

# Distributed Recoloring

**Marthe Bonamy**

CNRS, LaBRI, Université de Bordeaux, France  
marthe.bonamy@u-bordeaux.fr

**Paul Ouvrard**

LaBRI, CNRS, Université de Bordeaux, France  
paul.ouvrard@u-bordeaux.fr

**Mikaël Rabie**

Aalto University, Finland  
mikael.rabie@aalto.fi

**Jukka Suomela**

Aalto University, Finland  
jukka.suomela@aalto.fi

**Jara Uitto**

ETH Zürich, Switzerland; and University of Freiburg, Germany  
jara.uitto@inf.ethz.ch

---

## Abstract

Given two colorings of a graph, we consider the following problem: can we recolor the graph from one coloring to the other through a series of elementary changes, such that the graph is properly colored after each step?

We introduce the notion of *distributed recoloring*: The input graph represents a network of computers that needs to be recolored. Initially, each node is aware of its own input color and target color. The nodes can exchange messages with each other, and eventually each node has to stop and output its own recoloring schedule, indicating when and how the node changes its color. The recoloring schedules have to be globally consistent so that the graph remains properly colored at each point, and we require that adjacent nodes do not change their colors simultaneously.

We are interested in the following questions: How many communication rounds are needed (in the deterministic LOCAL model of distributed computing) to find a recoloring schedule? What is the length of the recoloring schedule? And how does the picture change if we can use *extra colors* to make recoloring easier?

The main contributions of this work are related to distributed recoloring with one extra color in the following graph classes: trees, 3-regular graphs, and toroidal grids.

**2012 ACM Subject Classification** Theory of computation → Distributed computing models, Theory of computation → Graph algorithms analysis

**Keywords and phrases** Distributed Systems, Graph Algorithms, Local Computations

**Digital Object Identifier** 10.4230/LIPIcs.DISC.2018.12

**Related Version** The full version of this work is available at arXiv:1802.06742.

**Funding** This work was supported in part by ERC Grant No. 336495 (ACDC) and ANR project DISTANCIA (ANR-17-CE40-0015).

**Acknowledgements** The authors would like to thank Nicolas Bousquet for helpful discussions regarding the proof of Lemma 14.



© Marthe Bonamy, Paul Ouvrard, Mikaël Rabie, Jukka Suomela, and Jara Uitto; licensed under Creative Commons License CC-BY

32nd International Symposium on Distributed Computing (DISC 2018).

Editors: Ulrich Schmid and Josef Widder; Article No. 12; pp. 12:1–12:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

## 1 Introduction

In classical graph problems, we are given a graph and the task is to *find* a feasible solution. In *reconfiguration problems*, we are given two feasible solutions – an input configuration and a target configuration – and the task is to find a sequence of moves that turns the input configuration into the target configuration.

**Recoloring problems.** Perhaps the most natural example of a reconfiguration problem is *recoloring*: we are given a graph  $G$  and two proper  $k$ -colorings of  $G$ , let us call them  $s$  and  $t$ , and the task is to find a way to turn  $s$  into  $t$  by changing the color of one node at a time, such that each intermediate step is a proper coloring. More formally, the task is to find a sequence of proper  $k$ -colorings  $x_0, x_1, \dots, x_L$  such that  $x_0 = s$  and  $x_L = t$ , and  $x_{i-1}$  and  $x_i$  differ only at one node. Such problems have been studied extensively from the perspective of graph theory and classical centralized algorithms, but the problems are typically inherently *global* and solutions are long, i.e.,  $L$  is large in the worst case.

In this work we introduce recoloring problems in a *distributed* setting. We show that there are natural relaxations of the problem that are attractive from the perspective of distributed graph algorithms: they admit solutions that are short and that can be found *locally* (e.g., in sublinear number of rounds). Distributed recoloring problems are closely related to classical symmetry-breaking problems that have been extensively studied in the area of distributed graph algorithms, but as we will see, they also introduce new kinds of challenges.

**Distributed recoloring.** We will work in the usual LOCAL model of distributed computing: Each node  $v \in V$  of the input graph  $G = (V, E)$  is a computer, and each edge  $e \in E$  represents a communication link between two computers. Computation proceeds in synchronous rounds: each node sends a message to each of its neighbors, receives a message from each of its neighbors, and updates its local state. Eventually, all nodes have to announce their local outputs and stop; the running time of the algorithm is the number of communication rounds until all nodes stop. We assume that the algorithm is deterministic, and each node is labeled with a unique identifier.


In *distributed recoloring*, each node  $v \in V$  is given two colors, an *input color*  $s(v)$  and a *target color*  $t(v)$ . It is guaranteed that both  $s$  and  $t$  form a proper coloring of  $G$ , that is,  $s(u) \neq s(v)$  and  $t(u) \neq t(v)$  for all  $\{u, v\} \in E$ . Each node  $v \in V$  has to output a finite *recoloring schedule*  $x(v) = (x_0(v), x_1(v), \dots, x_\ell(v))$  for some  $\ell = \ell(v)$ . For convenience, we define  $x_i(v) = x_\ell(v)$  for  $i > \ell(v)$ . We say that the node *changes its color at time*  $i > 0$  if  $x_{i-1}(v) \neq x_i(v)$ ; let  $C_i$  be the set of nodes that change their color at time  $i$ . Define  $L = \max_v \ell(v)$ ; we call  $L$  the *length* of the solution. A solution is feasible if the following holds:

1.  $x_0 = s$  and  $x_L = t$ ,
2.  $x_i$  is a proper coloring of  $G$  for all  $i$ ,
3.  $C_i$  is an independent set of  $G$  for all  $i$ .

The key differences between distributed recoloring and classical recoloring are:

1. Input and output are given in a distributed manner: no node knows everything about  $G$ ,  $s$ , and  $t$ , and no node needs to know everything about  $x_i$  or the length of the solution  $L$ .
2. We do not require that only one node changes its color; it is sufficient that adjacent nodes do not change their colors simultaneously.

See Figure 1 for a simple example of distributed recoloring steps.



■ **Figure 1** Distributed recoloring: the input coloring  $s$  can be seen on the left and the target coloring  $t$  on the very right. The illustration shows one possible way to reach the target coloring in three steps by, in each step, changing the colors of an independent set of nodes.

Note that a solution to distributed recoloring is locally checkable in the following sense: to check that a solution is feasible, it is enough to check independently for each edge  $\{u, v\} \in E$  that the recoloring sequences  $x(u)$  and  $x(v)$  are compatible with each other, and for each node  $v \in V$  that  $x(v)$  agrees with  $s(v)$  and  $t(v)$ . However, distributed recoloring is not necessarily an LCL problem [21] in the formal sense, as the length of the output per node is not a priori bounded.

