_v \notin T$  then permanently color  $v$  with  $\psi_v$ .
  - 4 Send and receive permanent colors, and remove the received one from  $\Psi(v)$
- 

**Auxiliary subroutine MultiTrial** MULTITRIAL is, informally speaking, a generalized version of TRYRANDOMCOLOR where each participating node selects  $x$  number of colors from its palette uniformly at random.

---

**Algorithm 4:** MULTITRIAL( $x$ ) from [HKNT22]

---

- 1  $v$  picks a set  $X_v$  of  $x$  random colors in its palette  $\Psi(v)$ , sends them to neighbours.
  - 2 **if**  $\exists \psi \in X_v$  such that  $\forall u \in N(v), \psi \notin X_u$  **then**
  - 3     | Permanently color  $v$  with some such  $\psi$  and broadcast to  $N(v)$ .
  - 4 **end**
- 

### Coloring sparse or even nodes

The first coloring task is to color all nodes which are *sparse* or *uneven* (that are in  $V_{\text{sparse}}$  or  $V_{\text{uneven}}$ ). The outline of the subroutine which does this COLORSPARSE (Algorithm 5). This is a recursively subroutine that uses an auxiliary subroutine GENERATESLACK (Algorithm 6).

---

**Algorithm 5:** COLORSPARSE from [HKNT22]

---

- 1 Identify set of nodes  $V_{\text{start}} \subset V_{\text{sparse}}$ .
  - 2 GENERATESLACK in  $G[(V_{\text{sparse}} \cup V_{\text{uneven}}) \setminus V_{\text{start}}]$ .
  - 3 SLACKCOLOR  $V_{\text{start}}$ .
  - 4 SLACKCOLOR  $V_{\text{sparse}} \setminus V_{\text{start}}$  and  $V_{\text{uneven}}$ .
- 

**Auxiliary subroutine GenerateSlack** The challenge in COLORSPARSE (Algorithm 5) is that all nodes at the start of the algorithm are only guaranteed to have constant slack: in the (easier)  $(\Delta + 1)$ -coloring problem, nodes of low-degree are guaranteed to have slack linear in their degree. The authors of [HKNT22] overcome this by identifying some subset of the nodes in  $V_{\text{sparse}}$  for which it is hard to generate slack: they call this set  $V_{\text{start}}$ . They then generate some slack using a simple primitive (GENERATESLACK, Algorithm 6) for nodes in  $(V_{\text{sparse}} \cup V_{\text{uneven}}) \setminus V_{\text{start}}$ .

---

**Algorithm 6:** GENERATESLACK from [HKNT22]

---

- 1  $S \leftarrow$  each node  $v$  is sampled into  $S$  independently with probability  $\frac{1}{10}$ .
  - 2 For all  $v \in S$  in parallel TRYRANDOMCOLOR( $v$ ).
- 

The authors of [HKNT22] show that nodes in  $V_{\text{start}}$  have many neighbours in  $(V_{\text{sparse}} \cup V_{\text{uneven}}) \setminus V_{\text{start}}$ , and so nodes in  $V_{\text{start}}$  have considerable *temporary slack* (slack obtained from neighbors being colored later). Since  $(V_{\text{sparse}} \cup V_{\text{uneven}}) \setminus V_{\text{start}}$  are exactly the nodes which do not struggle to obtain slack, they have that all sparse/uneven nodes have slack linear in their degree, and this is enough to color them in  $O(\log^* n)$  rounds (by the algorithm of [SW10] or by SLACKCOLOR).

### 2.2.3 Coloring dense nodes

Finally, it remains to color nodes in the almost-cliques  $C_1 \sqcup \dots \sqcup C_t$ . This task is solved using algorithm COLORDENSE (Algorithm 7) [HKNT22].

---

**Algorithm 7:** COLORDENSE from [HKNT22]

---

- 1 Compute the leader  $x_C$  and outliers  $O_C$  for each almost-clique  $C$ . Let  $O = \bigcup_{C \in \mathcal{C}} O_C$ , where  $\mathcal{C}$  denotes the set of all almost-cliques.
  - 2 GENERATESLACK.
  - 3  $P_C \leftarrow \text{PUTASIDE}(C)$  for each low slack almost-clique  $C$ . Let  $P$  denote the set of union of all such  $P_C$ 's.
  - 4 SLACKCOLOR  $O$ .
  - 5 SYNCHCOLORTRIAL  $V_{\text{dense}} \setminus P$ .
  - 6 SLACKCOLOR  $V_{\text{dense}} \setminus P$ .
  - 7 For each  $C \in \mathcal{C}$  with low slackability,  $x_C$  collects the palettes in  $P_C$  and colors the nodes locally.
- 

First, each almost-clique  $C$  selects a *leader*  $x_C$ , i.e., the node with minimum slackability (breaking ties arbitrarily).<sup>8</sup> The slackability of an almost-clique  $C$  is the slackability of the leader  $x_C$ .  $C$  is said to an almost-clique with low slackability if its slackability is at most  $\ell = \log^{2.1} \Delta$ ; otherwise  $C$  has high slackability. For each almost-clique with low slackability, we need to remove some of the nodes to help the coloring process. The details will be discussed later in this section.

The intuition is that the leader has similar palettes to a large fraction of other nodes in the almost-clique, a property which will be useful later. Then the almost-clique is split roughly in half, into *inliers* and *outliers*: inliers are the nodes that are similar to the leader (by some criteria that we will discuss later), and outliers are nodes which are dissimilar to the leader. In particular, since outliers are (slightly less than) half of the almost-clique, they have slack linear in their degree and can be colored in  $O(\log^* n)$  rounds by SLACKCOLOR (if we defer the nodes in the inliers to be colored later).

Next, inliers need to be colored. Since the leader is very similar to the inliers (in particular: it is connected to all inliers and has a similar palette to them), a surprisingly simple process, where the leader proposes one of its colors to each inlier (such that it proposes each color to at most one inlier), suffices. The pseudocode for this process is given in SYNCHCOLORTRIAL (Algorithm 8). After SYNCHCOLORTRIAL has completed, there is linear slack at each remaining inlier to apply SLACKCOLOR and color the remaining inliers in  $O(\log^* n)$  rounds.

---

**Algorithm 8:** SYNCHCOLORTRIAL for almost-clique  $C$  [HKNT22]

---

- 1  $x_C$  randomly permutes its palette  $\Psi(x_C)$ , sends each neighbor  $v \in I_C$  a distinct color  $c(v)$ . //  
 $I_C = C \setminus O_C$  denotes the set of inliers of the almost-clique  $C$ .
  - 2 Send  $c(u)$  to  $N(v)$ , and receive the set  $T = \{c(u) : u \in N^+(v)\}$ . //  $N^+(v)$  is the set of neighbours whose colors "conflict" with  $v$ .
  - 3 If  $c(v) \notin T$ , then permanently color  $v$  with  $c(v)$ .
  - 4 Send and receive permanent colors, and remove the received one from  $\Psi(v)$ .
- 

Finally, there is an additional challenge to overcome for some almost-cliques. If the "slackability" (intuition: a measure of easy it is to generate slack) is low, SYNCHCOLORTRIAL might not create enough slack for the remaining inliers. The authors of [HKNT22] get around this by computing a "put-aside" set: a set of inliers which we set aside and color at the end of the algorithm. The put-aside set is computed using PUTASIDE (Algorithm 9), and has three important properties: it is of polylogarithmic size (in each almost-clique that needs it); put-aside sets in different almost-cliques have no edges between them; and the inliers not assigned to the put-aside set each have sufficient neighbours in the put-aside sets which are to create sufficient slack to be colored ultrafast by SLACKCOLOR.

---

<sup>8</sup>See Definition 2 for the definition of slackability of a node.

---

**Algorithm 9:** PUTASIDE( $C$ ) from [HKNT22]

---

- 1 Sample each node in  $v \in I_C$  independently with probability  $p_s = \frac{\ell^2}{48\Delta_C}$ , and include in the set  $S_C$ , where  $\ell = \log^{2.1} \Delta$ .
  - 2 Let  $S$  denote the union of such  $S_C$ 's.
  - 3 return  $P_C$  as the set of nodes in  $S_C$  that does not have a neighbor in  $S$ .
- 

### 3 Randomized D1LC Algorithm in MPC

In this section, we argue that if the maximum degree of the input instance is not too large, then, when combined with a result of Czumaj et al. [CDP21c], the LOCAL algorithm, for D1LC due to Halldórsson et al. [HKNT22], can be efficiently implemented in exponentially fewer rounds in the MPC model. Since this *randomized* implementation is rather straightforward, we will only sketch it here. However, it is not at all clear that the algorithm of Halldórsson et al. [HKNT22] can be efficiently derandomized in MPC. The main contribution of this paper is a derandomization framework (presented in Section 4) and its use for a deterministic MPC algorithm for D1LC.

Recall that the LOCAL algorithm for D1LC due to Halldórsson et al. [HKNT22] handles the input graph in “ranges” of degree. It begins by coloring vertices with degrees in the range  $[\log^7 n, n]$ , followed by vertices with degrees in the range  $[\log^7 \log n, \log^7 n]$  vertices in the range  $[\log^7 \log \log n, \log^7 \log n]$ , and so on, giving  $O(\log^* n)$  ranges overall. Here, we consider the implementation of the (randomized) algorithm of [HKNT22] on a single range of degrees in the low-space MPC model. We will only sketch the approach here (since this is not our main result): we give this partial result to build intuition for which aspects of Theorem 1 are challenging to obtain.

