Simple, Deterministic, Constant-Round Coloring in the Congested Clique

Artur Czumaj [*]   Peter Davies [†]   Merav Parter [‡]

Abstract

We settle the complexity of the $(\Delta + 1)$-coloring and $(\Delta + 1)$-list coloring problems in the CONGESTED CLIQUE model by presenting a simple *deterministic* algorithm for both problems running in a constant number of rounds. This matches the complexity of the recent breakthrough *randomized* constant-round $(\Delta+1)$-list coloring algorithm due to Chang et al. (PODC'19), and significantly improves upon the state-of-the-art $O(\log \Delta)$-round deterministic $(\Delta + 1)$-coloring bound of Parter (ICALP'18).

A remarkable property of our algorithm is its simplicity. Whereas the state-of-the-art *randomized* algorithms for this problem are based on the quite involved local coloring algorithm of Chang et al. (STOC'18), our algorithm can be described in just a few lines. At a high level, it applies a careful derandomization of a recursive procedure which partitions the nodes and their respective palettes into separate bins. We show that after $O(1)$ recursion steps, the remaining uncolored subgraph within each bin has linear size, and thus can be solved locally by collecting it to a single node. This algorithm can also be implemented in the Massively Parallel Computation (MPC) model provided that each machine has linear (in $\mathfrak{n}$, the number of nodes in the input graph) space.

We also show an extension of our algorithm to the MPC regime in which machines have *sublinear* space: we present the first deterministic $(\Delta + 1)$-list coloring algorithm designed for sublinear-space MPC, which runs in $O(\log \Delta + \log \log \mathfrak{n})$ rounds.

1   Introduction

In this paper, we study deterministic complexity of vertex coloring in two fundamental models of distributed and parallel computation: the CONGESTED CLIQUE and the Massively Parallel Computation models.

The *CONGESTED CLIQUE* model is a classic and prominent model of distributing computing, introduced by Lotker et al. [16]. It is a variant of the standard CONGEST message-passing model that allows all-to-all communication, where in every single round, each node can send $O(\log n)$ bits to each other node. The *Massively Parallel Computation (MPC)* model, introduced by Karloff et al. [14], is a nowadays standard theoretical model for parallel algorithms. It shares many similarities to earlier models of parallel computation, but it also allows for unlimited local

computation, making it close to various models of distributed computing, e.g., to CONGESTED CLIQUE.

The $(\Delta+1)$-coloring problem and its variations are considered to be corner-stone problems of local computation and are arguably among the most fundamental graph problems in parallel and distributed computing with numerous implications. In the $(\Delta+1)$-*coloring problem*, where Δ is the maximum degree of the input graph $\mathfrak{G} = (V, E)$, the goal is to color each node in $\mathfrak{G}$ with a color in $\{1, \ldots, \Delta+1\}$ such that no two neighbors have the same color. We also consider its generalization, the $(\Delta+1)$-*list coloring problem*, where each node has a (possibly different) palette of $\Delta+1$ colors, and the goal is to compute a legal coloring in which each node is assigned a color from its own palette. Finally, the most restrictive variant of the problem is the $(\deg+1)$-*list coloring problem*, which is the same as the $(\Delta+1)$-list coloring problem except that the size of each node u's palette is $\deg(u) + 1$, which might be considerably smaller than $\Delta+1$. It is easy to see that in each of the variations of the $(\Delta+1)$ coloring problem above there is always a legal coloring, and the main task is to understand the complexity of finding the sought coloring.

Our main result is that we can solve the $(\Delta+1)$-list coloring problem in a constant number of rounds of CONGESTED CLIQUE, thereby settling the deterministic complexity of $(\Delta+1)$-coloring and $(\Delta+1)$-list coloring in this model.

Theorem 1. *Deterministic $(\Delta+1)$-list coloring can be performed in* $\mathbf{O(1)}$ *rounds in CONGESTED CLIQUE.*

We obtain this by presenting the following linear-space MPC result, and then applying known reductions between the CONGESTED CLIQUE and MPC models (cf. Section 1.2):

Theorem 2. *Deterministic $(\Delta+1)$-list coloring can be performed in* $\mathbf{O(1)}$ *rounds in MPC with $O(\mathfrak{n})$ local space per machine and $O(\mathfrak{n}\Delta)$ total space.*

We note that for general $(\Delta+1)$-list coloring, the input size is $\Theta(\mathfrak{n}\Delta)$ and therefore the global MPC space bound in Theorem 2 is optimal. However, for the special case of $(\Delta+1)$-coloring, where initial palettes are all $\{1, \ldots, \Delta+1\}$ (and so need not be specified as input), the input size is $\Theta(\mathfrak{m}+\mathfrak{n})$. We remark that we can also obtain optimal global space in this case.

Theorem 3. *Deterministic $(\Delta+1)$-coloring can be performed in* $\mathbf{O(1)}$ *rounds in MPC with $O(\mathfrak{n})$ local space per machine and $O(\mathfrak{m}+\mathfrak{n})$ total space.*

Finally, we extend our approach to the low-space MPC regime, and provide the first $(\deg+1)$-list coloring (and therefore $(\Delta+1)$-list coloring and $(\Delta+1)$-coloring) algorithm designed for that model.

Theorem 4. *For any constant $\varepsilon > 0$, deterministic $(\deg+1)$-list coloring can be performed in MPC in $O(\log\Delta + \log\log\mathfrak{n})$ rounds, using $O(\mathfrak{n}^\varepsilon)$ space per machine and $O(\mathfrak{m}+\mathfrak{n}^{1+\varepsilon})$ total space.*

1.1 Models of computation: CONGESTED CLIQUE and MPC

In the CONGESTED CLIQUE model of distributed computation, our input is a graph $\mathfrak{G}$ with $\mathfrak{n}$ nodes and $\mathfrak{m}$ edges, and the aim is to solve some graph problem by performing computation at the *nodes of the input graph*. Each nodes' initial input consists of its adjacent edges; we then proceed in synchronous *rounds*, in which each node can perform some local computation, and then send $O(\log\mathfrak{n})$-bit messages to every other node (i.e., unlike in the related LOCAL and CONGEST models, communication is *not* restricted to the edges of the input graph). We require

all nodes to eventually output their local part of the problem's solution (in the case of coloring problems, this means the color they have chosen), and our aim is to minimize the number of communication *rounds* required to do so.

The Massively Parallel Computation (MPC) model is a parallel computing model first introduced by Karloff et al. [14]. It recently gained a significant popularity as it provides a clean abstraction of a number of massively parallel computation frameworks extensively used in applications, such as MapReduce [8, 9], Hadoop [22], Dryad [13], or Spark [23]. In the MPC model, there are $\mathfrak{M}$ machines, each with $\mathfrak{s}$ available space. The input graph $\mathfrak{G}$ is initially distributed arbitrarily across the machines (i.e., we need at least enough machines so that the total space $\mathfrak{M} \cdot \mathfrak{s}$ is at least the input size). We progress again in synchronous rounds, in which each machine can perform some local computation, and then send a message to each other machine. Unlike in CONGESTED CLIQUE, we do not bound the size of each message; we instead require that the total information sent and received by each machine in each round must fit into the machine's local space (i.e., is of size at most $\mathfrak{s}$). This in particular implies that the total communication of the MPC model is bounded by $\mathfrak{M} \cdot \mathfrak{s}$ in each round. At the end of the computation, machines must collectively output the solution; any machine can output any part so long as the total output is a complete and consistent solution. Our aim is again to minimize the number of communication rounds; we also now have local space $\mathfrak{s}$ and total space $\mathfrak{M} \cdot \mathfrak{s}$ as model parameters. Space regimes of particular interest are the linear-space regime (where $\mathfrak{s} = \Theta(\mathfrak{n})$), and the low-space regime ($\mathfrak{s} = \Theta(\mathfrak{n}^\varepsilon)$ for any constant $\varepsilon \in (0, 1)$). Here we measure space in terms of $O(\log \mathfrak{n})$-bit words, so that $O(\mathfrak{n})$ space is sufficient to store e.g. a node's neighborhood.

1.2 Relationship between models

The two models we study here have their provenance in different fields: CONGESTED CLIQUE from distributed computing, and MPC from parallel computing. Because of this, CONGESTED CLIQUE is stated as a message-passing model, with the computation performed by *nodes of the input graph*, whose initial input consists of the adjacent edges to that node (but communication is *not* restricted to the input graph, and can be between any pair of nodes). In MPC, on the other hand, the computation is performed by *machines* that are not associated with any particular part of the input graph, and indeed the input is initially distributed arbitrarily across machines.

It has been noted, though, that (under certain conditions on algorithms) CONGESTED CLIQUE is equivalent to MPC using $O(\mathfrak{n})$ space per machine and $O(\mathfrak{n}^2)$ total space [3]. In particular, due to the constant round routing algorithm of Lenzen [15], information can be redistributed essentially arbitrarily, so there is no need to associate the computational entities with nodes in the input graph. This is in stark contrast to the related LOCAL and CONGEST distributed models, in which the link between computation and input graph locality is integral.

Since we give results both for CONGESTED CLIQUE and low-space MPC, we will adopt the MPC perspective, which can accommodate both. That is, we will consider our computation to be done on *machines* with $\mathfrak{s}$ space, and give results for the linear-space regime ($\mathfrak{s} = \Theta(\mathfrak{n})$) which directly apply also to CONGESTED CLIQUE, and then for the low-space regime ($\mathfrak{s} = \Theta(\mathfrak{n}^\varepsilon)$).