We emphasize that we keep the following aspects well-separated: what is the complexity of *finding* the schedule, and how *long* the schedules are. Hence it makes sense to ask, e.g., if it is possible to find a schedule of length  $O(1)$  in  $O(\log n)$  rounds (note that the physical reconfiguration of the color of the node may be much slower than communication and computation).

**Recoloring with extra colors.** Recoloring is computationally very hard, as solutions do not always exist, and deciding whether a solution exists is PSPACE-hard. It is in a sense analogous to problems such as finding an *optimal* node coloring of a given graph; such problems are not particularly interesting in the LOCAL model, as the complexity is trivially global. To make the problem much more interesting we slightly relax it.

We define a  $k + c$  recoloring problem (a.k.a.  $k$ -recoloring with  $c$  extra colors) as follows:

- We are given colorings with  $s(v), t(v) \in [k]$ .
- All intermediate solutions must satisfy  $x_i(v) \in [k + c]$ .

Here we use the notation  $[n] = \{1, 2, \dots, n\}$ .

The problem of  $k + c$  recoloring is meaningful also beyond the specific setting of distributed recoloring. For example, here is an example of a very simple observation:


► **Lemma 1.** *Recoloring with 1 extra color is always possible in any bipartite graph, with a distributed schedule of length  $L = 3$ .*

**Proof.** Let the bipartition be  $V = V_1 \cup V_2$ . First each node  $v \in V_1$  switches to  $k + 1$ , then each  $v \in V_2$  switches to color  $t(v)$ , and finally each  $v \in V_1$  switches to color  $t(v)$ . ◀

Incidentally, it is easy to extend this result to show that  $k$ -recoloring with  $c = \chi - 1$  extra colors is always possible with a schedule of length  $O(c)$  in a graph with chromatic number  $\chi$ , and in particular  $k$ -recoloring with  $c = k - 1$  extra colors is trivial. Figure 2 gives an illustration of recoloring a bipartite graph with one extra color.

As a corollary, we can solve distributed  $k + 1$  recoloring in trees in  $O(n)$  rounds, with a schedule of length  $O(1)$ : simply find a bipartition and apply the above lemma. However, is this optimal? Clearly finding a bipartition in a tree requires  $\Omega(n)$  rounds, but can we solve recoloring with 1 extra color strictly faster?

These are examples of problems that we study in this work. We initiate the study of distributed complexity of recoloring, with the ultimate objective of finding a complete



■ **Figure 2** In the input graph, a bipartition is given. Therefore, the target coloring can be reached by using one extra color in three steps.

characterization of graph families and parameters  $k$ ,  $c$ , and  $L$  such that distributed  $k + c$  recoloring with schedules of length  $L$  can be solved efficiently in a distributed setting.

As we will see, the problem turns out to be surprisingly rich already in very restricted settings such as grids or 3-regular trees. Many of the standard lower bound techniques fail; in particular, known results on the hardness of graph coloring do not help here, as we are already given two proper colorings of the input graph.

**Contributions.** Our main contribution is a comprehensive study of the complexity of distributed recoloring in various graph families; the results are summarized in Tables 1–5. The highlights of this work are the following results:

1. **An algorithm for 3 + 1 recoloring on trees.** On trees, 3-recoloring is inherently global: it is easy to see that the worst-case running time is  $\Theta(n)$  and the worst-case schedule length is  $\Theta(n)$ . With one extra color, we can trivially find a schedule of length  $O(1)$  in time  $O(n)$ . However, we show that we can do much better: it is possible to find a schedule of length  $O(1)$  in time  $O(\log n)$ .

Here the key component is a new algorithm that solves the following sub-problem in  $O(\log n)$  rounds: given a tree, find an independent set  $I$  such that the removal of  $I$  splits the tree in components of size 1 or 2. This subroutine may find applications in other contexts as well.

These results are presented in Section 5.

2. **An algorithm for 3 + 1 recoloring for graphs of degree at most 3.** In general graphs, 3 + 1 recoloring is not necessarily possible; we can construct a small 4-regular graph in which 3 + 1 recoloring is not solvable. However, we will show that if the maximum degree of the graph is at most 3 (i.e., we have a *subcubic* graph), 3 + 1 recoloring is always possible. Moreover, we can find a schedule of length  $O(\log n)$  in time  $\text{polylog}(n)$ .

This result is presented in Section 6.

3. **Complexity of 3 + 1 recoloring on toroidal grids.** We also give a complete characterization of 3 + 1 recoloring in one particularly interesting family of 4-regular graphs: 2-dimensional toroidal grids (a.k.a. torus grid graphs, Cartesian graph products of two cycles). While the case of 1-dimensional grids (cycles) is easy to characterize completely, the case of 2-dimensional grids turns out to be much more interesting.

Here our main contribution is the following graph-theoretic result: in an  $h \times w$  toroidal grid, 3 + 1 recoloring is possible for any input if and only if (i) both  $h$  and  $w$  are even, or (ii)  $h = 4$ , or (iii)  $w = 4$ . In all other cases we can find 3-colorings  $s$  and  $t$  such that  $t$  is not reachable from  $s$  even if we can use 1 extra color.

As a simple corollary, 3 + 1 recoloring is inherently global from the perspective of distributed computing, and it takes  $\Theta(n)$  rounds to solve even if we have the promise that e.g.  $h$  and  $w$  are even (and hence a schedule of length  $\Theta(1)$  trivially exists).

This result is presented in Section 7.

Additionally, several simple upper and lower bounds and corollaries are given in Sections 4 and 8 and in the full version of the article.

**Motivation.** As a simple application scenario, consider the task of reconfiguring a system of unmanned aerial vehicles. Here each node is an aircraft, the color corresponds to an altitude range, and an edge corresponds to a pair of aircraft whose paths might cross and hence need to be kept at different cruising altitudes to avoid collisions.

For each aircraft there are designated areas in which they can safely change their altitude. To reconfigure the entire system, we could take all aircraft to these areas simultaneously. However, this may be a costly maneuver.

Another possibility is to reserve a longer timespan during which a set  $X$  of aircraft may change their altitudes, whenever they happen to be at convenient locations. Now if we let two aircraft  $u, v \in X$  change their altitudes during the same timespan, we need to ensure that any intermediate configuration is safe, regardless of whether  $u$  or  $v$  happens to change its altitude first. Furthermore, we would like to complete reconfiguration in minimal time (short schedule), and we would like to waste precious airspace as little as possible and hence keep as few altitude levels as possible in reserve for reconfiguration (few extra colors).