Inspecting the algorithm presented in [HKNT22] unveils that, in the worst case, each of its steps either require a node  $v$  to send messages to all its neighbors (of size  $O(\deg(v))$ ), or to calculate some value or send some messages based on the full content of the 2-hop neighborhood of  $v$ . This requires  $s \geq \Delta^2$  (or equivalently, for a fixed  $s$  we need to have  $\Delta \leq \sqrt{s}$ ), since a single machine has to be able to store  $\deg(v)$  messages of size  $O(\deg(v))$ , and the 2-hop neighborhood of any vertex. Note also that the global space required is  $O(m + n^{1+\phi})$ .

If we have this amount of space relative to our maximum degree, then we can implement 1 round of the LOCAL algorithm of [HKNT22] in  $O(1)$  rounds of low-space MPC. Like the LOCAL algorithm, our MPC implementation would succeed with high probability for all nodes with degrees in  $[\log^7 n, \sqrt{s}]$ .

Also, as mentioned earlier, there is already an algorithm in low-space MPC (see, e.g., Czumaj et al. [CDP21c]) which colors D1LC instances with polylogarithmic maximum degree in  $O(\log \log \log n)$  rounds deterministically. Combining these two observations yields the following:

**Lemma 4.** *Let  $\phi \in (0, 1)$  be an arbitrary constant. There exists an algorithm that, for every  $n$ -node graph  $G = (V, E)$  with maximum degree  $\Delta = O(\sqrt{n^\phi})$ , solves the D1LC problem using  $O(\log \log \log n)$  rounds with high probability, in the sublinear local space MPC model with local space  $O(n^\phi)$  and global space  $O(m + n^{1+\phi})$ .*

Two enhancements are required to get from Lemma 4 to our main result, Theorem 1. In Section 4, we give our framework for derandomizing coloring algorithms in MPC, and in Section 5 we show that we can derandomize the implementation of the algorithm of [HKNT22] for the degree range  $[\log^7 n, s^c]$  for some suitable constant  $c \in (0, 1/2)$ . Finally, in Section 6 we show that we can, *deterministically*, reduce any input instance to a collection of instances with lower degree. We ensure that the palettes of these instances are mostly distinct, so that only a constant number of instances need to be colored sequentially.

### 4 Framework to Derandomize MPC Algorithms

In this section, we present our black-box derandomization technique for coloring problems in MPC. Our framework relies on a collection of suitable properties (normal distributed procedures, see Definition 5) that

allow us to fully derandomize randomized algorithm using a combination of pseudorandom generators (PRGs) and the method of conditional expectations.

## 4.1 Normal distributed procedures

Our framework relies on the notion of normal  $(\tau, \Delta)$ -round randomized distributed procedure. This notion is introduced to combat the issue that a PRG that fits on a machine in low-space MPC causes more nodes to fail than the underlying random process it is applied to. We wish to defer these failed nodes to deal with them later, but in some procedures this could cause a chain reaction of failures and result in an unsolvable instance. So, Definition 5 captures those procedures for which we will show we can safely defer failed nodes.

Definition 5 is quite technical, but is conceptually fairly simple, so we first explain its meaning before giving the formal notation. In general, the randomized distributed procedures we are interested in are algorithms that run for some number of rounds in a distributed model (in our case, LOCAL), and for which it is shown that upon termination, a particular desirable property (which we call the “strong success property”) holds for all nodes with high probability. In coloring problems, this property could be, for example, that all nodes have sufficient *slack* (difference between remaining palette size and number of uncolored neighbors). More generally, this could include properties such as having low remaining degree (i.e. few neighbors that have not yet terminated and given their final output for the graph problem of interest).

The defining property of a *normal* distributed procedure is that we are able to define a “weak success property” that is still sufficient for the overall algorithm to proceed, and is implied by the strong success property *even if an arbitrary subset of nodes are deferred, changing their outputs from the procedure to the special DEFER marker*. That is, if we were to run the procedure, causing all nodes to satisfy the strong success property (which will happen with high probability), and then an adversary were able to defer any subset of nodes, effectively nullifying their outputs, we would still be sure to satisfy the weak success property at all nodes. In this case, deferring does not substantially disrupt the overall algorithm, since if we continue to repeat the procedure on the deferred nodes, we will eventually reach a state in which all nodes have the required property to proceed to the next step.

This may seem like a very strong requirement for a randomized distributed procedure: indeed, it can be seen as a special case of the distributed Lovász Local Lemma, and much research has been devoted to the more general case where deferring nodes can cause their neighbors to no longer satisfy any success property (see e.g. [GHK18, Dav23]). However, as we will see, it applies to many useful procedures, and particularly those for coloring problems, where deferring nodes is almost always helpful and provides slack to neighbors.

**Definition 5.** A *normal*  $(\tau, \Delta)$ -round distributed procedure running on a graph  $G$ , of maximum degree at most  $\Delta$ , is a procedure in the randomized LOCAL model satisfying the following criteria:

- The procedure takes  $\tau$  rounds of LOCAL.
- At the beginning of the procedure, nodes  $v$  have  $O(\Delta^\tau)$ -word sets of input information  $\text{IN}_v$  associated with them.
- During the procedure, nodes only use information from their  $\tau$ -hop neighborhood (i.e., from  $\text{IN}_v \cup \bigcup_{u \in N^\tau(v)} \text{IN}_u$ ) and  $O(\Delta^{2\tau})$  random bits, and perform  $O(\Delta^{8\tau})$  computation.<sup>9</sup>
- The output of the procedure is a new  $O(\Delta^\tau)$ -word output information  $\text{OUT}_v$  for each node, from a set of possible outputs  $\text{OUT}$ .
- The procedure has a “**strong success property**” (computable with  $O(\Delta^{8\tau})$  computation) that determines whether it has been successful for a particular node, based on the output information of that node’s  $\tau$ -hop neighborhood (formally,  $\text{SSP}_v : \text{OUT}^{N^\tau(v)} \rightarrow \{T, F\}$ ).

---

<sup>9</sup> $N^\tau(v)$  denotes the set of vertices that are within  $\tau$ -hop neighborhood of  $v$ .

- At the end of the procedure, for any node  $v$ ,

$$\Pr \left[ \text{SSP}_v \left( \prod_{u \in N^\tau(v)} \text{OUT}_u \right) \right] \geq 1 - \frac{1}{2n} .$$

- The procedure also has a “**weak success property**” (also computable with  $O(\Delta^{8\tau})$  computation) that extends the output domain of each node to include a special DEFER marker (not to be conferred by the procedure itself) indicating that that node will be deferred until the end of the derandomization (formally,  $\text{WSP}_v : \{\text{OUT} \cup \text{DEFER}\}^{N^\tau(v)} \rightarrow \{T, F\}$ ).
- Denote the set of nodes that do not satisfy the strong success property as  $\overline{\text{SSP}}$ . The success properties are such that if a node  $v$  satisfies  $\text{SSP}_v$ , and nodes in  $\overline{\text{SSP}}$  are deferred,  $v$  still satisfies the weak success property. Formally:

$$\text{SSP}_v \left( \prod_{u \in N^\tau(v)} \text{OUT}_u \right) \implies \text{WSP}_v \left( \prod_{u \in N^\tau(v) \setminus \overline{\text{SSP}}} \text{OUT}_u \times \prod_{u \in N^\tau(v) \cap \overline{\text{SSP}}} \text{DEFER} \right) .$$

Let us now discuss how Definition 5 can be applied. The intuition behind Definition 5 is that it captures procedures whose output properties are not damaged too much by unsuccessful nodes deferring. As an example, consider Luby’s randomized algorithm for maximal independent set [Lub86]. Luby’s algorithm, in  $O(\log n)$  rounds of LOCAL, produces an output set that is certainly independent, and is maximal with high probability. We can define both success properties (strong and weak) for a node  $v$ , to be that  $v$  has a node within distance 1 in the resulting output set. Note that this only captures maximality, not independence, but that is sufficient since independence is guaranteed by the process, and a “failed” run only violates maximality. A standard analysis of Luby’s algorithm would give, as required, that  $\Pr \left[ \text{SSP}_v \left( \prod_{u \in N^\tau(v)} \text{OUT}_u \right) \right] \geq 1 - \frac{1}{2n}$ , i.e., that a node  $v$  is successful with high probability. Notice also that only nodes that are not in the output independent set can fail to satisfy SSP, by definition. Therefore, deferring such nodes does not remove any nodes from the output independent set, and so all nodes satisfying  $\text{SSP}_v$  also satisfy  $\text{WSP}_v$ .

The arguments above imply that Luby’s algorithm is a normal  $(O(\log n), \Delta)$ -round distributed procedure (the other necessary properties are trivial to check). In fact, we had something stronger: we used essentially the same condition for the strong and weak success property (with the only formal difference being that the weak success property has a domain extended to include DEFER tags). This indicates that, in this case, deferring unsuccessful nodes does not hurt us *at all*. This is something we will also see in our main application to D1LC. Definition 5 is more general, though, and allows WSP to be weaker than SSP in addition to incorporating DEFER tags, which can provide some leeway in derandomization.