1.3 Related work

Coloring in CONGESTED CLIQUE is a problem which has received considerable recent attention, with a succession of results on randomized algorithms:

1.3.1 State of the art for randomized coloring algorithms.

The starting point for the current $(\Delta + 1)$ coloring randomized CONGESTED CLIQUE algorithms is the following: provided that the maximum degree $\Delta = O(\sqrt{\mathfrak{n}})$, nodes can efficiently simulate the LOCAL $(\Delta + 1)$-list coloring algorithm of Chang et al. (CLP) within $O(\log^* \Delta)$ rounds. Building upon this observation, Parter [18] showed that the degree of the nodes can be reduced by applying $O(\log \log \Delta)$ steps of vertex partitioning. At each recursive partitioning step, the current graph was split into $poly(\Delta)$ vertex-disjoint subgraphs plus one additional left-over subgraph. These vertex-disjoint subgraphs were colored recursively in parallel, and the left-over subgraph was colored only once the coloring of the other subgraphs in its recursion level was complete. This approach led to an $O(\log \log \Delta \cdot \log^* \Delta)$-round algorithm. Parter and Su [19] improved this complexity to $O(\log^* \Delta)$, by modifying the internal details of the CLP algorithm to allow it to be simulated when $\Delta = \mathfrak{n}^{1/2+o(1)}$. In a recent breakthrough, Chang et al. [6] presented a $O(1)$ round procedure which was obtained by (i) modifying the recursive partitioning of [18], and (ii) modifying aspects of the CLP algorithm. The key innovative part in their recursive partitioning was to partition both the vertices and the *colors* in their respectability into independent coloring instances. This new partitioning idea was useful to support the more general $(\Delta+1)$-list coloring problem, as well as to provide efficient implementation low-space MPC, where (upon incorporating the subsequent polylogarithmic-round network decomposition algorithm of Rozhoň and Ghaffari [20]), the algorithm of Chang et al. [6] works in $O(\log \log \log \mathfrak{n})$ rounds. This round complexity is matched by a conditional lower bound due to Ghaffari et al. [10], which states that an $o(\log \log \log \mathfrak{n})$-round randomized component-stable coloring algorithm in low-space MPC would imply a $2^{\log^{o(1)} \log \mathfrak{n}}$-round deterministic LOCAL coloring algorithm, and a $\log^{o(1)} \log \mathfrak{n}$-round randomized one. (Note, though, that our algorithms are *not* component-stable, since they involve global agreement on hash functions.)

1.3.2 State of the art for deterministic coloring algorithms.

Censor-Hillel et al. [5] presented a quite general scheme for derandomization in the CONGESTED CLIQUE model by combining the methods of bounded independence with efficient computation of the conditional expectation. Via a reduction to the maximal independent set (MIS) problem, they gave a $(\Delta+1)$-coloring algorithm that runs in $O(\log \Delta)$ rounds provided that $\Delta = O(\mathfrak{n}^{1/3})$. Parter [18] showed a deterministic $(\Delta + 1)$-coloring in $O(\log \Delta)$ rounds, that works for any value of Δ. Barenboim and Khazanov [2] presented improved deterministic local algorithms as a function of the graph's *arboricity*. Finally, concurrently with our work, Bamberger et al. [1] give $(\deg +1)$-list coloring algorithms requiring $O(\log \Delta \log \log \Delta)$ rounds in CONGESTED CLIQUE, $O(\log^2 \Delta)$ in linear-space MPC, and $O(\log^2 \Delta + \log \mathfrak{n})$ in low-space MPC. These bounds are significantly weaker than ours, but we note that the low-space MPC algorithm of [1] has the advantage of using optimal $O(\mathfrak{m} + \mathfrak{n})$ global space, compared to our $O(\mathfrak{m} + \mathfrak{n}^{1+\varepsilon})$, and that for CONGESTED CLIQUE and linear-space MPC we present algorithms only for $(\Delta + 1)$-list coloring, rather than $(\deg +1)$-list coloring.

Prior to this, we are aware of no existing deterministic low-space MPC algorithms for the $(\Delta + 1)$-list coloring problem, though for the special case of $(\Delta + 1)$-coloring one can simulate the PRAM algorithm of Han [12] in $\tilde{O}(\log^2 \mathfrak{n})$ rounds.

1.4 Our approach

We present a simple, deterministic, constant-round CONGESTED CLIQUE algorithm for $(\Delta+1)$-list coloring, by carefully derandomizing a recursive partitioning procedure similar to those used by Parter [18] and Chang et al. [6].

The first step in our approach is to show a randomized procedure which maps nodes of the input graph and colors from nodes' palettes into bins. We show that if we leave only one of the bins without colors, most nodes in the remaining bins receive more colors from their palette into their bin than neighbors. This means that the graphs induced by the nodes in each bin can be colored recursively; we must also show that we can separately color the nodes in the bin which did not receive colors, and the nodes which did not receive more colors than neighbors (which can happen with small probability).

Then, we show that this randomized binning procedure requires only bounded independence. This means a *small random seed* suffices to provide all of the necessary randomness. We can then apply a distributed implementation of the *method of conditional expectations* in order to *deterministically* select such a seed that performs well.

The final step is to show that after each recursive application of this procedure, we significantly reduce the size of the graph induced by the nodes in each bin. In fact, after only *nine* recursive calls, we reduce the size of each instance from $O(\mathfrak{n}\Delta)$ to $O(\mathfrak{n})$, at which point it is small enough to simply collect onto a single machine and solve locally.

In low-space MPC only this final step must be altered; since now machines have only $O(\mathfrak{n}^\varepsilon)$ space, we cannot solve instances locally and instead apply a reduction to maximal independent set (MIS) in order to use an existing $O(\log \Delta + \log\log \mathfrak{n})$-round MIS algorithm of Czumaj et al. [7].

2 Preliminaries

In this section we detail some existing fundamentals on which our work relies.

2.1 Communication in **CONGESTED CLIQUE** and **MPC**

As noted in [7], prior work on MapReduce, a precursor to MPC, facilitates efficient deterministic communication in MPC, so long as the information sent and received by each machine adheres to the relevant space bounds of the model. In particular, the following results from [11] provide all the communication primitives we require (and use of features of MapReduce which are common to MPC):

Lemma 5 ([11], as stated in [7]). *For any positive constant δ, sorting and computing prefix sums of $\mathfrak{n}$ numbers can be performed deterministically in MapReduce (and therefore in MPC and CONGESTED CLIQUE) in $O(1)$ rounds using $\mathfrak{s} = \mathfrak{n}^\delta$ space per machine and $O(\mathfrak{n})$ total space.*

The computation of prefix sums here means the following: each machine $m \in [\mathfrak{M}]$ holds an input value x_m, and outputs $\sum_{i=1}^m x_i$.

This result essentially allows us to perform all of the communication we will need to do in a constant number of rounds. For example, if for each edge $\{u, v\}$ we create two entries (u, v) and (v, u) in memory, and the sort these lexicographically, we can ensure that the neighborhoods of all nodes are stored on contiguous blocks of machines. Then, by computing prefix sums, we can compute sums of values among a node's neighborhood, or indeed over the whole graph. This allows us to, e.g., compute objective functions for the method of conditional expectations (see Section 2.4). Where 2-hop neighborhoods fit in the space of a single machine, we can collect them by sorting edges to collect 1-hop neighborhoods onto machines, and then having each such machine send requests for the neighborhoods of all the nodes it stores.

An important point to note is that since Lemma 5 uses $\mathfrak{s} = n^\delta$ for *any* positive constant δ, by setting δ sufficiently smaller than our space parameter ε, we can perform $n^{\Omega(1)}$ *simultaneous*

sorting or prefix sum procedures. This will be especially useful to us for efficiently performing the method of conditional expectations.

The CONGESTED CLIQUE model contains the additional restriction of an $O(\log \mathfrak{n})$-bit bound on the information any node can send to any other in a round (as opposed to MPC, where only the *total* amount of information a node sends, to all others, is bounded). This restriction is circumvented by Lenzen's routing algorithm [15], which in $O(1)$ rounds allows all information to be routed to the correct destinations so long as all nodes obey an $O(\mathfrak{n})$ bound on the *total* amount of messages they send and receive.

So, henceforth we need not concern ourselves with routing information to the correct machines; so long as all machines obey the total space bounds at all times, we can do so in a generic fashion using these existing results.

2.2 Bounded-independence random variables

To convert randomized algorithms into deterministic algorithms, we follow a classic derandomization recipe: we show that a randomized algorithm can (or be modified to) work using only *bounded independence*, which means that it requires only a short random *seed*, and then we apply the *method of conditional expectations* to efficiently search the space of random seeds and deterministically fix one that performs well.

The first step of this process, showing that a randomized algorithm works using only bounded-independence variables, will make heavy use of the following *concentration bound* for sums of such variables.

Lemma 6 (Lemma 2.2 of [4]). *Let $c \geq 4$ be an even integer. Suppose $Z_1, \ldots, Z_t$ are c-wise independent random variables taking values in $[0,1]$. Let $Z = Z_1 + \cdots + Z_t$ and $\mu = \mathbf{E}[Z]$, and let $\lambda > 0$. Then,*

$$\mathbf{Pr}\left[|Z - \mu| \geq \lambda\right] \leq 2\left(\frac{ct}{\lambda^2}\right)^{c/2} \ .$$

Here, c-wise independence means that any c of the variables $Z_1, \ldots, Z_t$ can be considered to be independent, but groups of *more* than c cannot.