This scenario – as well as many similar scenarios, such as the task of reconfiguring the frequency bands of radio transmitters in a manner that never causes interference, even if the clocks are not perfectly synchronized – give rise to the following variant of distributed recoloring that we call *weak recoloring*: if two adjacent nodes  $u$  and  $v$  change their color simultaneously at time  $i$ , then  $\{x_{i-1}(u), x_i(u)\} \cap \{x_{i-1}(v), x_i(v)\} = \emptyset$ , that is, we have a proper coloring regardless of whether  $u$  or  $v$  changes its color first.

Let us now contrast weak recoloring with *strong recoloring*, in which adjacent nodes never change colors simultaneously. Trivially, strong recoloring solves weak recoloring. But the converse is also true up to constant factors: if we have  $k$  input colors and a solution to weak recoloring of length  $L$ , then we can also find a solution to strong recoloring of length  $kL$ . To see this, we can implement one weak recoloring step in  $k$  strong recoloring substeps such that in substep  $j$  nodes of input color  $j$  change their colors.

As our focus is on the case of a small number of input colors, we can equally well study strong or weak recoloring here; all of our results hold for either of them. While weak recoloring is closer to applications, we present our results using strong recoloring, as it has a more convenient definition.

## 2 Related work

**Reconfiguration and recoloring.** Recoloring, and more generally combinatorial reconfiguration has received attention over the past few years. Combinatorial reconfiguration problems consist of finding step-by-step transformations between two feasible solutions such that all intermediate results are also feasible. They model dynamic situations where a given solution is in place and has to be modified, but no disruption can be afforded. We refer the reader to the nice survey [24] for a full overview, and focus here on node coloring as a reference problem.

As mentioned earlier, we introduce distributed recoloring here, but centralized recoloring has been studied extensively before. Two main models are considered:

1. *Node recoloring*: at each step, we can recolor a node into a new color that does not appear on its neighborhood
2. *Kempe recoloring*: at each step, we can switch the colors in a bichromatic component (we operate a Kempe change).

The usual questions are of the form: Given a graph  $G$  and an integer  $k$ , are all its  $k$ -colorings equivalent (up to node or Kempe recolorings)? What is the complexity of deciding that? What is the maximum number of operations needed to go from the other?

All of those questions can also be asked for two specific  $k$ -colorings  $s$  and  $t$  of  $G$ . Are they equivalent (up to node or Kempe recolorings)? What is the complexity of deciding that? What is the maximum number of operations needed to go from  $s$  to  $t$  in  $G$ ?

While the complexity of questions related to Kempe recoloring remains elusive, the problems related to node recoloring are typically PSPACE-hard [6]. The related question of deciding equivalence when a bound on the length of an eligible recoloring sequence is given as part of the input has also been considered [7]. We know that the maximum number of operations needed to go from one 3-coloring to another in a tree is  $\Theta(n)$  [11]. While  $(\Delta + 1)$ -recoloring a graph with no node of degree more than  $\Delta$  is not always possible, having  $\Delta + 2$  colors always suffices [16], and there are also meaningful results to obtain for the problem of  $(\Delta + 1)$ -recoloring [14]. Two other settings have received special attention: characterizing fully when 3-recoloring is possible [11, 10], and guaranteeing short reconfiguration sequences in the case of sparse graphs for various notions of sparse [4, 8].

Kempe changes were introduced in 1879 by Kempe in his attempted proof of the Four Color Theorem [17]. Though this proof was fallacious, the Kempe change technique has proved useful in, for example, the proof of the Five Color Theorem and a short proof of Brooks' Theorem. Most works on the topic initially focused on planar graphs, but significant progress was recently obtained in more general settings. We know that all  $k$ -colorings of a graph with no node of degree more than  $k$  are equivalent (w.r.t. Kempe changes), except in the case of one very specific graph: the 3-prism [5, 15, 18].

Note that some other variants have also been studied, perhaps most notably the question of how many nodes to recolor at once so that the graph can be recolored [19].

While we will not discuss Kempe recoloring in our work, we point out that recoloring with extra colors is closely connected to Kempe recoloring: Kempe recolorability implies recolorability with one extra color (while the converse is not true). Hence the negative results related to one extra color also hold for Kempe recoloring.

**Distributed graph coloring.** Panconesi and Srinivasan [23] have used Kempe operations to design efficient distributed algorithms for graph coloring with  $\Delta$  colors. Other than that we are not aware of prior work on distributed recoloring. On the other hand, the literature on the standard distributed coloring is vast. The best overview on the topic is the book by Barenboim and Elkin [3]; the most important recent developments include the following results. There is a randomized  $O(\log^* n + 2\sqrt{\log \log n})$ -time algorithm for  $(\Delta + 1)$ -coloring by Chang et al. [13]. In the case of trees, the number of colors can be reduced to  $\Delta$  with the cost of increasing the runtime to  $O(\log_{\Delta} \log n)$  [12]. On the deterministic side, the best known  $(\Delta + 1)$ -coloring algorithm requires  $O(\Delta^{3/4} \log \Delta + \log^* n)$  communication rounds [2]. In the case of trees, the *rake-and-compress* -method by Miller and Reif gives a 3-coloring in time  $O(\log n)$  [20].

However, there seems to be surprisingly little technology that one can directly transfer between the coloring domain and recoloring domain. Toroidal grids are a good example: by prior work [9], 3-coloring is an inherently global problem, and by the present work,  $3 + 1$  recoloring is an inherently global problem, but the arguments that are used in these proofs are very different (despite the fact that both of them are related to the idea that a “parity” is preserved).

### 3 Preliminaries

In this article, each graph  $G = (V, E)$  is a simple undirected graph where  $V$  represents its node set and  $E$  its edge set. For a subset of nodes  $S \subseteq V$ , we denote by  $G[S]$  the subgraph induced by  $S$ . For a node  $u \in V$ , we denote by  $N(u)$  the *open neighborhood* of  $u$  that is the set of all the neighbors of  $u$  and by  $N[u]$  its *closed neighborhood* i.e. the set  $N(u) \cup \{u\}$ . For a subset  $S \subseteq V$ , its closed neighborhood corresponds to the set  $\bigcup_{u \in S} N[u]$ .

The *degree* of a node is the number of neighbors. A *k-regular graph* is a graph in which all nodes have degree  $k$ , a *cubic graph* is the same thing as a 3-regular graph, and a *subcubic graph* is a graph in which all nodes have degree at most 3. A *tree* is a connected acyclic graph, and a *k-regular tree* is a tree in which each node has degree 1 or  $k$ .

A *maximal independent set* (MIS)  $S \subseteq V$  is an independent set (i.e. a set of pairwise non-adjacent nodes) such that for each non-MIS node  $u \notin S$ ,  $N(u) \cap S \neq \emptyset$ .