We will argue (in Section 5.1) that we can only efficiently derandomize normal  $(\tau, \Delta)$ -round distributed procedures in MPC when  $\tau$  is low (ideally constant). This means that, in most applications, it will not suffice to show that entire algorithms for a problem are normal distributed procedures; we must instead consider those algorithms as series of short subroutines, and then show that each of those subroutines is a normal  $(\tau, \Delta)$ -round distributed procedure. This means defining success properties that capture an appropriate level of progress after each subroutine, not just correctness of the final output. Indeed, the ability to capture subroutines in this way is one of the main generalizations of our approach over that of [CDP21a].

Coloring algorithms are particularly amenable to this type of analysis, since they often consist mostly of subroutines to generate *slack*. That is, some short subroutine is performed, after which analysis shows that nodes’ remaining palette sizes will be larger than their remaining degrees by some amount. This measure of slack is then the success property for the subroutine. Importantly, deferring nodes (until the end of the entire coloring process) can only ever increase slack, since those deferred nodes are removed from neighbors’ neighborhoods, but do not block any colors from neighbors’ palettes. So, we can again define both  $\text{SSP}_v$  and  $\text{WSP}_v$  to be essentially the same property: that node  $v$  has at least some amount of slack at the end

of a subroutine. As in the example of Luby's algorithm, deferring unsuccessful nodes (in fact, in this case, even successful nodes) cannot hurt  $v$  at all, and  $\text{SSP}_v \implies \text{WSP}_v$  under any subset of deferrals.

Next, we discuss some known results on pseudorandom generators that we use in our derandomization framework.

## 4.2 PRGs and derandomization

A *Pseudorandom Generator (PRG)* is a function that takes a short *random seed* and produces a longer string of *pseudorandom* bits, which are computationally indistinguishable from truly random bits. We use the following definition from [Vad12] for indistinguishability:

**Definition 6** (Definition 7.1 in [Vad12]). *Random variables  $X$  and  $Y$  taking values in  $\{0,1\}^m$  are  $(t, \varepsilon)$  **indistinguishable** if for every non-uniform algorithm  $T : \{0,1\}^m \rightarrow \{0,1\}$  running in time at most  $t$ , we have  $|\Pr[T(X) = 1] - \Pr[T(Y) = 1]| \leq \varepsilon$ .*

Let  $U_k$  denote a random variable generated uniformly at random from  $\{0,1\}^k$ . Then pseudorandom generators are defined as follows:

**Definition 7** (PRG, Definition 7.3 in [Vad12]). *A deterministic function  $\mathcal{G} : \{0,1\}^d \rightarrow \{0,1\}^m$  is an  $(t, \varepsilon)$  **pseudorandom generator (PRG)** if (1)  $d < m$ , and (2)  $\mathcal{G}(U_d)$  and  $U_m$  are  $(t, \varepsilon)$  indistinguishable.*

A simple application of the probabilistic method can show the existence of PRGs with optimal parameters:

**Proposition 8** (Proposition 7.8 in [Vad12]). *For all  $t \in \mathbb{N}$  and  $\varepsilon > 0$ , there exists a (non-explicit)  $(t, \varepsilon)$  pseudorandom generator  $\mathcal{G} : \{0,1\}^d \rightarrow \{0,1\}^t$  with seed length  $d = \Theta(\log t + \log(1/\varepsilon))$ .*

As shown in [CDP21a], such a PRG can be computed using relatively low space (but exponential computation).

**Lemma 9** (Lemma 35 of [CDP21a], arXiv version). *For all  $t \in \mathbb{N}$  and  $\varepsilon > 0$ , there exists an algorithm for computing the  $(t, \varepsilon)$  PRG of Proposition 8 in time  $\exp(\text{poly}(t/\varepsilon))$  and space  $\text{poly}(t/\varepsilon)$ .*

Next, we show how to use PRGs to derandomize normal distributed procedures.

## 4.3 Derandomizing normal distributed procedures

In this section, we will sometimes be working on graphs  $G$  which are smaller induced subgraphs of the original graph. So, we use  $n_G$  to denote the number of nodes in  $G$ , as opposed to  $n$  which will denote the original number of nodes. This distinction is primarily because the space bounds of the MPC machines are still in terms of  $n$  even when working on a smaller graph  $G$ . Similarly, for a node  $v$ ,  $N_G(v)$  denotes the set of neighbors of  $v$  in  $G$  and  $d_G(v)$  denotes the degree of  $v$  in  $G$ .

For any integer  $\ell \in \mathbb{N}$ , let  $G^\ell$  be the  $\ell$ -power of  $G$ , that is,  $G^\ell$  is the graph having the same node set as  $G$  and any pair of nodes at a distance at most  $\ell$  in  $G$  form an edge in  $G^\ell$ .

We begin with the following lemma.

**Lemma 10.** *There is a constant  $C$  such that, given an  $O(\Delta^{8\tau})$ -coloring of  $G^{4\tau}$ , any normal  $(\tau, \Delta)$ -round distributed procedure on a graph  $G$  can be derandomized in  $O(\tau)$  rounds of MPC, using  $s = O(\Delta^{\tau^C})$  local space per machine and global space  $O(n_G \Delta^{\tau^C})$ , with the following properties:*

- At most  $(\frac{1}{2} + \Delta^{-11\tau})n_G$  nodes are deferred.
- All non-deferred nodes  $v$  satisfy the weak success property.

*Proof.* First, for each node  $v$  we collect the input information of its  $8\tau$ -hop neighborhood  $(\text{IN}_v \cup \bigcup_{u \in N_G^\tau(v)} \text{IN}_u)$  to a dedicated machine. This takes  $\tau$  rounds, and requires  $O(\Delta^{8\tau} \cdot \Delta^\tau) = O(\Delta^{11\tau})$  space per machine and  $O(n_G \Delta^{11\tau})$  global space. Our aim is then to simulate the procedure using randomness produced by

the  $(\Delta^{11\tau}, \Delta^{-11\tau})$  PRG implied by Proposition 8. This PRG has seed length  $d = \Theta(\log \Delta)$  and requires  $\text{poly}(\Delta^{11\tau})$  space to construct and store. We choose  $C$  so that this  $\text{poly}(\Delta^{22\tau})$  term is  $O(\Delta^{\tau^C})$ .

The PRG, when evaluated on a seed, produces a string of  $\Delta^{11\tau}$  pseudorandom bits. We use the provided  $O(\Delta^{8\tau})$ -coloring of  $G^{4\tau}$  to split this string into the input randomness for each node. By definition of a normal  $(\tau, \Delta)$ -round distributed procedure, each node requires  $O(\Delta^{2\tau})$  random bits, and we provide a node colored  $i$  in the  $O(\Delta^{8\tau})$ -coloring with the  $i^{\text{th}}$  chunk of  $O(\Delta^{2\tau})$  bits from the PRG's output. This means that any pair of nodes within distance  $4\tau$  receive disjoint chunks of pseudorandom bits.

The output of the PRG under a random seed is  $(\Delta^{11\tau}, \Delta^{-11\tau})$  indistinguishable from a uniform distribution. Consider the process of simulating the procedure for all nodes within distance  $\tau$  of a node  $v$ , and then evaluating the strong success property  $\text{SSP}_v(\bigcup_{u \in N_G^\tau(v)} \text{OUT}_u)$ . By definition of a normal  $(\tau, \Delta)$ -round distributed procedure, this combined process requires  $O(\Delta^{9\tau})$  computation, and depends on the input information and randomness of nodes up to distance  $2\tau$  from  $v$  (and note that all nodes within this radius receive different chunks of the PRG's output as their pseudorandom bits). This combined process can therefore be run on one MPC machine. Furthermore, it can be considered a non-uniform algorithm using at most  $\Delta^{9\tau}$  computation, and so is 'fooled' by the PRG. This means that the output ( $T$  or  $F$ , indicating whether the success property is satisfied) at  $v$  differs with probability at most  $\Delta^{-11\tau}$  from what it would be under full randomness. That is, the output will be  $F$  with probability at most  $\frac{1}{2n_G} + \Delta^{-11\tau}$ .

The expected number of nodes that do not satisfy the strong success property, when simulating the procedure using the PRG with a random seed, is therefore at most  $\frac{1}{2} + n_G \Delta^{-11\tau}$ . As this value is an aggregate of functions computable by individual machines, using the method of conditional expectations (as implemented for low-space MPC in [CDP21b, CDP21d]) we can deterministically select a seed for the PRG for which the number of nodes which do not satisfy the success property is at most its expectation (i.e.,  $\frac{1}{2} + n_G \Delta^{-11\tau}$ ), in  $O(1)$  rounds.

Then, we simply mark the nodes which do not satisfy the strong success property as deferred. By Definition 5, all non-deferred nodes still satisfy the weak success property. We therefore meet the conditions of the lemma.  $\square$

We can now iterate Lemma 10 in such a way that we can derandomize full algorithms for problems. To do so, we need these problems to satisfy a *self-reducibility property*, which informally means that partial solutions to the problem should be extendable to full solutions. As discussed, D1LC has this property, which is the reason why it emerges naturally in algorithms for  $(\Delta + 1)$ -coloring or  $(\Delta + 1)$ -list coloring, which themselves are not self-reducible.