2.3 Families of bounded-independence hash functions

To generate c-wise independent random variables from a short random seed, we use the concept of families of bounded-independence hash functions:

Definition 7. *For $N, L, c \in \mathbb{N}$ such that $c \leq N$, a family of functions $\mathcal{H} = \{h : [N] \to [L]\}$ is c-wise independent if for all distinct $x_1, \ldots, x_c \in [N]$, the random variables $h(x_1), \ldots, h(x_c)$ are independent and uniformly distributed in $[L]$ when h is chosen uniformly at random from $\mathcal{H}$.*

That is, the outputs of a uniformly random function from a c-wise independent family are c-wise independent random variables. Such families of small size (and therefore requiring a small seed to specify any particular element) are known to exist; we will use the following well-known lemma (see, e.g., Corollary 3.34 in [21]):

Lemma 8. *For every a, b, c, there is a family of c-wise independent hash functions $\mathcal{H} = \{h : \{0,1\}^a \to \{0,1\}^b\}$ such that choosing a random function from $\mathcal{H}$ takes $c \cdot \max\{a, b\}$ random bits, and evaluating a function from $\mathcal{H}$ takes $poly(a, b, c)$ computation.*

In particular, we will be using families $\mathcal{H}$ of hash functions with $\mathfrak{n}^{O(1)}$ range and domain, and some large constant c, so by Lemma 8, a uniformly random function from $\mathcal{H}$ can be specified using $O(\log \mathfrak{n})$ random bits (and this string of bits is what we call its *seed*). This also means that functions can be evaluated using polylogarithmic computation, and therefore space, which is possible locally even in low-space MPC.

The domain and range of the hash functions must be powers of 2, and therefore we incur some error when our desired range is not (for domain we simply ignore any excess). If, for example, we have some stated range r, we will actually use range $\{0,1\}^{\lceil \log(r\mathfrak{n}^3) \rceil}$, and then map intervals of this range as equally as possible (i.e., differing in size by at most 1) to $[r]$. We then have error in our resulting probabilities of $O(\mathfrak{n}^{-3})$, which for our purposes will be clearly negligible. Note, too, that outputs of our 'simulated' hash function with range r are still exactly c-wise independent (and therefore we can still use Lemma 6), but are no longer exactly uniform.

2.4 Method of conditional expectations

To deterministically choose a good hash function from our small families, we employ the classical *method of conditional expectations* (see, e.g., [21, Chapter 3.4] or [7, Section 2.4]). In our context, we will show that, over the choice of a uniformly random hash function $h \in \mathcal{H}$, the expectation of some cost function q (which is a sum of functions q_x calculable by individual machines) is at most some value Q. That is,

$$\mathbf{E}_{h \in \mathcal{H}}\left[q(h) := \sum_{\text{machines } x} q_x(h)\right] \leq Q \ .$$

The probabilistic method implies the existence of a hash function $h^* \in \mathcal{H}$ for which $q(h^*) \leq Q$, and our goal is to find one such $h^* \in \mathcal{H}$ in $O(1)$ rounds of CONGESTED CLIQUE or low-space MPC.

We will find the sought hash function h^* by deterministically fixing the $O(\log \mathfrak{n})$-bit seed that specifies it (cf. Lemma 8). To do so, we have all machines agree gradually on chunks of $\delta \log \mathfrak{n}$ bits at a time, for sufficiently small constant δ. That is, we iteratively extend a fixed, agreed, prefix h_{pre} of the final hash function h^*, $\delta \log \mathfrak{n}$ bits at a time, until we have fixed the entire seed.

For each chunk, and for each i, $1 \leq i \leq \mathfrak{n}^\delta$, we write i as a $\delta \log \mathfrak{n}$-length bit string. Then, each machine x calculates $\mathbf{E}_{h \in \mathcal{H}}\left[q_x(h)|\Xi_{h_{pre}+i}\right]$, where $\Xi_{h_{pre}+i}$ is the event that the seed specifying h is prefixed by the current fixed prefix h_{pre}, followed by i. We then sum these values over all machines for each i concurrently (using Lemma 5 and choosing δ sufficiently smaller than the space parameter ε), obtaining $\mathbf{E}_{h \in \mathcal{H}}\left[q(h)|\Xi_{h_{pre}+i}\right]$. By the probabilistic method, at least one of these values is at most $\mathbf{E}_{h \in \mathcal{H}}\left[q(h)|\Xi_{h_{pre}}\right]$. We fix a value of i such that this is the case, append this to h_{pre}, and continue.

After a constant number of iterations, we find the entire string of bits to define a hash function $h^* \in \mathcal{H}$. Since in every iteration we did not increase the conditional expectation $\mathbf{E}_{h \in \mathcal{H}}\left[q(h)|\Xi_{h_{pre}}\right]$, we have

$$q(h^*) \leq \mathbf{E}_{h \in \mathcal{H}}\left[q(h)\right] \leq Q \ .$$

Each iteration requires only a constant number of rounds in CONGESTED CLIQUE or MPC, so this process takes only a constant number of rounds in total.

3 Coloring in **CONGESTED CLIQUE** and Linear-Space **MPC**

We are now ready to present our constant-round algorithm for $(\Delta + 1)$-list coloring in CON-GESTED CLIQUE. Our algorithm will be a recursive procedure, which we call COLORREDUCE. To color an input graph $\mathfrak{G}$ we will call COLORREDUCE$(\mathfrak{G}, \Delta)$ (and in subsequent calls, ℓ will serve as an approximation of maximum degree). We will denote by $\mathfrak{n}$ the number of nodes in $\mathfrak{G}$, and when analyzing a subsequent recursive COLORREDUCE(G, ℓ) we denote by n_G the number of nodes in G.

Algorithm 1 COLORREDUCE(G, ℓ)

If G has size $O(\mathfrak{n})$: collect G onto a single machine and color G locally.
Otherwise: $G_0, \ldots, G_{\ell^{0.1}} \leftarrow$ PARTITION(G, ℓ).
Let $\ell' = \ell^{0.9} - \ell^{0.6}$.
For each $i = 1, \ldots, \ell^{0.1} - 1$, perform COLORREDUCE$(G_i, \ell')$ in parallel.
Update color palettes of $G_{\ell^{0.1}}$, perform COLORREDUCE$(G_{\ell^{0.1}}, \ell')$.
Update color palettes of G_0, collect G_0 onto a single machine and color locally.

The algorithm relies on a procedure for partitioning a coloring instance into multiple instances of lower size and lower maximum degree. Specifically, we partition the node set into $\ell^{0.1}$ bins, and then partition the colors into all but the final bin (and restrict nodes' palettes to the colors sent to their bin). Next, we show that (after possibly removing some *bad nodes*, to be defined shortly) each node (not in the final bin) has more colors from its palette assigned to its bin than neighbors. This means that we can recursively color the graphs induced by these bins in parallel, since the color palettes of any nodes from different bins are now completely disjoint. Then, we must afterwards color the final bin ℓ which was not previously assigned colors, and the graph of bad nodes. For that, in the last two steps in COLORREDUCE, before we recursively color $G_{\ell^{0.1}}$ and then color G_0, we first update color palettes of all nodes in $G_{\ell^{0.1}}$ and G_0, respectively, by removing from the palettes of every node v the colors used by all its neighbors (in $G_0, \ldots, G_{\ell^{0.1}}$) already colored before.

Algorithm 2 PARTITION(G, ℓ)

Let hash function $h_1 : [\mathfrak{n}] \to [\ell^{0.1}]$ map nodes v to a bin $h_1(v) \in [\ell^{0.1}]$.
Let hash function $h_2 : [\mathfrak{n}^2] \to [\ell^{0.1} - 1]$ map colors γ to a bin $h_2(\gamma) \in [\ell^{0.1} - 1]$.
Let G_0 be the graph induced by *bad* nodes.
Let $G_1, \ldots, G_{\ell^{0.1}}$ be the graphs induced by *good* nodes in bins $1, \ldots, \ell^{0.1}$ respectively.
Restrict palettes of nodes in $G_1, \ldots, G_{\ell^{0.1}-1}$ to colors assigned by h_2 to corresponding bins.
Return $G_0, \ldots, G_{\ell^{0.1}}$.

Since we consider list coloring, the total number of unique colors can be up to $\mathfrak{n}^2$, which necessitates the larger domain for h_2.

To fully specify the algorithm, we must define what we mean by *good* and *bad nodes* and by *good* and *bad bins*. Let $d(v)$ denote degree of v in G, and $d'(v)$ denote degree of v within its new bin $h_1(v)$, that is, in the subgraph of G induced by all nodes $u \in V$ with $h_1(u) = h_1(v)$. Let $p(v)$ be the size of the color palette of node v, and $p'(v)$ be the palette size of v after the call: for nodes mapped to bins in $[\ell^{0.1} - 1]$, $p'(v)$ denotes the number of colors from the color palette of v that are hashed to the bin $h_1(v)$; for nodes mapped to bin $\ell^{0.1}$, $p'(v)$ denotes the number of colors remaining from the color palette of v *after* v's neighbors *not* in bin $\ell^{0.1}$ (that is, in $G_1, \ldots, G_{\ell^{0.1}-1}$) have been colored by COLORREDUCE.

