

**LOCAL MULTICOLORING ALGORITHMS:
COMPUTING A NEARLY-OPTIMAL TDMA SCHEDULE IN CONSTANT TIME**

FABIAN KUHN¹

¹ MIT, Computer Science and Artificial Intelligence Lab
32 Vassar St, Cambridge, MA 02139, USA
E-mail address: fkuhn@csail.mit.edu

ABSTRACT. We are given a set V of autonomous agents (e.g. the computers of a distributed system) that are connected to each other by a graph $G = (V, E)$ (e.g. by a communication network connecting the agents). Assume that all agents have a unique ID between 1 and N for a parameter $N \geq |V|$ and that each agent knows its ID as well as the IDs of its neighbors in G . Based on this limited information, every agent v must autonomously compute a set of colors $S_v \subseteq C$ such that the color sets S_u and S_v of adjacent agents u and v are disjoint. We prove that there is a deterministic algorithm that uses a total of $|C| = \mathcal{O}(\Delta^2 \log(N)/\varepsilon^2)$ colors such that for every node v of G (i.e., for every agent), we have $|S_v| \geq |C| \cdot (1-\varepsilon)/(\delta_v+1)$, where δ_v is the degree of v and where Δ is the maximum degree of G . For $N = \Omega(\Delta^2 \log \Delta)$, $\Omega(\Delta^2 + \log \log N)$ colors are necessary even to assign at least one color to every node (i.e., to compute a standard vertex coloring). Using randomization, it is possible to assign an $(1-\varepsilon)/(\delta+1)$ -fraction of all colors to every node of degree δ using only $\mathcal{O}(\Delta \log |V|/\varepsilon^2)$ colors w.h.p. We show that this is asymptotically almost optimal. For graphs with maximum degree $\Delta = \Omega(\log |V|)$, $\Omega(\Delta \log |V|/\log \log |V|)$ colors are needed in expectation, even to compute a valid coloring.

The described multicoloring problem has direct applications in the context of wireless ad hoc and sensor networks. In order to coordinate the access to the shared wireless medium, the nodes of such a network need to employ some medium access control (MAC) protocol. Typical MAC protocols control the access to the shared channel by time (TDMA), frequency (FDMA), or code division multiple access (CDMA) schemes. Many channel access schemes assign a fixed set of time slots, frequencies, or (orthogonal) codes to the nodes of a network such that nodes that interfere with each other receive disjoint sets of time slots, frequencies, or code sets. Finding a valid assignment of time slots, frequencies, or codes hence directly corresponds to computing a multicoloring of a graph G . The scarcity of bandwidth, energy, and computing resources in ad hoc and sensor networks, as well as the often highly dynamic nature of these networks require that the multicoloring can be computed based on as little and as local information as possible.

1. Introduction

In this paper, we look at a variant of the standard vertex coloring problem that we name graph *multicoloring*. Given an n -node graph $G = (V, E)$, the goal is to assign a set S_v of colors to each node $v \in V$ such that the color sets S_u and S_v of two adjacent nodes $u \in V$ and $v \in V$ are disjoint while at the same time, the fraction of colors assigned to each node is as large as possible and the

Key words and phrases: distributed algorithms, graph coloring, local algorithms, medium access control, multicoloring, TDMA, wireless networks.

For space reasons, most proofs are omitted from this extended abstract. A full version can be received from the author's web site at <http://people.csail.mit.edu/fkuhn/publications/multicoloring.pdf>.



© F. Kuhn
© Creative Commons Attribution-NoDerivs License

total number of colors used is as small as possible. In particular, we look at the following *distributed* variant of this multicoloring problem. Each node has a unique identifier (ID) between 1 and N for an integer parameter $N \geq n$. The nodes are *autonomous agents* and we assume that every agent has only very limited, *local* information about G . Specifically, we assume that every node $v \in V$ merely knows its own ID as well as the IDs of all its neighbors. Based on this local information, every node v needs to compute a color set S_v such that the color sets computed by adjacent nodes are disjoint. Since our locality condition implies that every node is allowed to communicate with each neighbor only once, we call such a distributed algorithm a *one-shot algorithm*.