**Definition 11.** Consider a graph problem  $\mathcal{P}$  in which each graph node  $v$  takes some input information  $\text{IN}_v$  and must produce some output labelling. We call  $\mathcal{P}$  **self-reducible** if it has the following property:

- For any graph  $G$ , any subset  $S \subseteq V(G)$ , and any valid output labelling on  $G$ , nodes  $v \in S$  can compute in  $O(1)$  rounds of LOCAL new input information  $\widehat{\text{IN}}_v$  such that:
  - the graph induced on  $S$ , with new inputs  $\widehat{\text{IN}}$ , forms a valid instance of  $\mathcal{P}$ , and
  - replacing the output labelling of nodes in  $S$  with any valid output labelling for this induced problem still forms a valid output of the original problem on  $G$ .

Now, we can now state our main derandomization theorem.

**Theorem 12.** Consider a randomized algorithm  $\mathcal{A}$  for a self-reducible problem  $\mathcal{P}$  which consists of a series of  $k$  ( $k = n^{o(1)}$ ) normal  $(\tau, \Delta)$ -round distributed procedures (i.e., the final weak success property implies a valid output for  $\mathcal{P}$ ). Then, there is a constant  $C$  such that for any  $\delta \in (0, 1)$ , on any graph  $G$  with  $\Delta \leq n^{7\delta}$ ,  $\mathcal{A}$  can be derandomized in  $O(k\tau + \log^* n_G)$  rounds of MPC, using  $s = O(n^{7\delta\tau^C})$  space per machine and global space  $O(n_G \cdot n^{7\delta\tau^C})$ , with all nodes giving a valid output for  $\mathcal{P}$ .  $\mathcal{A}$  may also contain up to  $O(k\tau + \log^* n_G)$  deterministic steps of LOCAL or MPC adhering to the same space bounds.

*Proof.* We first note that any deterministic LOCAL or MPC steps contained in  $\mathcal{A}$  can simply be run as they are, and do not need any derandomization. For derandomization of the randomized procedure, we begin by computing an  $O(\Delta^{8\tau})$ -coloring of  $G^{4\tau}$  in  $O(\tau + \log^* n)$  rounds, by simulating round-by-round the  $O(\Delta^2)$ -coloring algorithm of Linial[Lin92] on the graph  $G^{4\tau}$ . Constructing this graph requires collecting  $4\tau$ -radius balls around each node onto machines, which can be done since we allow  $O(n^{7\delta\tau C})$  space per machine.

Then, we apply Lemma 10 to derandomize each of the  $k$  normal  $(\tau, \Delta)$ -round distributed procedures forming  $\mathcal{A}$  in order, in each case using  $\Delta = n^{7\delta}$  as an upper bound on the maximum degree (even though the actual maximum degree may be significantly lower). The result is that at most  $k(\frac{1}{2} + n_G \Delta^{-11\tau})$  nodes are deferred, and the other nodes satisfy the weak success property of the final procedure (have a valid output labelling). This takes  $O(k\tau)$  rounds of MPC, and uses  $s = O(n^{7\delta\tau C})$  space per machine and global space  $O(n_G \cdot n^{7\delta\tau C})$ .

Next, since the problem is self-reducible, the deferred nodes can compute, in  $O(1)$  rounds, new inputs such that it suffices to solve the problem on the induced graph of deferred nodes with these new inputs. To do so, we recursively apply the above process, again using  $\Delta = n^{7\delta}$  as an upper bound on the maximum degree. After  $r$  of these recursive applications, the number of remaining deferred nodes is at most  $\frac{k}{2} + n(k^r \cdot n^{-11\tau r \cdot 7\delta})$ . Taking  $r = 1/\delta$  (which is  $O(1)$ ) and  $\tau$  to be at least a suitable large constant, we then have  $n^{o(1)}$  remaining deferred nodes. We can greedily find valid output labels for them in  $O(1)$  further rounds collecting their induced graph onto a single machine, which greedily assigns them valid output labels in any order. Due to the self-reducibility of the problem, we now have a valid output for all nodes for the original problem on  $G$ .

The total number of rounds used is  $O(\tau + \log^* n + k\tau) = O(k\tau + \log^* n)$ , and the space used is  $s = O(n^{7\delta\tau C})$  per machine and  $O(n_G \cdot n^{7\delta\tau C})$  global space.  $\square$

## 5 D1LC Algorithm when degree is at most $n^{7\delta}$

In this section we show that our framework from Section 4, when combined with the D1LC algorithm of Halldórsson et al. [HKNT22], leads to efficient deterministic MPC algorithm for D1LC when the maximum degree is  $n^{7\delta}$ . Note that  $\delta \in (0, 1)$  is a constant set suitably low relative to the local space parameter  $\phi$  and is to be fixed later when we finally prove Theorem 1 in Section 6. Specifically, we show that in  $O(\log^* n)$  rounds, with  $s = O(n^\phi)$  and  $O(n^{1+\phi})$  global space, a subset of the vertices are colored deterministically and the graph induced on the nodes left uncolored will have maximum degree  $O(\log^7 n)$  (and will be a D1LC instance, due to the self-reducibility of D1LC). In Section 5.1, we argue how to derandomize the randomized subroutines of Halldórsson et al. , using the framework discussed in Section 4, when the maximum degree of the graph is at most  $n^{7\delta}$ . In Section 5.2, we will argue that the deterministic subroutines which the algorithm of Halldórsson et al. uses can be implemented in sublinear MPC, provided the maximum degree is at most  $\sqrt{s}$ . Due to a result of Czumaj et al. [CDP21c], this graph can be colored in  $O(\log \log \log n)$  rounds.

In this section, graph  $G$  is often not clear from the context. So, we use  $n_G$  to denote the number of nodes in  $G$ . For a node  $v$ ,  $N_G(v)$  denote the set of neighbors of  $v$  in  $G$  and  $d_G(v)$  denotes the degree of  $v$  in  $G$ . We denote the maximum degree of any node in  $G$  by  $\Delta_G$ . For a node  $v$  and  $\tau \in \mathbb{N}$ ,  $N_G^\tau(v)$  denotes the nodes in the  $\tau$ -hop neighborhood of  $v$ .

### 5.1 Derandomization of key randomized subroutines of [HKNT22]

Recall the randomized subroutines of [HKNT22] as discussed in Section 2.2. Note that all the subroutines succeeds with high probability when the degree of each node is at least  $\log^7 n$ . Here, we discuss the derandomization of the subroutines by proving that all of them are  $(O(1), \Delta)$ -round normal distributed procedures in Lemma 13 and arguing that such procedures can be efficiently derandomized by using the framework outlined in Section 4.

**Lemma 13.** *The pre-shattering part of the algorithm of [HKNT22] is a series of  $O(\log^* \Delta)$  normal  $(O(1), \Delta_G)$ -round distributed procedures, such that the final weak success property is that all nodes  $v$  with degree  $d_G(v) \geq \log^7 n$  are properly colored.*



*Proof.* The randomized subroutines used in the pre-shattering part of the algorithm are TRYRANDOMCOLOR (Algorithm 3), GENERATESLACK (Algorithm 6), PUTASIDE (Algorithm 9), SYNCHCOLORTRIAL (Algorithm 8), and SLACKCOLOR. We will show that TRYRANDOMCOLOR, GENERATESLACK, PUTASIDE, and SYNCHCOLORTRIAL are all normal  $(O(1), \Delta_G)$ -round distributed procedures, and SLACKCOLOR consists of a series of  $O(\log^* \Delta)$  normal  $(O(1), \Delta_G)$ -round distributed procedures.

First, we note that at the start of any algorithm, nodes may be assumed to have information about which sets they are members of ( $V_{\text{sparse}}$ ,  $V_{\text{uneven}}$ , etc.), as by earlier simulation arguments (see Observation 20) this can be computed in  $O(1)$  rounds in LOCAL (and MPC, for low enough degree). We may also assume that nodes have information about the value of any of the parameters in Definition 2 for themselves and any of their neighbours (see Lemma 18). Note that this is, in total,  $O(\Delta_G^2)$  words of information for each node  $v$ : a constant number of words for each parameter,  $O(\Delta_G)$  words for set membership, multiplied by the number of neighbors of  $v$ , which is  $O(\Delta_G)$ .

For each of the sub-procedures, we must define strong and weak success properties satisfying the conditions of Definition 5, that capture the notion of success for the subroutine. In each case, our weak success properties will be identical to the corresponding strong success property, other than the extension to deferred nodes. All of our success properties will also deem nodes of degree less than  $n^{7\delta}$  to always be successful, regardless of what happens during the randomized process. This means that in this section we will not show any constraints on these low-degree nodes - they will be dealt with afterwards using Lemma 14. We look at each procedure:

TRYRANDOMCOLOR (Algorithm 3): It takes  $O(1)$  rounds of LOCAL. Nodes need no other words of input information. Nodes only use information from their neighbors, and each node uses  $O(\log \Delta_G)$  random bits to select a color from its palette. Note that the computation is  $O(\Delta_G)$ . The output information is either a color with which  $v$  has permanently colored itself, or FAIL if it does not color itself: this is clearly  $O(\Delta_G)$  words of information. We set the success properties  $\text{SSP}_v$  and  $\text{WSP}_v$  to be that either the slack of  $v$  increases from  $c \cdot d_G(v)$  for some constant  $c$  to  $2 \cdot d_G(v)$  [HKNT22, Lemma 26], or  $d_G(v) < \log^7 n$ . Here, for the purposes of  $\text{WSP}_v$ , deferred neighbors are discounted from  $v$ 's degree and therefore contribute to its slack. So, deferring neighbors only increases  $v$ 's slack, and therefore the last condition of Definition 5 is satisfied. These properties are computable in time linear in the degree of a node and based only on the output of the immediate neighbors of a node. The property succeeds for each node with probability  $p = \exp(\Omega(s(v)))$  for nodes with  $\log^7 n$  and 1 otherwise, so this is with high probability in  $n$ . So, all necessary conditions are satisfied and TRYRANDOMCOLOR meets Definition 5.