Definition 9 (**Good and bad nodes and bins**)**.**

8

- *A node v in bins $[\ell^{0.1} - 1]$ is good if $|d'(v) - d(v)\ell^{-0.1}| \leq \ell^{0.6}$ and $p'(v) \geq p(v)\ell^{-0.1} + \ell^{0.7}$.*

- *A node v in bin $\ell^{0.1}$ is good if $|d'(v) - d(v)\ell^{-0.1}| \leq \ell^{0.6}$.*

- *A bin in $[\ell^{0.1}]$ is good if it contains fewer than $2n_G\ell^{-0.1} + \mathfrak{n}^{0.6}$ nodes.*

- *Nodes and bins are bad if they are not good.*

Intuitively, bad nodes and bins are those whose behavior differs significantly from what we would expect in a random partition of G into $\ell^{0.1}$ node-disjoint graphs, with random h_1 and h_2.

Overview of the analysis of ColorReduce. In order to analyze the complexity of Col-orReduce, we have to understand how to implement each single call to Partition, and how often we will be making the recursive calls. Further, in order the algorithm to work correctly, we have to ensure that every time we collect G_0 (containing bad nodes) onto a single machine in order to color it, the total size of G_0 is $O(\mathfrak{n})$ (since otherwise G_0 would not fit a single machine).

In Section 3.1, we describe a useful invariant for Partition call that is maintained throughout the entire run of ColorReduce. In Section 3.2, we estimate the number of bad nodes and bins for random hash functions h_1 and h_2 from appropriate families of hash functions. In Section 3.3, we describe how to deterministically choose hash functions h_1 and h_2 (from the same families of hash functions) so that these estimations are still ensured, and that we guarantee G_0 is sufficiently small. Next, in Section 3.4, we estimate the number of recursive calls performed in the call to ColorReduce(G, Δ) and subsequent recursive calls. Our main result (cf. Lemma 22) is that after a depth of 9 recursive calls, the graph induced by each bin is of size $O(\mathfrak{n})$. This implies that the call to ColorReduce$(\mathfrak{G}, \Delta)$ creates a recursion tree of depth at most 9, resulting in a constant number of groups of instances of size $O(\mathfrak{n})$ to be collected onto machines and colored sequentially. With these claims, in Section 3.5 we finalize our analysis of Theorem 1, and show that ColorReduce(G, Δ) deterministically returns $(\Delta + 1)$-list coloring in $O(1)$ rounds in CONGESTED CLIQUE and in $O(1)$ rounds on MPC with $O(\mathfrak{n})$ local space on any machine, and (optimal for $(\Delta + 1)$-list coloring) $O(\mathfrak{n}\Delta)$ total space. Then, in Section 3.6, we reduce the global MPC space for the $(\Delta + 1)$-coloring problem to the optimal bound of $O(\mathfrak{m} + \mathfrak{n})$.

3.1 Invariant maintained during all calls to Partition

We begin our analysis with presenting a central invariant of our Partition call that is maintained throughout the entire run of ColorReduce (recall that $\ell' = \ell^{0.9} - \ell^{0.6}$).

Lemma 10. *Assume that at the start of our Partition call, the following holds for all nodes v in G: (i) $\ell < p(v)$, (ii) $d(v) \leq \ell + \ell^{0.7}$, (iii) $d(v) < p(v)$. Then for all good nodes v in G,*

(i) $\ell' < p'(v)$,

(ii) $d'(v) \leq \ell' + \ell'^{0.7}$, and

(iii) $d'(v) < p'(v)$.

Note that since we assume $d(v) \leq \ell + \ell^{0.7}$, we can also assume ℓ is at least a sufficiently large constant, since otherwise maximum degree is also bounded by a constant, so the instance is of total size $O(\mathfrak{n})$ and would have already been collected and colored on a single machine.

Proof. Let us first consider any good node v with $h_1(v) \in [\ell^{0.1} - 1]$. Then, $d'(v) \leq d(v)\ell^{-0.1} + \ell^{0.6}$ and $p'(v) \geq p(v)\ell^{-0.1} + \ell^{0.7}$, by Definition 9. Therefore,

(i) since $\ell < p(v)$, we have $p'(v) \geq p(v)\ell^{-0.1} + \ell^{0.7} > p(v)\ell^{-0.1} > \ell^{0.9} > \ell'$;

(ii) since $d(v) \leq \ell + \ell^{0.7}$, we have $d'(v) \leq d(v)\ell^{-0.1} + \ell^{0.6} \leq (\ell^{0.9} + \ell^{0.6}) + \ell^{0.6} < \ell' + \ell'^{0.7}$;

(iii) since $d(v) < p(v)$, we have $p'(v) \geq p(v)\ell^{-0.1} + \ell^{0.7} > d(v)\ell^{-0.1} + \ell^{0.6} \geq d'(v)$.

Next, let us consider any good node v with $h_1(v) = \ell^{0.1}$. Then, by Definition 9, it holds $|d'(v) - d(v)\ell^{-0.1}| \leq \ell^{0.6}$. Further, notice that we update the color palette of v (and hence set $p'(v)$) only after coloring in ColorReduce all v's neighbors in $G_1, \ldots, G_{\ell^{0.1}-1}$. Therefore,

(i) since at most one color is removed from v's palette for each neighbor of v *not* in bin $\ell^{0.1}$, since $|d'(v) - d(v)\ell^{-0.1}| \leq \ell^{0.6}$, $d(v) < p(v)$, and $\ell < p(v)$, we have,

$$p'(v) \geq p(v) - (d(v) - d'(v)) \geq p(v) - (d(v) - (d(v)\ell^{-0.1} + \ell^{0.6})) = p(v) - d(v)(1 - \ell^{-0.1}) - \ell^{0.6}$$
$$> p(v) - p(v)(1 - \ell^{-0.1}) - \ell^{0.6} = p(v)\ell^{-0.1} - \ell^{0.6} > \ell^{0.9} - \ell^{0.6} = \ell' \ ;$$

(ii) as before: since $d(v) \leq \ell + \ell^{0.7}$, we have $d'(v) \leq d(v)\ell^{-0.1} + \ell^{0.6} \leq (\ell^{0.9} + \ell^{0.6}) + \ell^{0.6} < \ell' + \ell'^{0.7}$;

(iii) $d'(v) < p'(v)$ follows since $d(v) < p(v)$, and colors are only removed from v's palette when used by neighbors. $\square$

Notice that at the very beginning, with our input graph G of maximum degree Δ and color palettes satisfying $d(v) < p(v)$ for all v, our initial setting $\ell = \Delta$ ensures that before the first call of Partition all conditions of Lemma 10 (for all nodes v: (i) $\ell < p(v)$, (ii) $d(v) \leq \ell + \ell^{0.7}$, (iii) $d(v) < p(v)$) are satisfied, and hence, since Partition is only called on good nodes, Lemma 10 ensures that the properties in the invariant are preserved for future calls of Partition.

Corollary 11. *The input instance to any call to* Partition *satisfies the following three conditions for all nodes v: (i) $\ell < p(v)$, (ii) $d(v) \leq \ell + \ell^{0.7}$, (iii) $d(v) < p(v)$.*

3.2 Good and bad nodes and bins for *random* hash functions

In this section we analyze the likelihood of having many bad nodes and bins in a single call to Partition, assuming that h_1 and h_2 are hash functions *chosen at random* from the families of c-wise independent hash functions. (In Section 3.3 we will extend our analysis and show how to deterministically choose appropriate h_1 and h_2.)

Let $\mathcal{H}_1$ be a c-wise independent family of hash functions $h : [\mathfrak{n}] \to [\ell^{0.1}]$, and let $\mathcal{H}_2$ be a c-wise independent family of hash functions $h : [\mathfrak{n}^2] \to [\ell^{0.1} - 1]$, for sufficiently large constant c.

In order to understand basic properties of pairs of hash functions from $\mathcal{H}_1$ and $\mathcal{H}_2$, we define the *cost* function $\mathfrak{q}(h_1, h_2)$ of a pair of hash functions $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$, as follows:

$$\mathfrak{q}(h_1, h_2) = |\{\text{bad nodes under } h_1, h_2\}| + \mathfrak{n} \cdot |\{\text{bad bins under } h_1\}| \ . \tag{1}$$

We will analyze the likelihood of having many bad nodes and bins, assuming that h_1 and h_2 are chosen at random from the families of c-wise independent hash functions $\mathcal{H}_1$ and $\mathcal{H}_2$, and relying on Corollary 11 that at the start of our Partition for all nodes v: (i) $\ell < p(v)$, (ii) $d(v) \leq \ell + \ell^{0.7}$, (iii) $d(v) < p(v)$.

Good and bad bins. We will first prove that with high probability, no bins are bad.

Lemma 12. *With probability at least $1 - \mathfrak{n}^{-2}$, all bins are good.*

Proof. We apply Lemma 6 with $Z_1, \ldots, Z_{n_G}$ as indicator random variables for the events that each node is hashed to a particular bin $b \in [\ell^{0.1}]$. These variables are $(c-1)$-wise independent and each have expectation $\ell^{-0.1}$. Since we assume c to be a sufficiently high constant, by Lemma 6, we obtain $\mathbf{Pr}\left[|Z - \mathbf{E}\left[Z\right]| \geq n^{0.6}\right] \leq n^{-3}$. By the union bound over bins b, with probability at least $1 - \mathfrak{n}^{-2}$ all bins contain fewer than $n_G \ell^{-0.1} + \mathfrak{n}^{0.6}$ nodes, and are therefore good. $\square$

Good and bad nodes. In order to estimate the number of bad nodes, we begin with a proof that the probability that nodes' degrees differ too much from what would be expected (from a fully random partition) is low.