We prove nearly tight upper and lower bounds for deterministic and randomized algorithms solving the above distributed multicoloring problem. Let Δ be the largest degree of G . We show that for every $\varepsilon \in (0, 1)$, there is a deterministic multicoloring algorithm that uses $\mathcal{O}(\Delta^2 \log(N)/\varepsilon^2)$ colors and assigns a $(1 - \varepsilon)/(\delta + 1)$ -fraction of all colors to each node of degree δ . Note that because a node v of degree δ does not know anything about the topology of G (except that itself has δ neighbors), no one-shot multicoloring algorithm can assign more than a $1/(\delta + 1)$ -fraction of the colors to all nodes of degree δ (the nodes could be in a clique of size $\delta + 1$). The upper bound proof is based on the probabilistic method and thus only establishes the existence of an algorithm. We describe an algebraic construction yielding an explicit algorithm that achieves the same bounds up to polylogarithmic factors. Using $\mathcal{O}(\Delta^2 \log^2 N)$ colors, for a value $\varepsilon > 0$, the algorithm assigns a $\varepsilon/\mathcal{O}(\delta^{1+\varepsilon} \log N)$ -fraction of all colors to nodes of degree δ . At the cost of using $\mathcal{O}(\Delta^{\log^* N} \log N)$ colors, it is even possible to improve the fraction of colors assigned to each node by a factor of $\log N$. The deterministic upper bound results are complemented by a lower bound showing that if $N = \Omega(\Delta^2 \log \Delta)$, even for the standard vertex coloring problem, every deterministic one-shot algorithm needs to use at least $\Omega(\Delta^2 + \log \log N)$ colors.

If we allow the nodes to use randomization (and only require that the claimed bounds are obtained with high probability), we can do significantly better. In a randomized one-shot algorithm, we assume that every node can compute a sequence of random bits at the beginning of an algorithm and that nodes also know their own random bits as well as the random bits of the neighbors when computing the color set. We show that for $\varepsilon \in (0, 1)$, with high probability, $\mathcal{O}(\Delta \log(n)/\varepsilon^2)$ colors suffice to assign a $(1 - \varepsilon)/(\delta + 1)$ -fraction of all colors to every node of degree δ . If $\log n \leq \Delta \leq n^{1-\varepsilon}$ for a constant $\varepsilon > 0$, we show that every randomized one-shot algorithm needs at least $\Omega(\Delta \log n / \log \log n)$ colors. Again, the lower bound even holds for standard vertex coloring algorithms where every node only needs to choose a single color.

Synchronizing the access to a common resource is a typical application of coloring in networks. If we have a c -coloring of the network graph, we can partition the resource (and/or time) into c parts and assign a part to each node v depending on v 's color. In such a setting, it seems natural to use a multicoloring instead of a standard vertex coloring and assign more than one part of the resource to every node. This allows to use the resource more often and thus more efficiently.

The most prominent specific example of this basic approach occurs in the context of media access control (MAC) protocols for wireless ad hoc and sensor networks. These networks consist of autonomous wireless devices that communicate with each other by the use of radio signals. If two or more close-by nodes transmit radio signals at the same time, a receiving node only hears the superposition of all transmitted signals. Hence, simultaneous transmissions of close-by nodes interfere with each other and we thus have to control the access to the wireless channel. A standard way to avoid interference between close-by transmissions is to use a time (TDMA), frequency (FDMA), or code division multiple access (CDMA) scheme to divide the channel among the nodes. A TDMA protocol divides the time into time slots and assigns different time slots to conflicting nodes. When using FDMA, nodes that can interfere with each other are assigned different frequencies, whereas a CDMA scheme uses different (orthogonal) codes for interfering nodes. Classically,

TDMA, FDMA, and CDMA protocols are implemented by a standard vertex coloring of the graph induced by the interference relations. In all three cases, it would be natural to use the more general multicoloring problem in order to achieve a more effective use of the wireless medium. Efficient TDMA schedules, FDMA frequency assignments, or CDMA code assignments are all directly obtained from a multicoloring of the interference graph where the fraction of colors assigned to each nodes is as large as possible. It is also natural to require that the total number of colors is small. This keeps the length of a TDMA schedule or the total number of frequencies or codes small and thus helps to improve the efficiency and reduce unnecessary overhead of the resulting MAC protocols.