GENERATESLACK (Algorithm 6): This takes  $O(1)$  rounds of LOCAL and nodes need no other information at the beginning of the procedure. During the procedure, nodes only use information from their neighbors and  $\tilde{O}(\Delta_G)$  random bits (to determine whether the node is sampled and if so, what color is attempted). The output is either the color with which  $v$  permanently colored itself, or FAIL. The success properties  $\text{SSP}_v$  and  $\text{WSP}_v$  are that either that  $v$  generates *sufficient* slack, or  $d_G(v) < \log^7 n$ . By *sufficient* here we mean as described in the appropriate statements [HKNT22, Lemmas 10, 11, 13, 15, 17, 18]. These slack expressions are quite complicated: depending on the type of node, a different guarantee on the eventual slack is required. We note, however, that these guarantees all succeed with high probability for nodes with  $d_G(v) \geq \log^7 n$  (and again, for lower degree nodes we have defined  $\text{SSP}_v$  and  $\text{WSP}_v$  to always be satisfied) and are computable using only information in the immediate neighborhood of  $v$ . Deferring nodes again creates slack and so only helps nodes, so the success properties satisfy the necessary conditions of Definition 5.

PUTASIDE (Algorithm 9): The algorithm consists of 1 round of LOCAL, and nodes need no additional information. During the procedure, nodes only use information from their neighbours, and use  $O(\text{polylog}(n)) \in O(\Delta_G)$  (by construction) random bits to independently sample nodes into  $S_C$ . The output of the procedure is one word (whether  $v \in P_C$ ). The success properties  $\text{SSP}_v$  and  $\text{WSP}_v$  are that either: the put-aside set of the almost-clique containing  $v$  (say  $C$ ) satisfies  $|P_C| = \Omega(\ell^2)$  for some polylogarithmic  $\ell$  (see [HKNT22, Lemma 5]); or  $C$  contains  $\Omega(\ell^2)$  deferred nodes; or  $d_G(v) < \log^7 n$ . This can be checked using the 2-hop neighborhood of  $v$  because the diameter of  $C$  is at most 2. The strong success property succeeds with high probability, since [HKNT22] shows that  $|P_C| = \Omega(\ell^2)$  with high probability when  $d_G(v) \geq \log^7 n$ . If we defer nodes then this can reduce the size of the put-aside set by the same amount, but the deferred nodes

serve exactly the same purpose of providing slack to the almost clique and so this is still sufficient for the algorithm of [HKNT22]. Then,  $\text{SSP}_v$  and  $\text{WSP}_v$  must still be satisfied after deferrals, as required.

**SYNCHCOLORTRIAL** (Algorithm 8): This procedure takes 1 round of **LOCAL**. Nodes need no additional information about their neighborhood beforehand. Nodes only use information from their immediate neighbors during the procedure (since leader  $x_C$  is adjacent to all nodes in the set of inliers  $I_C$ ), and in the worst case (i.e. if the node in question is  $x_C$ ), nodes need  $\tilde{O}(\Delta_G)$  random bits to permute their palette. The output of the procedure for a node  $v \in I_C$  for some  $C$  is either a color in  $\Psi_v$  or **FAIL**. The success property for the procedure in [HKNT22] is that, for the almost-clique  $C$  that  $v$  is in, the number of nodes that return **FAIL** is bounded by  $O(t)$ , where  $t$  is some polylogarithmic value [HKNT22, Lemma 7]. So, we accordingly set  $\text{SSP}_v$  and  $\text{WSP}_v$  to be that either:  $C$  contains  $O(t)$  nodes that **FAIL**; or  $d_G(v) < \log^7 n$ . This happens with high probability ( $\exp(-t)$  when  $d_G(v) \geq \log^7 n$ , and 1 otherwise). Deferring nodes cannot increase the number of failed nodes, so  $\text{SSP}_v$  implies  $\text{WSP}_v$  after deferrals.

**SLACKCOLOR** (Algorithm 2): There are three parts to the algorithm. First, **TRYRANDOMCOLOR** is called in order to amplify the slack of each node (which is linear in the degree: a prerequisite of **SLACKCOLOR**). Then two loops of **MULTITRIAL** (Algorithm 4) instances are run. The  $O(1)$  calls to **TRYRANDOMCOLOR** are normal  $(O(1), \Delta_G)$ -distributed procedures by the item above. Each iteration of the first for-loop of executions of **MULTITRIAL** takes  $O(1)$  **LOCAL** rounds. Nodes need no additional input information. During the execution of the loop, nodes only use information from their immediate neighbors, and only use  $O(\Delta_G)$  random bits to select a subset of their palette. The output of the procedure for  $v$  is either a color in  $\Psi_v$  or **FAIL**, which is one word. Nodes succeed if  $d_i(v) \leq s(v)/\min(2^{x_i}, \rho^\kappa)$ , which occurs with high probability when  $d_G(v) \geq n^{7\delta}$  [HKNT22, Lemma 27]. This is, again, computable using only the output of the node and output information of the immediate neighborhood. Nodes deferring only reduces the degree, and so they cannot cause the property to be unsatisfied. So, setting  $\text{SSP}_v$  and  $\text{WSP}_v$  to be that  $v$  satisfies the above inequality or has  $d_G(v) < \log^7 n$  meets the conditions of Definition 5. The second for-loop of **MULTITRIAL** instances are normal distributed procedures for the same reason: the success property is  $d_i(v) \leq s(v)/\min(\rho^{(i+1)\cdot\kappa}, \rho)$ , and we can set  $\text{SSP}_v$  and  $\text{WSP}_v$  analogously to above for the same reason, satisfying the requirements of Definition 5. So, overall **SLACKCOLOR** consists of a sequence of  $O(\log^* \Delta)$  normal  $(O(1), \Delta_G)$ -round distributed procedures.  $\square$

We have shown that the pre-shattering algorithm of [HKNT22] consists of a series of  $(O(1), \Delta_G)$ -round distributed procedures as required. However, our success properties do not constrain nodes  $v$  with  $d_G(v) < \log^7 n$ , so to reach a full **D1LC** we must deal with these nodes afterwards. We can do so using the following lemma from [CDP21c]:

**Lemma 14** (Lemma 14 of [CDP21c]). *For any  $n$ -node graph  $G$  with maximum degree  $\Delta = \log^{O(1)} n$ , there exists an  $O(\log \log \log n)$ -round deterministic algorithm for computing **D1LC**, using  $O(n^\alpha)$  space per machine and  $O(n^{1+\alpha})$  global space, for any positive constant  $\alpha \in (0, 1)$ .*

Now, we can apply the pre-shattering algorithm of [HKNT22], followed by the low-degree algorithm of [CDP21c]. This fits into the framework of Theorem 12 and can therefore be derandomized:

**Lemma 15.** *There are constants  $c, C$  such that, for any constant  $\delta > 0$ , Algorithm 10 performs **D1LC** deterministically in  $O(\log \log \log n)$  rounds of MPC on any graph  $G$  with  $\Delta \leq n^{7\delta}$ , using  $s = O(n^{7\delta cC})$  space per machine and global space  $O(n_G \cdot n^{7\delta cC})$ .*

*Proof.* We have shown that the pre-shattering algorithm of [HKNT22] consists of a series of  $k = O(\log^* n) = O(\log \log \log n)$   $(c, \Delta_G)$ -round distributed procedures (for some sufficiently large constant  $c$ ), and some deterministic procedures which we can implement efficiently in MPC (see Section 5.2). We follow this with the  $O(\log \log \log n)$ -round deterministic low-space MPC procedure of Lemma 14. Setting  $\alpha$  (in Lemma 14) sufficiently lower than  $\delta$ , by Theorem 12, this algorithm can be derandomized in  $O(\log \log \log n)$  rounds of MPC using  $O(n^{7\delta cC})$  space per machine and  $O(n_G \cdot n^{7\delta cC})$  global space.  $\square$

---

**Algorithm 10:** DERANDOMIZEDMIDDEGREECOLOR( $G$ )

---

- 1 Let  $\mathcal{A}$  consist of the pre-shattering randomized LOCAL of [HKNT22] followed by the  $O(\log \log \log n)$ -round deterministic low-space MPC algorithm of [CDP21c].
  - 2 Derandomize  $\mathcal{A}$ , using the success properties from the proof of Lemma 13, by Theorem 12.
- 

## 5.2 Simulating deterministic LOCAL Subroutines of [HKNT22]

Here, we provide some arguments that the subroutines in [HKNT22] which are already deterministic can be simulated efficiently in the MPC model, provided that the maximum degree of the D1LC instance is below  $\sqrt{s}$ . In particular we show that:

**Lemma 16.** *In the sublinear local space MPC model with local space  $s = O(n^\phi) \geq \Delta^2$  and global space  $O(m+n^{1+\phi})$ , there exist algorithms running deterministically in  $O(1)$  rounds computing each of the following:*

- all the parameters (see Definition 2),
- a  $(\deg + 1)$ -ACD (see Definition 3),
- the set  $V_{\text{start}} \subseteq V_{\text{sparse}}$ <sup>10</sup> (whose purpose was described in Section 2.2.2),
- the leaders, inliers, and outliers from each almost-clique,
- and the put-aside sets for each almost-clique.