Lemma 13. *The probability that a node v satisfies $|d'(v) - d(v)\ell^{-0.1}| \geq \ell^{0.6}$ is at most ℓ^{-3}.*

Proof. We apply Lemma 6 with $Z_1, \ldots, Z_{d(v)}$ as indicator random variables for the events that each of v's neighbors is hashed to the same bin as v. These variables are $(c-1)$-wise independent and each have expectation $\ell^{-0.1}$. Since we assume c to be a sufficiently high constant and (cf. Corollary 11) have $d(v) \leq \ell + \ell^{0.7}$, by Lemma 6, we obtain $\mathbf{Pr}\left[|Z - \mathbf{E}\left[Z\right]| \geq \ell^{0.6}\right] \leq \ell^{-3}$. Therefore, since $\mathbf{E}\left[Z\right] = d(v) \cdot \ell^{-0.1}$ and $Z = d'(v)$, we can conclude the required claim. $\square$

We now prove that the probability that nodes in bins $1, \ldots, \ell^{0.1} - 1$ do not receive enough colors to color themselves is low.

Lemma 14. *The probability that a node v with $h_1(v) \neq \ell^{0.1}$ satisfies $p'(v) \leq p(v)\ell^{-0.1} + \ell^{0.7}$ is at most ℓ^{-3}.*

Proof. We apply Lemma 6, with $Z_1, \ldots, Z_{p(v)}$ as indicator random variables for the events that each color in v's palette is hashed to its bin. These variables are $(c-1)$-wise independent and each have expectation $\frac{1}{\ell^{0.1}-1}$. Since we assume c to be a sufficiently high constant, by Lemma 6, we obtain $\mathbf{Pr}\left[|Z - \mathbf{E}\left[Z\right]| \geq p(v)^{0.6}\right] \leq p(v)^{-3}$. Therefore, since $Z = p'(v)$ and $\mathbf{E}\left[Z\right] = \frac{p(v)}{\ell^{0.1}-1}$, the probability that $p'(v) \leq \frac{p(v)}{\ell^{0.1}-1} - p(v)^{0.6}$ is at most $p(v)^{-3}$. Furthermore, since $\ell < p(v)$ (cf. Corollary 11), we get,

$$\frac{p(v)}{\ell^{0.1} - 1} - \frac{p(v)}{\ell^{0.1}} = \frac{p(v)}{\ell^{0.1}(\ell^{0.1} - 1)} \geq \frac{p(v)}{\ell^{0.2}} \geq \frac{1}{2}\left(\frac{p(v)}{p(v)^{0.2}} + \frac{\ell}{\ell^{0.2}}\right) > p(v)^{0.6} + \ell^{0.7} \ .$$

Hence, with probability at least $1 - p(v)^{-3} \geq 1 - \ell^{-3}$ we obtain the following,

$$p'(v) > \frac{p(v)}{\ell^{0.1} - 1} - p(v)^{0.6} > \frac{p(v)}{\ell^{0.1}} + \ell^{0.7} \ . \qquad \square$$

We can now combine Lemmas 13 and 14 to bound the probability that nodes are bad:

Lemma 15. *A node is bad with probability at most $2\ell^{-3}$.*

Proof. For a node v to be bad, at least one of the bad events described by Lemmas 13 and 14 must occur. By the union bound, the probability of this is at most $2\ell^{-3}$. $\square$

This provides a bound on the cost of random hash function pairs:

Lemma 16. *The expected cost of a random hash function pair from $\mathcal{H}_1 \times \mathcal{H}_2$ is at most $\frac{\mathfrak{n}}{\ell^2}$.*

Proof. Combining Lemmas 12 and 15 gives the following:

$$\mathbf{E}\left[\mathfrak{q}(h_1, h_2)\right] = \mathbf{E}\left[|\{\text{bad nodes}\}| + \mathfrak{n} \cdot |\{\text{bad bins}\}|\right] \leq 2\mathfrak{n}\ell^{-3} + \mathfrak{n} \cdot \mathfrak{n}^{-2} \leq \frac{\mathfrak{n}}{\ell^2} \ . \qquad \square$$

3.3 Choosing appropriate hash functions

In this section, we discuss how to deterministically choose hash functions $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$ to ensure the desired properties. First, let us recall that by Lemma 8, we can assume that families of hash functions $\mathcal{H}_1$ and $\mathcal{H}_2$ have size $\mathfrak{n}^{O(1)}$. Next, by Lemma 16, upon choosing a random pair of hash functions $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$, we have $\mathbf{E}\left[\mathfrak{q}(h_1, h_2)\right] \leq \frac{\mathfrak{n}}{\ell^2}$. Therefore, by the method of conditional expectations (cf. Section 2.4), we can find a pair of hash functions $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$ with cost $\mathfrak{q}(h_1, h_2) \leq \frac{\mathfrak{n}}{\ell^2}$ in $O(1)$ rounds. By (1), this ensures that there are no bad bins and that there are at most $\frac{\mathfrak{n}}{\ell^2}$ bad nodes.

In order to implement the method of conditional expectations to find such an $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$, we require that $\mathfrak{q}(h_1, h_2)$ is the sum of functions computable by individual machines. To achieve this, we will distribute information between the machines as follows: each node v will be assigned a machine x_v, which will store all of its adjacent edges and its palette. The machine will then be able to determine whether a particular pair of hash functions h_1, h_2 results in the node being *bad*, since this depends only on v's palette, neighbor IDs, and the hash functions. The local cost function $q_{x_v}(h_1, h_2)$ will then simply be an indicator variable for the event that v is *bad* under h_1 and h_2.

We will also assign a special machine x^* to check if any bins are bad (contain at least $2n_G \ell^{-0.1} + \mathfrak{n}^{0.6}$ nodes) for any hash function h_1, for which it needs only to store the IDs of each node. The local cost function $q_{x^*}(h_1, h_2)$ is the number of bad bins under h_1, multiplied by $\mathfrak{n}$, which is computable by x^*.

The global cost function $\mathfrak{q}(h_1, h_2) = |\{\text{bad nodes under } h_1, h_2\}| + \mathfrak{n} \cdot |\{\text{bad bins under } h_1\}|$ is then the sum of these local cost functions computable by individual machines. If we in fact use machines for multiple nodes where space allows, then (even among all concurrent instances) we will require only $O(1 + \frac{\mathfrak{m}}{\mathfrak{n}})$ machines, i.e., total space of $O(\mathfrak{m} + \mathfrak{n})$.

Lemma 17. *In $O(1)$ MPC rounds with $O(\mathfrak{n})$ local space per machine and $O(\mathfrak{m} + \mathfrak{n})$ total space, one can select hash functions $h_1 \in \mathcal{H}_1$, $h_2 \in \mathcal{H}_2$ such that in a single call to* PARTITION,

- *there are no bad bins, and*

- *there are at most $\frac{\mathfrak{n}}{\ell^2}$ bad nodes.*

This latter property allows us to bound the size of G_0, the graph induced by bad nodes encountered in any recursive call:

Corollary 18. *Consider an arbitrary call of* COLORREDUCE(G, ℓ). *Then G_0, the graph induced by bad nodes in* PARTITION(G, ℓ), *is of size $O(\mathfrak{n})$.*

Proof. By Lemma 17, in the call to PARTITION(G, ℓ), G_0 has at most $\frac{\mathfrak{n}}{\ell^2}$ nodes. Furthermore, by Corollary 11 (ii), the degree of each node in the call to PARTITION(G, ℓ) is at most $\ell + \ell^{0.7} < \ell^2$. Therefore, the graph of the bad nodes is of size at most $\frac{\mathfrak{n}}{\ell^2} \cdot \ell^2 = O(\mathfrak{n})$, and can be collected onto a single MPC machine and properly colored. $\qquad\square$

3.4 Recursive calls to COLORREDUCE and recursion depth

Our algorithm is recursive, and it starts with the call to COLORREDUCE$(\mathfrak{G}, \Delta)$ for the input graph $\mathfrak{G}$ of maximum degree Δ, where each node v has assigned a color palette of size $p(v)$, $p(v) > d(v)$. Starting with our initial call on a graph with $n_0 = \mathfrak{n}$ nodes and $\ell_0 = \Delta$, we will upper bound the number of nodes and degree of graphs in subsequent recursive calls of depth i. We will denote these values by n_i and Δ_i respectively, and by ℓ_i the value of ℓ in a recursive call of depth i.

We begin with the following bound for the values of ℓ_i in a recursive call of depth i.

Lemma 19. $\frac{1}{2}\Delta^{0.9^i} < \ell_i \le \Delta^{0.9^i}$.

Proof. The proof is by induction on i. Clearly the claim is true for $\ell_0 = \Delta$.