Given a graph  $G = (V, E)$ , a *list-assignment* is a function which assigns to each node  $v \in V$  a list of colors  $L(v)$ . An *L-coloring* of  $G$  is a function  $c$  that assigns to each node  $v \in V$  a color  $c(v) \in L(v)$  such that for any two adjacent nodes  $u, v \in V$ , we have  $c(u) \neq c(v)$ . A graph  $G$  is *k-list-colorable* if it admits an *L-coloring* for every list-assignment where the list of each node is of size at least  $k$ . Therefore, list-coloring generalizes node-coloring if we consider the special case where each node receives the same input list. The notion of *L-recoloring* is the natural generalization of *k-recoloring*: the same elementary steps are considered, and every intermediate coloring must be an *L-coloring*.

In order to output a recoloring schedule, it is convenient to consider the question of recoloring a graph  $G$  from a coloring  $s$  to a coloring  $t$ , rather than the more symmetric question of whether the two colorings are equivalent in the given setting. We take this opportunity to note that we can reverse time and hence recoloring schedule from  $s$  to  $t$  also yields a recoloring schedule from  $t$  to  $s$ . In the rest of the paper, we therefore address the two questions as one.

### 4 Simple upper bounds

We will start by providing some simple upper bounds of recoloring problems:

► **Lemma 2.** *In any graph,  $k + c$  recoloring for  $c = k - 1$  is possible in 0 communication rounds, with a schedule of length  $O(k)$ .*

**Proof.** Generalize the idea of Lemma 1; note that the schedule of node  $v$  depends only on  $s(v)$  and  $t(v)$ , and not on the colors of any other node around it. ◀

► **Lemma 3.** *Let  $G$  be a graph of maximum degree at most  $\Delta$ , and let  $k \geq \Delta + 2$ . Then  $k$ -recoloring with  $c$  extra colors is at least as easy as  $(k - 1)$ -recoloring with  $c + 1$  extra colors.*

**Proof.** Given a  $k$ -coloring  $x$ , we can construct a  $(k - 1)$ -coloring  $x'$  as follows: all nodes of color  $k$  pick a new color from  $\{1, 2, \dots, k - 1\}$  that is not used by any of their neighbors. Note that  $x \rightarrow x'$  is a valid step in distributed recoloring (nodes of color  $k$  form an independent set), and by reversing the time, also  $x' \rightarrow x$  is a valid step.

Hence to recolor  $s \rightarrow t$  with  $c$  extra colors, it is sufficient to recolor  $s' \rightarrow t'$  with  $c + 1$  extra colors (color  $k$  no longer appears in the input and target colorings and can be used as an auxiliary color during recoloring). Then we can put everything together to form a recoloring schedule  $s \rightarrow s' \rightarrow t' \rightarrow t$ , with only constant overhead in the running time and schedule length. ◀

Please see the full version for more examples of simple upper and lower bounds.



## 5 Recoloring algorithm for trees

In this section, we provide two efficient algorithms for recoloring and list-recoloring trees. Note that Theorem 5 is tight; see the full version for more details.

► **Theorem 4.** *For any  $k \in \mathbb{N}$ , for every tree  $T$  on  $n$  nodes, for any two  $k$ -colorings  $\alpha, \beta$  of  $T$ , we can compute in  $O(\log n)$  rounds how to recolor  $T$  from  $\alpha$  to  $\beta$  with 1 extra color and a schedule of length  $O(1)$ .*

► **Theorem 5.** *For every tree  $T$  on  $n$  nodes and any list assignment  $L$  of at least 4 colors to every node of  $T$ , for any two  $L$ -colorings  $\alpha, \beta$  of  $T$ , we can compute in  $O(\log n)$  rounds how to  $L$ -recolor  $T$  from  $\alpha$  to  $\beta$  with schedule of length  $O(\log n)$ .*

We first discuss how to compute efficiently an independent set with some desirable properties. For this, we use a simple modification of the *rake and compress* method by Reif and Miller [20]. More precisely, we iterate rake and compress operations, and label nodes based on the step at which they are reached. We then use the labels to compute an independent set satisfying given properties. We finally explain how to make use of the special independent set to obtain an efficient recoloring algorithm, in each case.

► **Definition 6.** A *light  $h$ -labeling* is a labeling  $V \rightarrow [h]$  such that for any  $i \in [h]$ :

1. Any node labeled  $i$  has at most two neighbors with label  $\geq i$ , at most one of which with label  $\geq i + 1$ .
2. No two adjacent nodes labeled  $i$  both have a neighbor with label  $\geq i + 1$ .

► **Lemma 7.** *There is an  $O(\log n)$ -round algorithm that finds a light  $(2 \log n)$ -labeling of a tree.*

**Proof.** As discussed above, we merely use a small variant of the *rake and compress* method. At step  $i$ , we remove all nodes of degree 1 and all nodes of degree 2 that belong to a chain of at least three nodes of degree 2, and assign them label  $i$ .

One can check that this yields a light labeling. It remains to discuss how many different labels are used, i.e. how many steps it takes to delete the whole tree. Let us argue that no node remains after  $2 \log n$  rounds. Let  $T$  be a tree, let  $V_1$  (resp.  $V_2, V_3$ ) be the number of nodes of degree 1 (resp. 2,  $\geq 3$ ) in the tree, and let  $T'$  be the tree obtained from  $T$  by replacing any maximal path of nodes of degree 2 with an edge. Note that  $|V(T')| = |V_1| + |V_3|$ . Let  $W$  be the set of nodes in  $T$  that have degree 2 with both neighbors of degree 2. Note that  $|V_2 \setminus W| \leq 2|E(T')| = 2(|V_1| + |V_3| - 1)$ . Note also that  $|V_1| \geq |V_3|$ , simply by the fact that there are fewer edges than nodes in a tree. It follows that  $|W| \geq |V_2| - 2(|V_1| + |V_3| - 1) = |V(T)| - |V_1| - |V_3| - 2(|V_1| + |V_3| - 1) \geq |V(T)| - 6|V_1|$ . Consequently, we obtain  $|W| + |V_1| \geq \frac{|V|}{6}$ . In other words, at every step, we remove in particular  $W \cup V_1$ , hence at least a sixth of the nodes. It follows that after  $k$  steps, the number of remaining nodes is at most  $n \cdot (\frac{5}{6})^k$ . Note that this is less than 1 once  $k \geq 2 \log n$ . ◀

We now discuss how to make use of light  $h$ -labelings.