The rest of the section is concerned with the proof of each of the components of Lemma 16.

We begin with a useful tool that says (in essence) that nodes can send and receive  $O(\sqrt{s})$  messages to each of their neighbours in one round<sup>11</sup>, provided that their degree is at most  $\sqrt{s}$ .

**Lemma 17.** *Let  $G = (V, E)$  be a graph  $G$  such that for all  $v \in V$ ,  $d(v) \leq \sqrt{s}$ .*

*Suppose each node  $v$  is assigned some unique machine  $M(v)$ : informally,  $M(v)$  is responsible for  $v$ . Then, the following subroutines can be performed in parallel for each  $v$  (and its assigned machine  $M(v)$ ) in  $O(1)$  rounds in MPC with  $s = O(n^\phi)$ , and global space  $O(m + n^{1+\phi})$ :*

- $M(v)$  can send  $d(v)$  messages to each machine  $M'$  such that  $M(u)$  is responsible for node  $u \in N(v)$
- $M$  can collect all of the edges between neighbours of  $v$

*Proof.* Suppose each node is allocated a machine (giving global space of  $O(m + n^{1+\phi})$ ). The machine responsible for node  $v$  can prepare  $d(v)$  messages for each of the neighbors of  $v$ . This takes  $\leq s$  local space, so these messages can be sent in  $O(1)$  rounds. The second subroutine follows from the first: the  $d(v)$  words that each node broadcasts consists of the list of its neighbors.  $\square$

We can now argue that provided the maximum degree of our graph is low enough, the parameters described in Definition 2 can be efficiently computed in parallel for each node.

**Lemma 18** (Computing Parameters of Definition 2). *For a graph  $G = (V, E)$  with maximum degree  $\Delta \leq O(\sqrt{s})$ , all of the parameters listed in Definition 2 can be computed for every node in parallel in MPC with  $s = O(n^\phi)$  and global space  $O(m + n^{1+\phi})$  in  $O(1)$  rounds.*

*Proof.* We extensively use the subroutines in Lemma 17. We deal with the parameters in the order they appear in Definition 2.

---

<sup>10</sup>A full definition of  $V_{\text{start}}$  will be given later in this section.

<sup>11</sup>We remark that the MPC model is not node-centric, but machine-centric, and that when we talk about node  $u$  “sending messages” to node  $v$ , we suppose that each node has a machine responsible for it, and mean that the machine responsible for node  $u$  is sending a message to the machine responsible for node  $v$ .

- **Slack:** In  $O(1)$  rounds, we can sort tuples corresponding to edges incident to  $v$  and colors in  $\Psi(v)$  such that they appear on consecutive machines: we can then count how many of each there are and compute  $s(v)$ .
- **Sparsity:** Since we can compute  $d(v)$  for all nodes in constant rounds (machines can collect all edges incident to  $v$  and count them), to compute  $\zeta_v$  it remains to count the edges between neighbors of  $v$ . By the second subroutine in Lemma 17 we can collect the 2-hop neighborhood of  $v$  on a single machine. We can then remove duplicates if necessary and count the number of edges, and then calculate  $\zeta_v$ .
- **Disparity:** By the first subroutine of Lemma 17 we can collect the palette of each of the neighbors of  $v$  in constant rounds, allowing us to calculate  $\bar{\eta}_{u,v}$  for each neighbor  $u$  of  $v$ .
- **Discrepancy:** Follows immediately from the ability to collect  $\bar{\eta}_{u,v}$  for all neighbors  $u$  of  $v$ .
- **Unevenness:** It suffices to collect  $d(u)$  from each neighbor  $u$  of  $v$ , which can be done using the first subroutine of Lemma 17.
- **Slackability / Strong Slackability:** Note that these are additions of already computed parameters, and hence can be computed in constant rounds.

Mainly due to Lemma 17 and from the above description, it is easy to see that each of the above parameters can be computed in  $O(1)$  rounds for all the nodes in parallel. For the global space complexity for computing the above parameters, note that each node  $v$  collects its 2-neighborhood which is of size  $O(s)$  as maximum degree of any node is  $O(\sqrt{s})$ .  $\square$

Next, we show how to compute an  $(\deg + 1)$ -ACD in  $O(1)$  rounds.

**Lemma 19** (Computing  $(\deg + 1)$ -ACD). *For a graph  $G = (V, E)$  with maximum degree  $\Delta \leq \sqrt{s}$ , an  $(\deg + 1)$ -ACD can be computed in MPC with  $s = O(n^\phi)$  and global space  $O(m + n^{1+\phi})$  in  $O(1)$  rounds.*

*Proof.* First we argue that the diameter of the subgraph induced by any almost-clique  $C$ , is at most 2. This is the case because, by (iv) of Definition 3, the number of neighbors any node  $v \in C$  has in  $C$  is  $|N(v) \cap C| \geq \frac{|C|}{1+\varepsilon_{ac}} > \frac{|C|}{2}$ . Hence, any two nodes  $u$  and  $v$  in  $C$  have a common neighbor in  $C$ .

By Lemma 18, all nodes ( $v$ ) can compute their sparsity ( $\zeta_v$ ) and unevenness ( $\eta_v$ ) in  $O(1)$  rounds in parallel. From the values of  $\zeta_v$  and  $\eta_v$ , we can decide whether  $v \in V_{\text{sparse}}$ , or  $v \in V_{\text{uneven}}$ , or neither, giving the following:

**Observation 20.**  $V_{\text{sparse}}$  and  $V_{\text{uneven}}$  can be determined in  $O(1)$  rounds.

Note that each node  $v \in (V \setminus (V_{\text{sparse}} \sqcup V_{\text{uneven}}))$  is present in some almost-clique  $C(v)$ . As the diameter of  $G[C(v)]$  is at most 2,  $C(v)$  is a subset of the 2-hop neighborhood of  $v$ . As  $d(v) \leq \sqrt{s}$ , by Lemma 17 the 2-hop neighborhood of  $v$  can be found in  $O(1)$  rounds and stored on a single machine. Once this is done,  $M(v)$  can determine which almost-clique  $v$  is in. Putting everything together, the ACD can be computed in  $O(1)$  rounds and the global space used is  $O(m + n^{1+\phi})$ .  $\square$

We now consider the identification of  $V_{\text{start}}$ , the subset of  $V_{\text{sparse}} \cup V_{\text{uneven}}$  for which it is hard to generate slack. To identify this set, we consider a breakdown of  $V_{\text{sparse}} \cup V_{\text{uneven}}$  into several sets based on the parameters defined in Definition 2. We first explain the notion of a “heavy” color. A heavy color  $c$  with respect to a node  $v$  is a color such that if all of the neighbors of  $v$  were to pick a color from their palettes uniformly at random, the expected number of neighbors of  $v$  which would pick  $c$  (which we denote  $H(c)$ ) is at least some suitable constant. We denote by  $\mathcal{C}_v^{\text{heavy}}$  the set of heavy colors with respect to  $v$ . Now  $V_{\text{start}}$  is defined as follows (taken from [HKNT22]; the  $\varepsilon_i$  are all constants):

$$\begin{aligned}
V_{\text{balanced}} &= \{v \in V_{\text{sparse}} : |\{u \in N(v) : d(u) > 2d(v)/3\}| \geq \varepsilon_1 d(v)\}. \\
V_{\text{disc}} &= \{v \in V_{\text{sparse}} : \bar{\eta}_v \geq \varepsilon_2 d(v)\}. \\
V_{\text{easy}} &= V_{\text{balanced}} \cup V_{\text{disc}} \cup V_{\text{uneven}} \\
&\quad \cup \{v \in V_{\text{sparse}} : |N(v) \cap V_{\text{dense}}| \geq \varepsilon_3 d(v)\}. \\
V_{\text{heavy}} &= \{v \in V_{\text{sparse}} \setminus V_{\text{easy}} : \sum_{c \in \mathcal{C}_v^{\text{heavy}}} H(c) \geq \varepsilon_4 d(v)\}. \\
V_{\text{start}} &= \{v \in V_{\text{sparse}} \setminus (V_{\text{easy}} \cup V_{\text{heavy}}) : |N(v) \cap V_{\text{easy}}| \geq \varepsilon_5 d(v)\}.
\end{aligned}$$

**Lemma 21** (Identifying  $V_{\text{start}}$ ). *Given a graph  $G = (V, E)$  with maximum degree  $\Delta \leq \sqrt{s}$ , the set  $V_{\text{start}}$  can be identified on an MPC with  $s = O(n^\phi)$  and global space  $O(m + n^{1+\phi})$  in  $O(1)$  rounds.*

*Proof.* Observe that  $V_{\text{balanced}}$ ,  $V_{\text{disc}}$ , and  $V_{\text{easy}}$  are computable in  $O(1)$  rounds by applications of Lemma 18 and Lemma 17. We briefly note that, again using an application of the first subroutine of Lemma 17, nodes can be made aware of which neighbors are in which sets in  $O(1)$  rounds.