Assuming that the claim holds for $i = j$, for $i = j + 1$ we have the following,

$$\ell_{j+1} = \ell_j^{0.9} - \ell_j^{0.6} \le \ell_j^{0.9} \le \Delta^{0.9 \cdot 0.9^j} = \Delta^{0.9^{j+1}}$$

and (assuming Δ is greater than a sufficiently large constant)

$$\ell_{j+1} = \ell_j^{0.9} - \ell_j^{0.6} > \frac{1}{2^{0.9}}\Delta^{0.9 \cdot 0.9^j} - \Delta^{0.6 \cdot 0.9^j} > \frac{1}{2}\Delta^{0.9^{j+1}} \quad . \qquad \square$$

Next, we bound the number of nodes n_i in subsequent recursive calls of depth i.

Lemma 20. $n_i \le 3^i(\mathfrak{n}\Delta^{0.9^i - 1} + \mathfrak{n}^{0.6})$.

Proof. The proof is by induction on i. As a base case, $n_0 = \mathfrak{n} \le 3^0(n\Delta^{0.9^0 - 1} + \mathfrak{n}^{0.6}) = \mathfrak{n} + \mathfrak{n}^{0.6}$.

Assume that the claim holds for $i = j$. Then for $i = j + 1$, by induction, by the definition of good bins (cf. Definition 9), and by Lemma 19, we obtain

$$n_{j+1} \le 2n_j\ell_j^{-0.1} + \mathfrak{n}^{0.6} \le 2\left(3^j(\mathfrak{n}\Delta^{0.9^j - 1} + \mathfrak{n}^{0.6})\right)\left(2^{0.1}\Delta^{-0.1 \cdot 0.9^j}\right) + \mathfrak{n}^{0.6}$$

$$\le 3^{j+1}\mathfrak{n}\Delta^{0.9^j - 1 - 0.1 \cdot 0.9^j} + 3^{j+1}\mathfrak{n}^{0.6} = 3^{j+1}\mathfrak{n}\Delta^{0.9^{j+1} - 1} + 3^{j+1}\mathfrak{n}^{0.6} \quad . \qquad \square$$

We can similarly prove a bound on degree of graphs in recursive calls of depth i.

Lemma 21. $\Delta_i \le 2^i\Delta^{0.9^i}$.

Proof. The proof is by induction on i. As a base case, $\Delta_0 = \Delta = 2^0\Delta^{0.9^0}$.

Assume the claim holds for $i = j$. By the definition of good bins (cf. Definition 9), we have $\Delta_{j+1} \le \Delta_j\ell_j^{-0.1} + \ell_j^{0.6}$, and if we combine this bound with Lemma 19, by induction we obtain,

$$\Delta_{j+1} \le \Delta_j\ell_j^{-0.1} + \ell_j^{0.6} \le \left(2^j\Delta^{0.9^j}\right)\left(2^{0.1}\Delta^{-0.1 \cdot 0.9^j}\right) + \Delta^{0.6 \cdot 0.9^j} < 2^{j+1}\Delta^{0.9^{j+1}} \quad . \qquad \square$$

Let us first give a simple bound for the recursion depth of our algorithm.

Lemma 22. *After a depth of 9 recursive calls, the graph induced by each bin is of size $O(\mathfrak{n})$.*

Proof. By Lemmas 20 and 21, the total size of any graph G' induced by any bin after a depth of i recursive calls satisfies the following,

$$|G'| \le n_i \cdot \Delta_i < 3^i(\mathfrak{n}\Delta^{0.9^i - 1} + \mathfrak{n}^{0.6}) \cdot 2^i\Delta^{0.9^i} = 6^i(\mathfrak{n}\Delta^{0.9^i - 1} + \mathfrak{n}^{0.6}) \cdot \Delta^{0.9^i} \quad .$$

Plugging in $i = 9$, we obtained the following bound for the total size of any graph G' induced by any bin after a depth of 9 recursive calls,

$$|G'| < 6^9(\mathfrak{n}\Delta^{0.9^9 - 1} + \mathfrak{n}^{0.6}) \cdot \Delta^{0.9^9} < 6^9(\mathfrak{n}\Delta^{0.4 - 1} + \mathfrak{n}^{0.6}) \cdot \Delta^{0.4} \le 6^9(\mathfrak{n}\Delta^{-0.2} + \mathfrak{n}) = O(\mathfrak{n}) \quad . \quad \square$$

3.5 Finalizing the analysis for CONGESTED CLIQUE: Proof of Theorem 1

Now we are ready to complete our analysis and prove our main theorem, Theorem 1.

Proof of Theorem 1. COLORREDUCE$(\mathfrak{G}, \Delta)$ creates a recursion tree of depth at most 9, resulting in $O(1)$ groups of instances of size $O(\mathfrak{n})$ to be collected onto machines and colored sequentially. All operations within each recursive call (including coloring graphs G_0, see Corollary 18) can be implemented in a constant number of rounds in CONGESTED CLIQUE and on MPC with $O(\mathfrak{n})$ local space on any machine (and $O(\mathfrak{n}\Delta)$ total space), so the total number of rounds required is also constant. $\qquad \square$

3.6 Note on optimal global space for $(\Delta + 1)$-coloring (Theorem 3)

For the special case of $(\Delta + 1)$-coloring, where initial palettes are all $[\Delta + 1]$ and therefore need not be given as input, we remark that explicitly maintaining palettes for all nodes costs $\Theta(\mathfrak{n}\Delta)$ global space, which can exceed the optimal global space bound of $O(\mathfrak{m} + \mathfrak{n})$. We address this issue as follows:

We require palettes in two places: when applying PARTITION, and when coloring instances of size $O(\mathfrak{n})$ locally. In this latter case, we can afford to arbitrarily drop colors from a node v's palette until its palette size is $d(v) + 1$, so the space used in total is $O(\mathfrak{m})$.

In the former case, consider applying PARTITION to an instance with n nodes and m edges. Nodes will be distributed over $O(1 + \frac{m}{\mathfrak{n}})$ machines in order to apply the method of conditional expectations, for which they require palette access.

Notice that during the run of our algorithm, node palettes are updated in two ways: by restriction based on partition by some hash function, and by removing colors when they are used to color neighbors. We will maintain the palette of each node v by explicitly storing colors used by its neighbors, and implicitly storing colors removed during partitioning by storing the chosen hash function. In this way, nodes' palettes are fully specified and can be queried. For any node v, we explicitly store up to one color per neighbor, which requires $d(v)$ space for v and therefore total $O(\mathfrak{m})$ for all nodes.

So, it remains to analyze the space cost of storing the necessary hash functions. In our PARTITION call on n nodes and m edges, we have $O(1 + \frac{m}{\mathfrak{n}})$ machines which must each store the $O(1)$ hash functions which have previously been applied to this instance. These hash functions are specified by $O(\log \mathfrak{n})$ bits, so the total space required is $O(\log \mathfrak{n} + \frac{m \log \mathfrak{n}}{\mathfrak{n}})$.

We may run up to $O(\sqrt{\mathfrak{n}})$ calls to PARTITION concurrently (we cannot have more, because we must have $m = \Omega(\mathfrak{n})$, so $n = \Omega(\sqrt{\mathfrak{n}})$). Furthermore, each edge in $\mathfrak{G}$ is present in at most one call. So, the total space used to store hash functions over *all* concurrent PARTITION calls is $O(\sqrt{\mathfrak{n}} \log \mathfrak{n} + \frac{m \log \mathfrak{n}}{\mathfrak{n}}) = O(\mathfrak{n} + \mathfrak{m})$, which yields Theorem 3.

4 Coloring in Low-Space MPC

We now extend our methods to the more restrictive regime of low-space MPC. Again, our algorithm will be a recursive procedure which relies on a derandomized partitioning of nodes and colors into smaller instances. However, reducing to instances of size $O(\mathfrak{n})$ is no longer sufficient to allow collection onto single machines for a constant-round solution. So, we instead reduce until degree is $\mathfrak{n}^{7\delta}$ for some sufficiently small constant δ, and then we can apply a reduction to maximal independent set (MIS) in order to use an existing algorithm of [7], requiring $O(\log \Delta + \log \log \mathfrak{n})$ rounds. This is the reason for the difference in round complexity compared to Algorithms 1, 2.

4.1 Reduction to MIS in low degree instances

When degree is $O(n^{7\delta})$, we apply the standard (due to Luby [17]) reduction to maximal independent set (MIS): a new graph is created in which nodes v in the original graph correspond to cliques of size $p(v)$ in the new graph, with each clique node representing a color from the palette of node v. If two neighboring original nodes share a color between their palettes, an edge is drawn between their corresponding clique nodes. An MIS in this new graph necessarily indicates a complete proper coloring of the original graph, since it can easily be seen that exactly one clique node from each clique joins the MIS, and this determines the color for the original node. When the reduction is applied to a graph with $\hat{\mathfrak{n}}$ vertices and maximum degree $\mathfrak{n}^{7\delta}$, the new reduction graph has at most $O(\hat{\mathfrak{n}} \cdot \mathfrak{n}^{7\delta})$ vertices, and maximum degree at most $\mathfrak{n}^{14\delta}$.

14

Then, we can solve $(\deg+1)$-list coloring by applying the MIS algorithm of [7] to this reduction graph. This algorithm requires $O(\log \Delta + \log \log \mathfrak{n})$ rounds, and can be applied to our reduction graphs using $O(\mathfrak{n}^\delta)$ and $O(\hat{\mathfrak{n}}^{1+\delta}\mathfrak{n}^{21\delta})$ local and global space respectively.

4.2 Main low-space MPC algorithm

We now present our main algorithm for coloring in low-space MPC, which recursively reduces instances until they have degree $O(\mathfrak{n}^{7\delta})$ and then applies the MIS reduction.