► **Lemma 8.** *For any graph  $T$ , any 3-coloring  $\alpha$  of  $T$ , and any integer  $h$ , let  $L$  be a light  $h$ -labeling of  $T$ . There is an  $O(h)$ -round algorithm that finds a maximal independent set  $S$  such that  $T \setminus S$  only has connected components on 1 or 2 nodes.*



**Proof.** In brief, we proceed as follows: at step  $i = h, h - 1, \dots, 1$ , we first add all nodes of label  $i$  which have a neighbor of label  $\geq i + 1$  that is not in  $S$  (they form an independent set by definition of a light label), then use the 3-coloring to obtain a fast greedy algorithm to make  $S$  maximal on the nodes of label  $\geq i$ . The detailed algorithm can be found in the full version.

The fact that the output  $S$  is an independent set follows directly from the construction, as does the fact that the running time is  $O(h)$  rounds. We note that no connected component of  $T \setminus S$  contains nodes of different labels, due to the first operation at step  $i$ .

It remains to argue that for any  $i$ , the nodes of label  $i$  that do not belong to  $S$  only form connected components of size 1 or 2. Assume for a contradiction that there is a node  $u$  of label  $i$  which has two neighbors  $v$  and  $w$ , also of label  $i$ , such that none of  $\{u, v, w\}$  belongs to  $S$ . By definition of a light label, the node  $u$  has no other neighbor of label  $\geq i$ , a contradiction to the fact that we build  $S$  to be an MIS among the nodes of label  $\geq i$ . ◀

Combining Lemmas 7 and 8, and observing that a 3-coloring of a tree can be obtained in  $O(\log n)$  rounds, we immediately obtain the following.

► **Lemma 9.** *There is an  $O(\log n)$ -round algorithm that finds an MIS in a tree, such that every component induced by non-MIS nodes is of size one or two.*

We are now ready to prove Theorem 4.

**Proof of Theorem 4.** First, we use Lemma 9 to obtain in  $O(\log n)$  rounds an MIS  $S$  such that  $T \setminus S$  only has connected components of size 1 or 2. We recolor each node in  $S$  with the extra color. Remove  $S$ , and recolor each component from  $\alpha$  to  $\beta$  without using any extra colors; this can be done in  $O(1)$  recoloring rounds. Each node in  $S$  can then go directly to its color in  $\beta$ . ◀

Moving on to the list setting, we have to use a more convoluted approach since there is no global extra color that we can use. Before discussing 4-list-recoloring, we discuss 3-list-recoloring. For the sake of intuition, we start by presenting an algorithm for 3-recoloring trees, and explain afterwards how to adapt it for the list setting.

► **Lemma 10.** *For every tree  $T$  with radius at most  $p$  and for any two 3-colorings  $\alpha, \beta$  of  $T$ , we can compute in  $O(p)$  rounds how to 3-recolor  $T$  from  $\alpha$  to  $\beta$  with a schedule of length  $O(p)$ .*

**Proof.** Let  $c: V \rightarrow [3]$  be a 3-coloring of  $T$ . We introduce an identification operation: Given a leaf  $u$  and a node  $v$  such that  $u$  and  $v$  have a common neighbor  $w$ , we recolor  $u$  with  $c(v)$ , and from then on we pretend that  $u$  and  $v$  are a single node. In other words, we delete  $u$  from the tree we are considering, and reflect any recoloring of  $v$  to the node  $u$ . Note that these operations can stack up: the recoloring of a single node might be reflected on an arbitrarily large independent set in the initial tree.

We now briefly describe an algorithm to recolor a 3-coloring into a 2-coloring  $c'$  in  $O(p)$  rounds, with schedule  $O(p)$ . First, root  $T$  on a node  $r$  which is at distance at most  $p$  of any node of  $T$ . Any node of  $T$  which is not adjacent to the root has a *grandparent*, which is defined as its parent's parent.

Then, at each step, we consider the set  $A$  of leaves of  $T$  which have a grandparent, if any. We identify each leaf in  $A$  with its grandparent (note that the notion of grandparent guarantees that this operation is well-defined, and that the operation results in  $A$  being deleted).

## 12:10 Distributed Recoloring

This process stops when  $T$  consists only of the root  $r$  and its children. We select one of the children arbitrarily and identify the others with it. This results in  $T$  being a single edge. Note that the color partition of  $c'$  is compatible with the identification operations, as we only ever identify nodes at even distance of each other.

We then recolor  $T$  into  $c'$ : this is straightforward in the realm of 3-recoloring.

We can now choose a 2-coloring of  $T$  (this can be done in  $O(p)$  rounds), and apply the above algorithm to 3-recolor both  $\alpha$  and  $\beta$  to that 2-coloring. This results in a 3-recoloring between  $\alpha$  and  $\beta$  with schedule  $O(p)$ . ◀

The same idea can be adapted to list coloring; we give a proof of the following result in the full version of the article:

► **Lemma 11.** *For every tree  $T$  with radius at most  $p$ , for any list assignment  $L$  of at least 3 colors to each node, for any two  $L$ -colorings  $\alpha, \beta$  of  $T$ , we can compute in  $O(p)$  rounds how to  $L$ -recolor  $T$  from  $\alpha$  to  $\beta$  with schedule  $O(p)$ .*

To prove Theorem 5, we first split the tree in small components. We slightly adapt the proof of Lemma 8; see the full version for the details:

► **Lemma 12.** *For any tree  $T$ , any 3-coloring  $\alpha$  of  $T$ , and any integer  $h$ , let  $L$  be a light  $h$ -label of  $T$ . There is a  $O(h)$ -round algorithm that finds a maximal independent set  $S$  such that no node has two neighbors in  $S$  and  $T \setminus S$  only has connected components of radius  $O(h)$ .*

Now we are ready to prove Theorem 5.

**Proof of Theorem 5.** Compute (in  $O(\log n)$  rounds) an independent set  $S$  such any two elements of  $S$  are at distance at least 2 of each other and every connected component of  $T \setminus S$  has radius  $O(\log n)$ . By Lemmas 7 and 12 and the fact that a 3-coloring of a tree can be computed in  $O(\log n)$  rounds, we compute (in  $O(\log n)$  rounds) an  $L$ -coloring  $\gamma$  of  $T \setminus S$  such that every node adjacent to an element  $u \in S$  has a color different from  $\alpha(u)$  and  $\beta(u)$ . Note that this coloring exists since any tree is 2-list-colorable. Use Lemma 11 to recolor each connected component of  $T \setminus S$  from  $\alpha$  to  $\gamma$ . Recolor every element of  $S$  with its color in  $\beta$ . Use Lemma 11 to recolor each connected component  $T \setminus S$  from  $\gamma$  to  $\beta$ . Note that this yields an  $L$ -recoloring of  $T$  from  $\alpha$  to  $\beta$  with schedule  $O(\log n)$ . ◀

Note that a direct corollary of Theorem 5 is that for any  $k$ -coloring  $\alpha, \beta$  of a trees with  $k \geq 4$ , a schedule of length  $\Theta(\log n)$  can be found in  $\Theta(\log n)$  rounds.