It remains to compute which colors are heavy (and the set  $V_{\text{heavy}}$ ): by the first subroutine of Lemma 17 we can gather the palettes of all neighbors of  $v$  on  $M(v)$ , and then  $M(v)$  can compute this information in  $O(1)$  rounds.  $\square$

In the following lemma, we show that we can identify leaders, inliers, and outliers of almost-cliques, and establish which almost-cliques require put-aside sets, in  $O(1)$  rounds.

**Lemma 22** (Finding leaders, outliers and inliers). *Given a graph  $G = (V, E)$  with maximum degree  $\Delta \leq \sqrt{s}$ , we can find the leader, the set of inliers and the set of outliers for each almost-clique in MPC with  $s = O(n^\phi)$  and global space  $O(m + n^{1+\phi})$  in  $O(1)$  rounds. Moreover, we can detect all the almost-cliques for which we need to find put-aside sets.*

*Proof.* First, we explain how we can find the leader  $x_C$  and the outliers  $O_C$  for each almost-clique  $C \in \mathcal{C}$ . The leader  $x_C$  is the node in  $C$  with minimum slackability. As we can determine the slackability of all the nodes (in  $C$ ) by using  $O(1)$  rounds in parallel (see Lemma 18), we can determine  $x_C$  for each  $i \in [t]$  in parallel in  $O(1)$  rounds. The slackability of an almost-clique  $C$  is defined as the slackability of its leader  $x_C$ .  $C$  is said to be a *low slack* almost-clique if its slackability is at most  $\ell = \log^{2.1} \Delta$ . Recall that these are the almost-cliques (with low slackability) for which we need to find put-aside sets.

Let  $O_C$  be the set of *outliers* in an almost-clique  $C$ .  $O_C$  is comprised of the union of the  $\frac{\max\{d(x_C), |C|\}}{3}$  nodes with the fewest common neighbors with  $x_C$ , the  $\frac{|C|}{6}$  nodes of largest degree, and the nodes in  $C$  that are not neighbors of  $x_C$ . The nodes in  $C$  that are not outliers are called *inliers* of  $C$  (denoted by  $I_C$ ). As each almost-clique  $C$  is stored in one machine and we can determine the degrees of the nodes in  $O(1)$  rounds, we can clearly find  $O_C$  and  $I_C$  for all almost-cliques in  $O(1)$  rounds.  $\square$

Finally, we conclude by observing that we have shown each component of Lemma 16:

*Proof of Lemma 16.* Follows from Lemmas 18, 19, 21 and 22 and observation 20.  $\square$

## 6 Overall Deterministic D1LC Algorithm

In this section we extend our (now derandomized) algorithm for D1LC (as discussed in Algorithm 10) to handle instances with maximum degree  $\Delta \geq n^{7\delta}$ . The idea is to reduce our D1LC instance to a collection of instances with lower degree so that constant-radius neighborhoods of any node fit onto a single machine. Instead of guaranteeing that there are not too many instances, we ensure that the sequential dependency between instances is not too long. That is, many of the instances resulting from our decomposition can be colored in parallel; there are only  $O(1)$  sets of base-case instances which must be solved sequentially.

We use a recursive structure `LOWSPACECOLORREDUCE` (Algorithm 11) similar to [CDP21d, CDP21c]. The recursive structure in Algorithm 11 relies on a partitioning procedure (Algorithm 12) to divide the nodes and colors in the input instance into *bins*. We can follow the analysis of [CDP21d] to analyze the partitioning process, since it is the base case that changed. Lemma 23 provides the important properties of the partitioning.

---

**Algorithm 11:** `LOWSPACECOLORREDUCE( $G$ )`

---

- 1  $G_{\text{mid}}, G_1, \dots, G_{n^\delta} \leftarrow \text{LOWSPACEPARTITION}(G)$ .
  - 2 For each  $i = 1, \dots, n^\delta - 1$  in parallel: call `LOWSPACECOLORREDUCE( $G_i$ )`.
  - 3 Update color palettes of  $G_{n^\delta}$ , call `LOWSPACECOLORREDUCE( $G_{n^\delta}$ )`.
  - 4 Update color palettes of  $G_{\text{mid}}$ .
  - 5 Color  $G_{\text{mid}}$  using `DERANDOMIZEDMIDDEGREECOLOR( $G_{\text{mid}}$ )`.
- 

---

**Algorithm 12:** `LOWSPACEPARTITION( $G$ )`

---

- 1 Let  $G_{\text{mid}}$  be the graph induced by the set of nodes  $v$  with  $d(v) \leq n^{7\delta}$ .
  - 2 Let hash function  $h_1 : [n] \rightarrow [n^\delta]$  map each node  $v$  to a bin  $h_1(v) \in [n^\delta]$ .
  - 3 Let hash function  $h_2 : [n^2] \rightarrow [n^\delta - 1]$  map colors  $\gamma$  to a bin  $h_2(\gamma) \in [n^\delta - 1]$ .
  - 4 Let  $G_1, \dots, G_{n^\delta}$  be the graphs induced by bins  $1, \dots, n^\delta$  respectively, minus the nodes in  $G_{\text{mid}}$ .
  - 5 Restrict palettes of nodes in  $G_1, \dots, G_{n^\delta-1}$  to colors assigned by  $h_2$  to corresponding bins.
  - 6 Return  $G_{\text{mid}}, G_1, \dots, G_{n^\delta}$ .
- 

**Lemma 23** (Lemma 4.6 of [CDP21d]). *Assume that, at the beginning of a call to `LOWSPACEPARTITION`, we have  $d(v) < p(v)$  for all nodes  $v$ . Then, in  $O(1)$  MPC rounds with  $O(n^{7\delta})$  local space per machine and  $O(n + m)$  global space (over all parallel instances), one can deterministically select hash functions  $h_1, h_2$  such that after the call,*

- for any node  $v \notin G_{\text{mid}}$ ,  $d'(v) < 2d(v)n^{-\delta}$ , and
- for any node  $v$ ,  $d'(v) < p'(v)$ .

Here  $d'(v)$  denotes the degree of  $v$  in the subgraph induced by the nodes present in the same bucket as  $v$  and  $p'(v)$  denotes the number of  $v$ 's palette colors that are in the same bucket as  $v$ .

Now, we are ready to prove our main result (Theorem 1):

*Proof of Theorem 1.* As in [CDP21c], calling `LOWSPACECOLORREDUCE` on our input graph creates a recursion tree of  $O(1)$  depth (since each recursive call reduces the maximum degree by a  $n^{-\delta}$  factor). It therefore creates  $O(1)$  sequential sets of base-case instances to solve concurrently, which in Algorithm 11 are solved by `DERANDOMIZEDMIDDEGREECOLOR`. Furthermore, each set of concurrent instances has at most  $n$  nodes in total, since all nodes are only partitioned into one instance, and each instance has maximum degree  $n^{7\delta}$ .

By Lemma 15, each such instance  $G$  is colored in  $O(\log \log \log n)$  rounds using  $O(n^{7\delta cC})$  space per machine and  $O(n_G \cdot n^{7\delta cC})$  global space. The global space used by all concurrent instances is therefore  $O(n^{1+7\delta cC})$ . Setting  $\delta \leq \frac{\phi}{7cC}$ , this is  $O(n^\phi)$  space per machine and  $O(n^{1+\phi})$  global space. Since receiving the input and the first call to `LOWSPACEPARTITION` also requires  $O(m)$  global space, the overall space bound is  $O(m + n^{1+\phi})$ .  $\square$