Algorithm 3 LowSpaceColorReduce(G)

$G_0, \ldots, G_{\mathfrak{n}^\delta} \leftarrow$ LowSpacePartition(G).
For each $i = 1, \ldots, \mathfrak{n}^\delta - 1$, perform LowSpaceColorReduce(G_i) in parallel.
Update color palettes of $G_{\mathfrak{n}^\delta}$, perform LowSpaceColorReduce($G_{\mathfrak{n}^\delta}$).
Update color palettes of G_0, color G_0 using MIS reduction.

Again, we employ a partitioning procedure to divide the input instance into *bins*, which are then solved recursively:

Algorithm 4 LowSpacePartition(G)

Let G_0 be the graph induced by the set V_0 of nodes v with $d(v) \leq \mathfrak{n}^{7\delta}$.
Let hash function $h_1 : [\mathfrak{n}] \to [\mathfrak{n}^\delta]$ map each node $v \notin V_0$ to a bin $h_1(v) \in [\mathfrak{n}^\delta]$.
Let hash function $h_2 : [\mathfrak{n}^2] \to [\mathfrak{n}^\delta - 1]$ map colors γ to a bin $h_2(\gamma) \in [\mathfrak{n}^\delta - 1]$.
Let $G_1, \ldots, G_{\mathfrak{n}^\delta}$ be the graphs induced by bins $1, \ldots, \mathfrak{n}^\delta$ respectively.
Restrict palettes of nodes in $G_1, \ldots, G_{\mathfrak{n}^\delta - 1}$ to colors assigned by h_2 to corresponding bins.
Return $G_0, \ldots, G_{\mathfrak{n}^\delta}$.

The structure of this algorithm is very similar to Algorithms 1 and 2. The main structural difference is that G_0 is no longer a graph of *bad* nodes (we will show that we can choose hash functions so that, in essence, no nodes are bad), but instead a graph of the nodes with low enough degree to apply the MIS reduction. By dealing with low-degree nodes throughout the algorithm in this way, we can account for differences in palette sizes, and thereby solve $(\deg+1)$-list coloring instead of just $(\Delta + 1)$-list coloring.

The only invariant we maintain during the course of LowSpaceColorReduce is that all nodes have sufficient colors to color themselves, i.e. $p(v) > d(v)$. This is clearly true for our initial call on input graph $\mathfrak{G}$, and later (in Lemma 28) it is preserved by LowSpaceColorPartition.

We would like to (analogously to Definition 9) define *good* nodes whose behavior does not differ too much from what we would expect from a random partition (and *bad* nodes for whom the opposite is true). However, one of the restrictions of the low-space regime is that we cannot store a node's palette, or all of its adjacent edges, on a single machine, and therefore machines could not determine locally whether a node is good or bad. So, we will instead divide nodes' neighbors and palettes into smaller sets which *do* fit onto single machines, and define a notion of a *machine* being good.

Specifically, for each node $v \notin G_0$, we will create a set M_N^v of machines which will be responsible for v's neighbors. We partition v's neighbor set among machines in M_N^v, with each machine receiving between $\mathfrak{n}^{7\delta}$ and $2\mathfrak{n}^{7\delta}$ neighbors. Similarly, for each node $v \notin G_0 \cup G_{\mathfrak{n}^\delta}$, we create a set M_C^v of machines responsible for the colors in its palette, and partition the palette among the machines, with each machine receiving between $\mathfrak{n}^{7\delta}$ and $2\mathfrak{n}^{7\delta}$ colors. In both cases, this is possible since for $v \notin G_0$, $p(v) > d(v) > \mathfrak{n}^{7\delta}$.

15

For consistency of notation, for a machine $x \in M_N^v$, we define $d(x)$ to be the number of neighbors it receives, and $d'(x)$ to be the number of such neighbors which are hashed to the same bin as v. Similarly, for a machine $x \in M_C^v$ holding colors for a node v, we define $p(x)$ to be the number of colors received, and $p'(x)$ to be the number of such colors hashed by h_2 to the same bin as v.

Definition 23 (Good and bad machines).

- *A machine $x \in M_N^v$ is* good *if $|d'(x) - d(x)\mathfrak{n}^{-\delta}| \le d(x)^{0.6}$.*

- *A machine $x \in M_C^v$ is* good *if $p'(x) > p(x)\mathfrak{n}^{-\delta} + p(x)^{0.7}$.*

- *Machines are* bad *if they are not* good.

Our next step is to show that, under a *random* choice of hash functions during partitioning, with high probability there are no bad machines.

4.3   Good and bad machines for *random* hash functions

We apply the bounded-independence concentration bound to show that node degrees and palette sizes within bins to not differ too much from their expectation:

Lemma 24. *For each machine $x \in M_N^v$, $|d'(x) - d(x)\mathfrak{n}^{-\delta}| \le d(x)^{0.6}$ with probability at least $1 - \mathfrak{n}^{-2}$.*

Proof. We apply Lemma 6 with $Z_1, \ldots, Z_{d(x)}$ as indicator random variables for the events that each neighbor held is hashed to v's bin. These variables are $(c-1)$-wise independent and each have expectation $\mathfrak{n}^{-\delta}$. Since we assume c to be a sufficiently high constant and we have $d(x) \ge \mathfrak{n}^{7\delta}$, by Lemma 6, we obtain $\mathbf{Pr}\left[|Z - \mathbf{E}[Z]| \ge d(x)^{0.6}\right] \le d(x)^{-0.2c} \le \mathfrak{n}^{-2}$. Since $\mathbf{E}[Z] = d(x)\mathfrak{n}^{-\delta}$, we have the required claim. $\square$

Lemma 25. *For each machine $x \in M_C^v$, $p'(v) > p(v)\mathfrak{n}^{-\delta} + p(v)^{0.7}$ with probability at least $1 - \mathfrak{n}^{-2}$.*

Proof. We apply Lemma 6, with $Z_1, \ldots, Z_{p(x)}$ as indicator random variables for the events that each color held is hashed to v's bin. These variables are $(c-1)$-wise independent and each have expectation $\frac{1}{\mathfrak{n}^\delta - 1}$. Since we assume c to be a sufficiently high constant and have $p(x) \ge \mathfrak{n}^{7\delta}$, by Lemma 6, we obtain $\mathbf{Pr}\left[|Z - \mathbf{E}[Z]| \ge p(x)^{0.6}\right] \le p(x)^{-0.2c} \le \mathfrak{n}^{-2}$.

Therefore, since $Z = p'(x)$ and $\mathbf{E}[Z] = \frac{p(x)}{\mathfrak{n}^\delta - 1}$, the probability that $p'(x) \le \frac{p(x)}{\mathfrak{n}^\delta - 1} - p(x)^{0.6}$ is at most $p\mathfrak{n}^{-2}$.

Furthermore, since $p(v) \ge \mathfrak{n}^{7\delta}$,

$$\frac{p(v)}{\mathfrak{n}^\delta - 1} - \frac{p(v)}{\mathfrak{n}^\delta} = \frac{p(v)}{\mathfrak{n}^\delta(\mathfrak{n}^\delta - 1)} \ge \frac{p(v)}{\mathfrak{n}^{2\delta}} > \frac{2p(v)}{p(v)^{0.3}} = 2p(v)^{0.7} \ .$$

Hence, with probability at least $1 - \mathfrak{n}^{-2}$ we obtain the following,

$$p'(v) > \frac{p(v)}{\mathfrak{n}^\delta - 1} - p(v)^{0.6} > \frac{p(v)}{\mathfrak{n}^\delta} + p(v)^{0.7} \ . \qquad \square$$

We will now define a cost function for pairs of hash functions that we would like to minimize, analogously to our definition in Section 3.2. This time our cost function is simpler, since we can weight all of our bad events equally, and with high probability none of them occur.

We define the *cost* function $\mathfrak{q}(h_1, h_2)$ of a pair of hash functions $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$, as follows:

$$\mathfrak{q}(h_1, h_2) = |\{\text{bad machines under } h_1, h_2\}| \ . \tag{2}$$

We can then bound the cost of random hash function pairs:

Lemma 26. *The expected cost of a random hash function pair from $\mathcal{H}_1 \times \mathcal{H}_2$ is less than 1.*

Proof. By Lemmas 24 and 25, any machine is bad with probability at most $\mathfrak{n}^{-2}$. The number of machines we require is $O(\mathfrak{n} + \frac{\mathfrak{m}}{\mathfrak{n}^{7\delta}}) = o(\mathfrak{n}^2)$. Therefore,

$$\mathbf{E}\left[\mathfrak{q}(h_1, h_2)\right] = \mathbf{E}\left[|\{\text{bad machines}\}|\right] \leq \mathfrak{n}^{-2} \cdot o(\mathfrak{n}^2) < 1 \ . \qquad \square$$

Applying the method of conditional expectations therefore gives the following:

Lemma 27. *In $O(1)$ MPC rounds with $O(\mathfrak{n}^{7\delta})$ local space per machine and $O(\mathfrak{m}+\mathfrak{n})$ total space, one can select hash functions $h_1 \in \mathcal{H}_1$, $h_2 \in \mathcal{H}_2$ such that in a single call to LOWSPACEPARTITION,*

- *for any node $v \notin G_0$, $d'(v) < 2d(v)\mathfrak{n}^{-\delta}$, and*

- *for any node $v \notin G_0 \cup G_{\mathfrak{n}^\delta}$, $d'(v) < p'(v)$.*

Proof. By the method of conditional expectations (using, as local functions q_x, the indicator variables for the event that machine x is bad), we can select a hash function pair h_1, h_2 with cost $q(h_1, h_2) < 1$; since cost must be a integer, this means that $q(h_1, h_2) = 0$, i.e. all machines are good.