## 6 Recoloring algorithm for subcubic graphs

In this section we study recoloring in subcubic graphs (graphs of maximum degree at most 3); our main result is summarized in the following theorem:

► **Theorem 13.** *For every subcubic graph  $G$  on  $n$  nodes, for any two 3-colorings  $\alpha, \beta$  of  $G$ , we can compute in  $O(\log^2 n)$  rounds how to recolor  $G$  from  $\alpha$  to  $\beta$  with 1 extra color and a schedule of length  $O(\log n)$ .*

A *theta* is formed of three node-disjoint paths between two nodes. Note that in particular if a graph contains two cycles sharing at least one edge, then it contains a theta. We note  $B^k(u)$  the set of nodes at distance at most  $k$  to  $u$ .

We show here, roughly, that there is around every node a nice structure that we can use to design a valid greedy algorithm for the whole graph. This proof is loosely inspired by one in [1]. The proofs of Lemmas 14 and 15 are given in the full version.

**Algorithm 1** DECOMPOSING INTO A SMALL FOREST AND AN INDEPENDENT SET.**Require:** A subcubic graph  $G$ .**Ensure:** A decomposition  $(F, S)$  of  $V(G)$  such that  $G[S]$  is an independent set and every connected component of  $G[F]$  has radius at most  $\log n$ .

- 1: **for**  $u$  in  $V(G)$  (in parallel) **do**
- 2:   Acquire knowledge on  $B^{2\log n}(u)$
- 3:   Select in the node set of  $B^{2\log n}(u)$  a configuration  $C(u)$  that is a minimal theta or a node of degree 1 or 2
- 4: **end for**
- 5: Compute a  $(4\log n, 8\log n)$ -ruling set  $X$  in  $G$
- 6: Define  $\mathcal{A} = \cup_{u \in X} \{C(u)\}$
- 7: Compute the distance of every node in  $G$  to an element of  $\mathcal{A}$
- 8: Let  $F = S = \emptyset$
- 9: **for**  $i = 8\log n$  downto 1 **do**
- 10:   Extend the partition  $(F, S)$  to the nodes at distance  $i$  from  $\mathcal{A}$ , more precisely:
- 11:   Each connected component is a path or cycle where no internal node has an already assigned neighbor, let  $U_i$  be the set of the internal nodes
- 12:   Assuming a pre-computed MIS on each layer for the sets  $U_i$ , assign that MIS to  $S$
- 13:   Extend greedily on the remaining nodes (which form bounded-size components), assigning nodes to  $S$  when possible, to  $F$  when not
- 14: **end for**
- 15: Extend the partition  $(F, S)$  to the nodes belonging to an element of  $\mathcal{A}$  using Lemma 15

► **Lemma 14.** *For every subcubic graph  $G$  on  $n$  nodes, for every node  $u \in V(G)$ , there is a node  $v$  with degree at most 2 or a theta that is contained in  $B^{2\log n}(u)$ .*

► **Lemma 15.** *Let  $G$  be a subcubic graph, let  $p$  be an integer, and let  $\mathcal{A}$  be a collection of thetas and nodes of degree  $\leq 2$  in  $G$  each at distance at least 2 of each other. Let  $r \geq 1$  be such that no element of  $\mathcal{A}$  has diameter more than  $\frac{r}{2}$ . If the nodes of  $G \setminus (\cup_{A \in \mathcal{A}} A)$  can be partitioned into  $S$  and  $F$  such that  $G[S]$  is an independent set and  $G[F]$  is a forest of radius at most  $p$ , then there is a partition  $(S', F')$  of  $\cup_{A \in \mathcal{A}} A$  such that  $G[S \cup S']$  is an independent set and  $G[F \cup F']$  is a forest of radius at most  $p + r$ .*

► **Lemma 16.** *Let  $G$  be a subcubic graph on  $n$  nodes. We can compute in  $O(\log^2 n)$  rounds a partition  $(S, F)$  of the nodes of  $G$  that  $G[S]$  is an independent set and  $G[F]$  is a forest of radius  $O(\log n)$ .*

**Proof.** To that purpose, we combine the previous lemmas in Algorithm 1. The algorithm computes a decomposition as desired and runs in  $O(\log n) + RS(n)$  rounds, where  $RS(n)$  is the number of rounds necessary to compute a  $(4\log n, 8\log n)$ -ruling set in a subcubic graph. We derive from [22] that  $RS(n) = O(\log^2(n))$ , hence the conclusion. ◀

We are now ready to prove Theorem 13, which we do in a similar fashion as Theorem 4.

**Proof.** Use Lemma 16, and obtain a decomposition  $(S, F)$  as stated. Recolor all of  $S$  to the extra color, then use Lemma 11 on each connected component of  $G[F]$  so that all nodes of  $F$  reach their target color (remember that each connected component of  $G[F]$  has radius  $O(\log n)$ ). Finally recolor each node of  $S$  with its target color. ◀

## 7 Recoloring in toroidal grids

In this section we study toroidal grids (torus grid graphs). Throughout this section, an  $h \times w$  toroidal grid is the Cartesian graph product of cycles of lengths  $h$  and  $w$ ; we assume  $h \geq 3$  and  $w \geq 3$ . A toroidal grid can be constructed from an  $h \times w$  grid by wrapping both boundaries around into a torus. In the full version, we show that e.g.  $2 + 0$ ,  $3 + 0$ , and  $4 + 0$  recoloring is not always possible, and by Lemma 2 e.g.  $2 + 1$ ,  $3 + 2$ , and  $4 + 3$  recoloring is trivial. The first nontrivial case is  $3 + 1$  recoloring; in this section we give a complete characterization of  $3 + 1$  recolorability in toroidal grids:

► **Theorem 17.** *Let  $G$  be the  $h \times w$  toroidal grid graph. Then  $3 + 1$  recoloring is possible for any source and target coloring in the following cases: (i) both  $h$  and  $w$  are even, or (ii)  $h = 4$ , or (iii)  $w = 4$ . For all other cases it is possible to construct 3-colorings  $s$  and  $t$  such that  $t$  is not reachable from  $s$  by valid recoloring operations using 1 extra color.*

This also shows that  $3 + 1$  recoloring is an inherently global problem in toroidal grids, even if we have a promise that recoloring is possible. For example, if there was a sublinear-time distributed recoloring algorithm  $A$  for  $6 \times w$  grids for an even  $w$ , we could apply the same algorithm in a  $6 \times w$  grid with an odd  $w$  (the algorithm cannot tell the difference between these two cases in time  $o(w)$ ), and hence we could solve recoloring in  $6 \times w$  grids for all  $w$ , which contradicts Theorem 17. By a similar argument, distributed recoloring in non-toroidal grids is also an inherently global problem.