## References

- [AA20] Noga Alon and Sepehr Assadi. Palette sparsification beyond  $(\Delta + 1)$  vertex coloring. In *Proceedings of the 24th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM'20)*, pages 6:1–6:22, 2020.
- [ANOY14] Alexandr Andoni, Aleksandar Nikolov, Krzysztof Onak, and Grigory Yaroslavtsev. Parallel algorithms for geometric graph problems. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC'14)*, pages 574–583, 2014.
- [ASS<sup>+</sup>18] Alexandr Andoni, Zhao Song, Clifford Stein, Zhengyu Wang, and Peilin Zhong. Parallel graph connectivity in log diameter rounds. In *Proceedings of the 59th IEEE Symposium on Foundations of Computer Science (FOCS'18)*, pages 674–685, 2018.
- [BBD<sup>+</sup>19] Soheil Behnezhad, Sebastian Brandt, Mahsa Derakhshan, Manuela Fischer, MohammadTaghi Hajiaghayi, Richard M. Karp, and Jara Uitto. Massively parallel computation of matching and MIS in sparse graphs. In *Proceedings of the 38th ACM Symposium on Principles of Distributed Computing (PODC'19)*, pages 481–490, 2019.
- [BE13] Leonid Barenboim and Michael Elkin. *Distributed Graph Coloring: Fundamentals and Recent Developments*. Morgan & Claypool Publishers, 2013.
- [BHH19] Soheil Behnezhad, MohammadTaghi Hajiaghayi, and David G. Harris. Exponentially faster massively parallel maximal matching. In *Proceedings of the 60th IEEE Symposium on Foundations of Computer Science (FOCS'19)*, pages 1637–1649, 2019.
- [BKM20] Philipp Bamberger, Fabian Kuhn, and Yannic Maus. Efficient deterministic distributed coloring with small bandwidth. In *Proceedings of the 39th ACM Symposium on Principles of Distributed Computing (PODC'20)*, pages 243–252, 2020.
- [BKS17] Paul Beame, Paraschos Koutris, and Dan Suciu. Communication steps for parallel query processing. *J. ACM*, 64(6):40:1–40:58, 2017.
- [CC22] Sam Coy and Artur Czumaj. Deterministic massively parallel connectivity. In *Proceedings of the 52th Annual ACM Symposium on Theory of Computing (STOC'22)*, pages 162–175, 2022.
- [CCDM23] Sam Coy, Artur Czumaj, Peter Davies, and Gopinath Mishra. Optimal  $(\text{degree} + 1)$ -coloring in Congested Clique. In *Proceedings of the 50th International Colloquium on Automata, Languages and Programming (ICALP'23)*, to appear, 2023.
- [CDP21a] Artur Czumaj, Peter Davies, and Merav Parter. Component stability in low-space massively parallel computation. In *Proceedings of the 40th ACM Symposium on Principles of Distributed Computing (PODC'21)*, pages 481–491, 2021.
- [CDP21b] Artur Czumaj, Peter Davies, and Merav Parter. Graph sparsification for derandomizing massively parallel computation with low space. *ACM Transactions on Algorithms*, 17(2), May 2021.
- [CDP21c] Artur Czumaj, Peter Davies, and Merav Parter. Improved deterministic  $(\Delta + 1)$  coloring in low-space MPC. In *Proceedings of the 40th ACM Symposium on Principles of Distributed Computing (PODC'21)*, pages 469–479, 2021.
- [CDP21d] Artur Czumaj, Peter Davies, and Merav Parter. Simple, deterministic, constant-round coloring in congested clique and MPC. *SIAM J. Comput.*, 50(5):1603–1626, 2021.
- [CFG<sup>+</sup>19] Yi-Jun Chang, Manuela Fischer, Mohsen Ghaffari, Jara Uitto, and Yufan Zheng. The complexity of  $(\Delta + 1)$  coloring in congested clique, massively parallel computation, and centralized local computation. In *Proceedings of the 38th ACM Symposium on Principles of Distributed Computing (PODC'19)*, pages 471–480, 2019.

- [CHPS20] Keren Censor-Hillel, Merav Parter, and Gregory Schwartzman. Derandomizing local distributed algorithms under bandwidth restrictions. *Distributed Computing*, 33(3):349–366, 2020.
- [CKP19] Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. An exponential separation between randomized and deterministic complexity in the LOCAL model. *SIAM J. Comput.*, 48(1):122–143, 2019.
- [CLM<sup>+</sup>18] Artur Czumaj, Jakub Łącki, Aleksander Mądry, Slobodan Mitrović, Krzysztof Onak, and Piotr Sankowski. Round compression for parallel matching algorithms. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC’18)*, pages 471–484, 2018.
- [CLP18] Yi-Jun Chang, Wenzheng Li, and Seth Pettie. An optimal distributed  $(\Delta + 1)$ -coloring algorithm? In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC’18)*, pages 445–456, 2018.
- [Dav23] Peter Davies. Improved distributed algorithms for the Lovász local lemma and edge coloring. In *Proceedings of the 34th ACM-SIAM Symposium on Discrete Algorithms (SODA’23)*, pages 4273–4295, 2023.
- [DG08] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.
- [FGG22] Manuela Fischer, Jeff Giliberti, and Christoph Grunau. Improved deterministic connectivity in massively parallel computation. In *Proceedings of the 36th International Symposium on Distributed Computing (DISC’22)*, pages 22:1–22:17, 2022.
- [FHM23] Manuela Fischer, Magnús M. Halldórsson, and Yannic Maus. Fast distributed Brooks’ theorem. In *Proceedings of the 34th ACM-SIAM Symposium on Discrete Algorithms (SODA’23)*, pages 2567–2588, 2023.
- [GG23] Mohsen Ghaffari and Christoph Grunau. Faster deterministic distributed MIS and approximate matching. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1777–1790. ACM, 2023.
- [GGK<sup>+</sup>18] Mohsen Ghaffari, Themis Gouleakis, Christian Konrad, Slobodan Mitrović, and Ronitt Rubinfeld. Improved massively parallel computation algorithms for MIS, matching, and vertex cover. In *Proceedings of the 37th ACM Symposium on Principles of Distributed Computing (PODC’18)*, pages 129–138, 2018.
- [Gha16] Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms (SODA’16)*, pages 270–277, 2016.
- [GHK18] Mohsen Ghaffari, David G. Harris, and Fabian Kuhn. On derandomizing local distributed algorithms. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 662–673, 2018.
- [GK21] Mohsen Ghaffari and Fabian Kuhn. Deterministic distributed vertex coloring: Simpler, faster, and without network decomposition. In *Proceedings of the 62nd IEEE Symposium on Foundations of Computer Science (FOCS’21)*, pages 1009–1020, 2021.
- [GKU19] Mohsen Ghaffari, Fabian Kuhn, and Jara Uitto. Conditional hardness results for massively parallel computation from distributed lower bounds. In *Proceedings of the 60th IEEE Symposium on Foundations of Computer Science (FOCS’19)*, pages 1650–1663, 2019.
- [GSZ11] Michael T. Goodrich, Nodari Sitchinava, and Qin Zhang. Sorting, searching, and simulation in the MapReduce framework. In *Proceedings of the 22nd International Symposium on Algorithms and Computation (ISAAC’11)*, pages 374–383, 2011.



- [GU19] Mohsen Ghaffari and Jara Uitto. Sparsifying distributed algorithms with ramifications in massively parallel computation and centralized local computation. In *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms (SODA'19)*, pages 1636–1653, 2019.
- [HKNT22] Magnús M. Halldórsson, Fabian Kuhn, Alexandre Nolin, and Tigran Tonoyan. Near-optimal distributed degree+1 coloring. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC'22)*, pages 450–463, 2022.
- [HNT22] Magnús M. Halldórsson, Alexandre Nolin, and Tigran Tonoyan. Overcoming congestion in distributed coloring. In *Proceedings of the 41st ACM Symposium on Principles of Distributed Computing (PODC'22)*, pages 26–36, 2022.
- [IBY<sup>+</sup>07] Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. Dryad: Distributed data-parallel programs from sequential building blocks. *SIGOPS Operating Systems Review*, 41(3):59–72, March 2007.
- [Kar85] Howard J. Karloff. *Fast Parallel Algorithms for Graph-theoretic Problems, Matching, Coloring and Partitioning*. University of California, Berkeley, 1985. PhD thesis.
- [KSV10] Howard J. Karloff, Siddharth Suri, and Sergei Vassilvitskii. A model of computation for MapReduce. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA'10)*, pages 938–948, 2010.
- [Kuh20] Fabian Kuhn. Faster deterministic distributed coloring through recursive list coloring. In *Proceedings of the 31st ACM-SIAM Symposium on Discrete Algorithms (SODA'20)*, pages 1244–1259, 2020.
- [Lin87] Nathan Linial. Distributive graph algorithms — global solutions from local data. In *Proceedings of the 28th IEEE Symposium on Foundations of Computer Science (FOCS'87)*, pages 331–335, 1987.
- [Lin92] Nathan Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
- [ŁMOS20] Jakub Łącki, Slobodan Mitrović, Krzysztof Onak, and Piotr Sankowski. Walking randomly, massively, and efficiently. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC'20)*, pages 364–377, 2020.
- [Lub86] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986.
- [LW10] Christoph Lenzen and Roger Wattenhofer. Brief announcement: Exponential speed-up of local algorithms using non-local communication. In *Proceedings of the 29th ACM Symposium on Principles of Distributed Computing (PODC'10)*, pages 295–296, 2010.
- [RG20] Václav Rozhoň and Mohsen Ghaffari. Polylogarithmic-time deterministic network decomposition and distributed derandomization. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC'20)*, pages 350–363, 2020.
- [RVW18] Tim Roughgarden, Sergei Vassilvitski, and Joshua R. Wang. Shuffles and circuits (on lower bounds for modern parallel computation). *Journal of the ACM*, 65(6):41:1–41:24, November 2018.
- [SW10] Johannes Schneider and Roger Wattenhofer. A new technique for distributed symmetry breaking. In *Proceedings of the 29th ACM Symposium on Principles of Distributed Computing (PODC'10)*, pages 257–266, 2010.

- [Vad12] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012.
- [Whi15] Tom White. *Hadoop: The Definitive Guide: Storage and Analysis at Internet Scale*. O’Reilly Media, Sebastopol, CA, 4th edition, 2015.
- [ZCF<sup>+</sup>10] Matei Zaharia, Mosharaf Chowdhury, Michael J. Franklin, Scott Shenker, and Ion Stoica. Spark: Cluster computing with working sets. In *Proceedings of the 2nd USENIX Workshop on Hot Topics in Cloud Computing (HotCloud’10)*, 2010.