In contrast to many wired networks, wireless ad hoc and sensor networks typically consist of small devices that have limited computing and storage capabilities. Because these devices operate on batteries, wireless nodes also have to keep the amount of computation and especially communication to a minimum in order to save energy and thus increase their lifetime. As the nodes of an ad hoc or sensor network need to operate without central control, everything that is computed, has to be computed by a distributed algorithm by the nodes themselves. Coordination between the nodes is achieved by exchanging messages. Because of the resource constraints, these distributed algorithms need to be as simple and efficient as possible. The messages transmitted and received by each node should be as few and as short as possible. Note that because of interference, the bandwidth of each local region is extremely limited. Typically, for a node v , the time needed to even receive a single message from all neighbors is proportional to the degree of v (see e.g. [19]). As long as the information provided to each node is symmetric, it is clear that every node needs to know the IDs of all adjacent nodes in G in order to compute a reasonably good multicoloring of G . Hence, the one-shot multicoloring algorithms considered in this paper base their computations on the minimum information needed to compute a non-trivial solution to the problem. Based on the above observations, even learning the IDs of all neighbors requires quite a bit of time and resources. Hence, acquiring significantly more information might already render an algorithm inapplicable in practice.¹

As a result of the scarcity of resources, the size and simplicity of the wireless devices used in sensor networks, and the dependency of the characteristic of radio transmissions on environmental conditions, ad hoc and sensor networks are much less stable than usual wired networks. As a consequence, the topology of these networks (and of their interference graph) can be highly dynamic. This is especially true for ad hoc networks, where it is often even assumed that the nodes are mobile and thus can move in space. In order to adapt to such dynamic conditions, a multicoloring needs to be recomputed periodically. This makes the resource and time efficiency of the used algorithms even more important. This is particularly true for the locality of the algorithms. If the computation of every node only depends on the topology of a close-by neighborhood, dynamic changes also only affect near-by nodes.

The remainder of the paper is organized as follows. In Section 2, we discuss related work. The problem is formally defined in Section 3. We present the deterministic and randomized upper bounds in Section 4 and the lower bounds in Section 5.

2. Related Work

There is a rich literature on distributed algorithms to compute classical vertex colorings (see e.g. [1, 4, 11, 15, 16, 21]). The paper most related to the present one is [15]. In [15], deterministic algorithms for the standard coloring problem in the same distributed setting are studied (i.e., every

¹It seems that in order to achieve a significant improvement on the multicolorings computed by the algorithms presented in this paper, every node would need much more information. Even if every node knows its complete $O(\log \Delta)$ -neighborhood, the best deterministic coloring algorithm that we are aware of needs $\Theta(\Delta^2)$ colors.

node has to compute its color based on its ID and the IDs of its neighbors). The main result is a $\Omega(\Delta^2/\log^2 \Delta)$ lower bound on the number of colors. The first paper to study distributed coloring is a seminal paper by Linial [16]. The main result of [16] is an $\Omega(\log^* n)$ -time lower bound for coloring a ring with a constant number of colors. As a corollary of this lower bound, one obtains an $\Omega(\log \log N)$ lower bound on the number of colors for deterministic one-shot coloring algorithms as studied in this paper. Linial also looks at distributed coloring algorithms for general graph and shows that one can compute an $\mathcal{O}(\Delta^2)$ -coloring in time $\mathcal{O}(\log^* n)$. In order to color a general graph with less colors, the best known distributed algorithms are significantly slower.² Using randomization, an $\mathcal{O}(\Delta)$ -coloring can be obtained in time $\mathcal{O}(\sqrt{\log n})$ [14]. Further, the fastest algorithm to obtain a $(\Delta + 1)$ -coloring is based on an algorithm to compute a maximal independent set by Luby [17] and on a reduction described in [16] and has time complexity $\mathcal{O}(\log n)$. The best known deterministic algorithms to compute a $(\Delta + 1)$ -coloring have time complexities $2^{\mathcal{O}(\sqrt{\log n})}$ and $\mathcal{O}(\Delta \log \Delta + \log^* n)$ and are described in [21] and [15], respectively. For special graph classes, there are more efficient deterministic algorithms. It has long been known that in rings [4] and bounded degree graphs [11, 16], a $(\Delta + 1)$ -coloring can be computed in time $\mathcal{O}(\log^* n)$. Very recently, it has been shown that this also holds for the much larger class of graphs with bounded local independent sets [26]. In particular, this graph class contains all graph classes that are typically used to model wireless ad hoc and sensor networks. Another recent result shows that graphs of bounded arboricity can be colored with a constant number of colors in time $\mathcal{O}(\log n)$ [3].