**Existence.** To prove Theorem 17, let us start with the positive results. If  $h$  and  $w$  are even, the graph is bipartite and recoloring is always possible by Lemma 1. The remaining cases are covered by the following lemma.

► **Lemma 18.** *Let  $G$  be a  $4 \times w$  toroidal grid for any  $w \geq 3$ , and let  $s$  and  $t$  be any 3-colorings. Then there exists a recoloring from  $s$  to  $t$  with one extra color.*

**Proof.** We first take an MIS  $S$  over pairs of consecutive columns, i.e. a set of indices of the form  $(i, i + 1)$  such that every column  $j \notin S$  is such that at least one of  $j - 1$  and  $j + 2$  belongs to  $S$ , every column  $i \in S$  is such that precisely one of  $i - 1$  and  $i + 1$  is in  $S$ . Note that indices are taken modulo  $w$ . For every pair in  $S$ , we select a maximal independent set of the corresponding columns. The resulting union yields an independent set  $R$ . We then greedily make  $R$  maximal columnwise away from  $S$ . We recolor  $R$  with the extra color. It remains to argue that  $G \setminus R$  can reach its targeted coloring. We note that since leaves are not problematic, removing  $R$  essentially boils down to removing the columns with index in  $S$ . Note that the remaining connected components are cycles of length 4. Cycles of length 4 can be always 3-recolored.

Note that the above proof yields in fact an  $O(\log n)$  rounds algorithm that outputs an  $O(1)$  schedule. We can improve it into an  $O(1)$ -round algorithm, simply by pointing out that there is only a finite number of possible colorings for a column, and two adjacent columns cannot have the same coloring. This allows us to compute  $S$  in constant time. ◀

**Non-existence.** Let us now prove the negative result. Our high-level plan is as follows. Let  $G$  be an  $h \times w$  toroidal grid. We will look at all *tiles* of size  $2 \times 2$ . If  $G$  is properly colored with  $k$  colors, so is each tile. The following two tiles are of special importance to us; we call these tiles of *type A*:

$$\begin{bmatrix} 2 & 3 \\ 3 & 1 \end{bmatrix}, \quad \begin{bmatrix} 1 & 3 \\ 3 & 2 \end{bmatrix}.$$

We are interested in the *number* of type-A tiles. For example, consider the following colorings of the  $3 \times 3$  toroidal grid:

$$s = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \\ 3 & 1 & 2 \end{bmatrix}, \quad t = \begin{bmatrix} 3 & 1 & 2 \\ 2 & 3 & 1 \\ 1 & 2 & 3 \end{bmatrix}.$$

Here  $s$  contains 3 tiles of type A (recall that we wrap around at the boundaries), while  $t$  does not have any tiles of type A. In particular,  $s$  has an odd number of type-A tiles and  $t$  has an even number of type-A tiles. In brief, we say that the *A-parity* of  $s$  is odd and the A-parity of  $t$  is even. It turns out that this is sufficient to show that recoloring from  $s$  to  $t$  with one extra color is not possible (see the full version of the article for the proof of this lemma):

► **Lemma 19.** *Let  $G$  be a toroidal grid, and let  $s$  and  $t$  be two 3-colorings. If  $s$  and  $t$  have different A-parities, then it is not possible to recolor  $G$  from  $s$  to  $t$  with 1 extra color.*

Hence the A-parity of a coloring partitions the space of colorings in two components that are not connected by  $3 + 1$  recoloring operations. To complete the proof of Theorem 17, it now suffices to construct a pair of 3-colorings with different A-parities for each relevant combination of  $h$  and  $w$ . We give the details in the full version.

## 8 Simple corollaries

► **Lemma 20.** *Assume that we are given a graph  $G$  and input and target colorings with  $k \geq 3$  colors. Assume that in  $O(f(n))$  rounds we can find an independent set  $I$  of  $G$  such that  $V \setminus I$  induces a forest of trees of depth at most  $O(d(n))$ . Then in  $O(f(n) + d(n))$  rounds we can solve  $k + 1$  recoloring, with a schedule of length  $O(d(n))$ .*

**Proof.** Each node in  $I$  switches to color  $k + 1$ . We then use the algorithm described in the proof of Lemma 10 to find a recoloring with schedule of length  $O(d(n))$  for each connected component after the removal of  $I$ . After that, each node of  $I$  can switch to its final color. ◀

► **Lemma 21.** *In cycles and paths,  $3 + 1$  recoloring is possible in  $O(1)$  rounds, with a schedule of length  $O(1)$ .*

**Proof.** Use the input coloring to find a maximal independent set  $I$ . Nodes of  $V \setminus I$  induce paths of length  $O(1)$ , apply Lemma 20. ◀

► **Lemma 22.** *In subcubic graphs,  $4 + 1$  recoloring is possible in  $O(1)$  rounds, with a schedule of length  $O(1)$ .*

**Proof.** Use the input coloring to find a maximal independent set  $I$  in constant time. Nodes of  $I$  switch to color 5. Delete  $I$ ; we are left with a graph  $G'$  that consists of paths and isolated nodes. Apply Lemmas 3 and 21 to solve  $4 + 0$  recoloring in each connected component of  $G'$ . Finally nodes of  $I$  can switch to their target colors. ◀

## 12:14 Distributed Recoloring

■ **Table 1** Results: distributed recoloring in cycles (C) and paths (P).

graph family	input colors	extra colors	schedule length	communication rounds	reference
C/P	2	0	$\infty$		see full version
C/P	2	1	$O(1)$	0	Lemma 2
C	3	0	$\infty$		see full version
P	3	0	$\Theta(n)$	$\Theta(n)$	see full version
C/P	3	1	$O(1)$	$O(1)$	Lemma 21
C/P	3	2	$O(1)$	0	Lemma 2
C/P	4	0	$O(1)$	$O(1)$	Lemmas 21 and 3
C/P	4	3	$O(1)$	0	Lemma 2

■ **Table 2** Results: distributed recoloring in 3-regular trees.

input colors	extra colors	schedule length	communication rounds	reference
2	0	$\infty$		see full version
2	1	$O(1)$	0	Lemma 2
3	0	$\Theta(n)$	$\Theta(n)$	see full version
3	1	$O(1)$	$O(\log n)$	Theorem 4
3	2	$O(1)$	0	Lemma 2
4	0	$\Theta(\log n)$	$\Theta(\log n)$	Theorem 5 and see full version
4	1	$O(1)$	$O(1)$	Lemma 22
4	3	$O(1)$	0	Lemma 2
5	0	$O(1)$	$O(1)$	Lemmas 22 and 3

■ **Table 3** Results: distributed recoloring in trees.