By Definition 23, for any node $v \notin G_0$, we therefore have:

$$d'(v) = \sum_{x \in M_N^v} d'(x) \leq \sum_{x \in M_N^v} \left(d(x)\mathfrak{n}^{-\delta} + d(x)^{0.6}\right) \leq d(v)\mathfrak{n}^{-\delta} + |M_N^v|(2\mathfrak{n}^{7\delta})^{0.6} < d(v)\mathfrak{n}^{-\delta} + \mathfrak{n}^{4.5\delta}\frac{d(v)}{\mathfrak{n}^{7\delta}}$$
$$= d(v)\mathfrak{n}^{-\delta} + d(v)\mathfrak{n}^{-2.5\delta} \ .$$

Since $d(v)\mathfrak{n}^{-\delta} + d(v)\mathfrak{n}^{-2.5\delta} < 2d(v)\mathfrak{n}^{-\delta}$, this satisfies the first part of the lemma.

For any node $v \notin G_0 \cup G_{\mathfrak{n}^\delta}$, we also have:

$$p'(v) = \sum_{x \in M_C^v} p'(x) > \sum_{x \in M_C^v} \left(p(x)\mathfrak{n}^{-\delta} + p(x)^{0.7}\right) \geq p(v)\mathfrak{n}^{-\delta} + |M_C^v|(\mathfrak{n}^{7\delta})^{0.7}$$

$$\geq p(v)\mathfrak{n}^{-\delta} + \mathfrak{n}^{4.9\delta}\frac{p(v)}{2\mathfrak{n}^{7\delta}} > p(v)\mathfrak{n}^{-\delta} + p(v)\mathfrak{n}^{-2.5\delta} > d(v)\mathfrak{n}^{-\delta} + (v)\mathfrak{n}^{-2.5\delta} \geq d'(v) \ . \qquad \square$$

We can therefore show that whenever we call LOWSPACEPARTITION, we maintain our invariant that all nodes have sufficient colors, i.e. $p'(v) > d'(v)$.

Lemma 28. *Assume that at the start of our LOWSPACEPARTITION(G) call, all nodes v in G have $p(v) > d(v)$. Then, after the call, we have $p'(v) > d'(v)$.*

Proof. By Lemma 27, the claim holds for all $v \notin G_0 \cup G_{\mathfrak{n}^\delta}$. For a node $v \in G_0 \cup G_{\mathfrak{n}^\delta}$, v's palette is not restricted by partitioning and colors are only subsequently removed when used by v's neighbors, so it follows that $p'(v) > d'(v)$. $\qquad \square$

Now we are ready to prove Theorem 4, which states that Algorithm 3 applied to the input graph $\mathfrak{G}$ performs $(\deg +1)$-list coloring in $O(\log \Delta + \log \log \mathfrak{n})$ rounds of MPC, with $O(\mathfrak{n}^\varepsilon)$ space per machine and $O(\mathfrak{m} + \mathfrak{n}^{1+\varepsilon})$ total space.

Proof of Theorem 4. Each call of LowSpaceColorReduce makes two sequential sets of parallel recursive calls, on instances with maximum degree a factor of $\frac{1}{2}\mathfrak{n}^\delta$ lower than input maximum degree, by Lemma 27. (It also makes a call to the MIS reduction taking $O(\log \Delta + \log \log \mathfrak{n})$ rounds.) So, after a recursion depth of $\log_{\frac{1}{2}\mathfrak{n}^\delta} \Delta = O(1)$, the maximum degree is at most $O(\mathfrak{n}^{7\delta})$, and no further recursive calls are made. Since the depth of the recursion tree is $O(1)$, so is the number of LowSpaceColorReduce calls performed sequentially. So, the total round complexity is asymptotically equivalent to the complexity of a single LowSpaceColorReduce call, which is dominated by the $O(\log \Delta + \log \log \mathfrak{n})$ rounds of the MIS reduction.

Calls to LowSpacePartition use $O(\mathfrak{n}^{7\delta})$ local space, and all concurrent calls use $O(\mathfrak{n} + \mathfrak{m})$ global space in total, since we never store more than $O(1)$ copies of any node, edge or palette entry.

The global space complexity is dominated by the calls to MIS reduction, which as mentioned requires $O(\mathfrak{n}^\delta)$ and $O(\hat{\mathfrak{n}}^{1+\delta}\mathfrak{n}^{21\delta})$ local and global space respectively, where $\hat{\mathfrak{n}}$ is the number of nodes in the graph on which the call is made. We make many concurrent calls to the MIS reduction, but each node of $\mathfrak{G}$ is part of only one MIS reduction call. Therefore, the total space requirements for all concurrent calls are at most $O(\mathfrak{n}^\delta)$ local space and $O(\mathfrak{n}^{1+22\delta})$ global space. Setting $\delta \le \frac{\varepsilon}{22}$, this is $O(\mathfrak{n}^\varepsilon)$ and $O(\mathfrak{n}^{1+\varepsilon})$ respectively. $\qquad\square$

5 Conclusions

In this paper we have presented a constant-round deterministic $(\Delta+1)$-list coloring algorithm for the CONGESTED CLIQUE and linear-space MPC models, which not only greatly improves the deterministic complexity of the problem to optimality, but is also significantly simpler than the recent constant-round *randomized* algorithm [6]. Our algorithm relies on the derandomization of a recursive graph and palette sparsification procedure. We also extended this approach to low-space MPC; here the $O(\log \Delta + \log \log n)$ round complexity arises from the fact that we cannot collect instances onto single machines, and instead reduce to the problem to MIS and apply an existing algorithm once degree is sufficiently low.

With this work we settle the deterministic complexity of coloring in CONGESTED CLIQUE; the complexity in low-space MPC, however, remains open. A corresponding constant-round algorithm there seems unlikely, due to the $\Omega(\log \log \log \mathfrak{n})$ conditional lower bound [10], but there is no reason to believe one could not improve the dependency on Δ to sub-logarithmic. Another area for possible improvement is our $O(\mathfrak{m} + \mathfrak{n}^{1+\varepsilon})$ global space bound, which is slightly weaker than the optimal $O(\mathfrak{m} + \mathfrak{n})$.

References

[1] 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)*, pages 243–252, 2020.

[2] Leonid Barenboim and Victor Khazanov. Distributed symmetry-breaking algorithms for Congested Cliques. In *Proceedings of the 13th International Computer Science Symposium in Russia (CSR)*, pages 41–52, 2018.

[3] Soheil Behnezhad, Mahsa Derakhshan, and MohammadTaghi Hajiaghayi. Brief announcement: Semi-MapReduce meets Congested Clique. *CoRR abs/1802.10297*, 2018.

[4] Mihir Bellare and John Rompel. Randomness-efficient oblivious sampling. In *Proceedings of the 35th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 276–287, 1994.

[5] Keren Censor-Hillel, Merav Parter, and Gregory Schwartzman. Derandomizing local distributed algorithms under bandwidth restrictions. In *Proceedings of the 31st International Symposium on Distributed Computing (DISC)*, pages 11:1–11:16, 2017.

[6] 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)*, pages 471–480, 2019.

[7] Artur Czumaj, Peter Davies, and Merav Parter. Graph sparsification for derandomizing massively parallel computation with low space. In *Proceedings of the 32nd Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 175–185, 2020.

[8] Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified data processing on large clusters. In *Proceedings of the 6th Conference on Symposium on Opearting Systems Design & Implementation (OSDI)*, pages 10–10, 2004.

[9] Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified data processing on large clusters. *Commununication of the ACM*, 51(1):107–113, January 2008.

[10] 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)*, pages 1650–1663, 2019.

[11] 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)*, pages 374–383, 2011.

[12] Yijie Han. A fast derandomization scheme and its applications. *SIAM Journal on Computing*, 25(1):52–82, 1996.

[13] 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.

[14] Howard J. Karloff, Siddharth Suri, and Sergei Vassilvitskii. A model of computation for MapReduce. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 938–948, 2010.

[15] Christoph Lenzen. Optimal deterministic routing and sorting on the Congested Clique. In *Proceedings of the 32nd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 42–50, 2013.

[16] Zvi Lotker, Boaz Patt-Shamir, and Seth Pettie. Improved distributed approximate matching. *Journal of the ACM*, 62(5):38:1–38:17, November 2015.

[17] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.

[18] Merav Parter. $(\Delta + 1)$ coloring in the Congested Clique model. In *Proceedings of the 45th Annual International Colloquium on Automata, Languages and Programming (ICALP)*, pages 160:1–160:14, 2018.

[19] Merav Parter and Hsin-Hao Su. Randomized $(\Delta + 1)$-coloring in $O(\log^* \Delta)$ Congested Clique rounds. In *Proceedings of the 32nd International Symposium on Distributed Computing (DISC)*, pages 39:1–39:18, 2018.

[20] 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)*, pages 350–363, 2020.

[21] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012.

[22] Tom White. *Hadoop: The Definitive Guide.* O'Reilly Media, Inc., 2012.

[23] 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)*, 2010.