input colors	extra colors	schedule length	communication rounds	reference
2	0	$\infty$		see full version
2	1	$O(1)$	0	Lemma 2
3	0	$\Theta(n)$	$\Theta(n)$	see full version
3	1	$O(1)$	$O(\log n)$	Theorem 4
3	2	$O(1)$	0	Lemma 2
4	0	$\Theta(\log n)$	$\Theta(\log n)$	Theorem 5 and see full version
4	1	$O(1)$	$O(\log n)$	Theorem 4
4	3	$O(1)$	0	Lemma 2

■ **Table 4** Results: distributed recoloring in toroidal grids. The distributed complexity of  $4 + 1$  recoloring is left as an open question. However, by prior work it is known that  $4 + 1$  recoloring is always possible: grids are 4-regular graphs, therefore they are 4-recolorable with Kempe operations, and hence also with 1 extra color.

input colors	extra colors	schedule length	communication rounds	reference
2	0	$\infty$		see full version
2	1	$O(1)$	0	Lemma 2
3	0	$\infty$		see full version
3	1	$\infty$		Theorem 17
3	2	$O(1)$	0	Lemma 2
4	0	$\infty$		see full version
4	1	?	?	
4	2	$O(1)$	$O(1)$	see full version
4	3	$O(1)$	0	Lemma 2
5	0	$\infty$		see full version
5	1	$O(1)$	$O(1)$	see full version
5	4	$O(1)$	0	Lemma 2
6	0	$O(1)$	$O(1)$	Lemma 3 and see full version

■ **Table 5** Results: distributed recoloring in subcubic graphs.

input colors	extra colors	schedule length	communication rounds	reference
2	0	$\infty$		see full version
2	1	$O(1)$	0	Lemma 2
3	0	$\infty$		see full version
3	1	$O(\log n)$	$O(\log^2 n)$	Theorem 13
3	2	$O(1)$	0	Lemma 2
4	0	$\infty$		see full version
4	1	$O(1)$	$O(1)$	Lemma 22
4	3	$O(1)$	0	Lemma 2
5	0	$O(1)$	$O(1)$	Lemma 3



## References

- 1 Pierre Aboulker, Marthe Bonamy, Nicolas Bousquet, and Louis Esperet. Distributed coloring in sparse graphs with fewer colours. *arXiv preprint arXiv:1802.05582*, 2018.
- 2 Leonid Barenboim. Deterministic  $(\Delta + 1)$ -coloring in sublinear (in  $\Delta$ ) time in static, dynamic and faulty networks. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 345–354, 2015.
- 3 Leonid Barenboim and Michael Elkin. Distributed graph coloring: Fundamentals and recent developments. *Synthesis Lectures on Distributed Computing Theory*, 4(1):1–171, 2013.
- 4 Marthe Bonamy and Nicolas Bousquet. Recoloring graphs via tree decompositions. *European Journal of Combinatorics*, 69:200–213, 2018.
- 5 Marthe Bonamy, Nicolas Bousquet, Carl Feghali, and Matthew Johnson. On a conjecture of Mohar concerning Kempe equivalence of regular graphs. *arXiv preprint arXiv:1510.06964*, 2015.
- 6 Paul Bonsma and Luis Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theoretical Computer Science*, 410(50):5215–5226, 2009.
- 7 Paul Bonsma, Amer E Mouawad, Naomi Nishimura, and Venkatesh Raman. The complexity of bounded length graph recoloring and CSP reconfiguration. In *International Symposium on Parameterized and Exact Computation*, pages 110–121. Springer, 2014.
- 8 Nicolas Bousquet and Guillem Perarnau. Fast recoloring of sparse graphs. *European Journal of Combinatorics*, 52:1–11, 2016.
- 9 Sebastian Brandt, Juho Hirvonen, Janne H Korhonen, Tuomo Lempiäinen, Patric RJ Östergård, Christopher Purcell, Joel Rybicki, Jukka Suomela, and Przemysław Uznański. LCL problems on grids. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 101–110. ACM, 2017.
- 10 Luis Cereceda, Jan Van den Heuvel, and Matthew Johnson. Mixing 3-colourings in bipartite graphs. *European Journal of Combinatorics*, 30(7):1593–1606, 2009.
- 11 Luis Cereceda, Jan Van Den Heuvel, and Matthew Johnson. Finding paths between 3-colorings. *Journal of graph theory*, 67(1):69–82, 2011.
- 12 Yi-Jung Chang, Tsvi Kopelowitz, and Seth Pettie. An exponential separation between randomized and deterministic complexity in the LOCAL model. In *Foundations of Computer Science (FOCS)*, pages 615–624, 2016.
- 13 Yi-Jung Chang, Wenzheng Li, and Seth Pettie. An optimal distributed  $(\Delta + 1)$ -coloring algorithm? In *Proceedings of the 50th ACM Symposium on Theory of Computing (STOC)*, 2018.
- 14 Carl Feghali, Matthew Johnson, and Daniël Paulusma. A reconfigurations analogue of Brooks’ theorem and its consequences. *Journal of Graph Theory*, 83(4):340–358, 2016.
- 15 Carl Feghali, Matthew Johnson, and Daniël Paulusma. Kempe equivalence of colourings of cubic graphs. *European Journal of Combinatorics*, 59:1–10, 2017.
- 16 Mark Jerrum. A very simple algorithm for estimating the number of  $k$ -colorings of a low-degree graph. *Random Structures & Algorithms*, 7(2):157–165, 1995.
- 17 Alfred B Kempe. On the geographical problem of the four colours. *American Journal of Mathematics*, 2(3):193–200, 1879.
- 18 Michel Las Vergnas and Henri Meyniel. Kempe classes and the Hadwiger conjecture. *Journal of Combinatorial Theory, Series B*, 31(1):95–104, 1981.
- 19 Daniel C McDonald. Connectedness and Hamiltonicity of graphs on vertex colorings. *arXiv preprint arXiv:1507.05344*, 2015.
- 20 Gary L. Miller and John H. Reif. Parallel tree contraction part 1: Fundamentals. *Advances in Computing Research*, 5:47–72, 1989.

- 21 Moni Naor and Larry Stockmeyer. What can be computed locally? *SIAM Journal on Computing*, 24(6):1259–1277, 1995.
- 22 Alessandro Panconesi and Aravind Srinivasan. Improved distributed algorithms for coloring and network decomposition problems. In *Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 581–592. ACM, 1992.
- 23 Alessandro Panconesi and Aravind Srinivasan. The local nature of  $\Delta$ -coloring and its algorithmic applications. *Combinatorica*, 15(2):255–280, 1995.
- 24 Jan van den Heuvel. The complexity of change. *Surveys in Combinatorics*, 409(2013):127–160, 2013.