Closely related to vertex coloring algorithms are distributed algorithms to compute edge colorings [5, 12, 22]. In a seminal paper, Naor and Stockmeyer were the first to look at distributed algorithms where all nodes have to base their decisions on constant neighborhoods [20]. It is shown that a weak coloring with $f(\Delta)$ colors (every node needs to have a neighbor with a different color) can be computed in time 2 if every vertex has an odd degree. Another interesting approach is taken in [9] where the complexity of distributed coloring is studied in case there is an oracle that gives some nodes a few bits of extra information.

There are many papers that propose to use some graph coloring variant in order to compute TDMA schedules and FDMA frequency or CDMA code assignments (see e.g. [2, 10, 13, 18, 24, 25, 27]). Many of these papers compute a vertex coloring of the network graph such that nodes at distance at most 2 have different colors. This guarantees that no two neighbors of a node use the same time slot, frequency, or code. Some of the papers also propose to construct a TDMA schedule by computing an edge coloring and using different time slots for different edges. Clearly, it is straight-forward to use our algorithms for edge colorings, i.e., to compute a multicoloring of the line graph. With the exception of [13] all these papers compute a coloring and assign only one time slot, frequency, or code to every node or edge. In [13], first, a standard coloring is computed. Based on this coloring, an improved slot assignment is constructed such that in the end, the number of slots assigned to a node is inversely proportional to the number of colors in its neighborhood.

3. Formal Problem Description

3.1. Mathematical Preliminaries

Throughout the paper, we use $\log(\cdot)$ to denote logarithms to base 2 and $\ln(\cdot)$ to denote natural logarithms, respectively. By $\log^{(i)} x$ and by $\ln^{(i)} x$, we denote the i -fold applications of the logarithm functions \log and \ln to x , respectively³. The log star function is defined as $\log^* n :=$

²In [6], it is claimed that an $\mathcal{O}(\Delta)$ coloring can be computed in time $\mathcal{O}(\log^*(n/\Delta))$. However, the argumentation in [6] has a fundamental flaw that cannot be fixed [23].

³We have $\log^{(0)} x = \ln^{(0)} x = x$, $\log^{(i+1)} x = \log(\log^{(i)} x)$, and $\ln^{(i+1)} x = \ln(\ln^{(i)} x)$. Note that we also use $\log^i x = (\log x)^i$ and $\ln^i x = (\ln x)^i$.

$\min_i \{\log^{(i)} n \leq 1\}$. We also use the following standard notations. For an integer $n \geq 1$, $[n] = \{1, \dots, n\}$. For a finite set Ω and an integer $k \in \{0, \dots, |\Omega|\}$, $\binom{\Omega}{k} = \{S \in 2^\Omega : |S| = k\}$. The term with high probability (w.h.p.) means with probability at least $1 - 1/n^c$ for a constant $c \geq 1$.

3.2. Multicoloring

The multicoloring problem that was introduced in Section 1 can be formally defined as follows.

Definition 3.1 (Multicoloring). An $(\rho(\delta), k)$ -multicoloring γ of a graph $G = (V, E)$ is a mapping $\gamma : V \rightarrow 2^{[k]}$ that assigns a set $\gamma(v) \subset [k]$ of colors to each node v of G such that $\forall \{u, v\} \in E : \gamma(u) \cap \gamma(v) = \emptyset$ and such that for every node $v \in V$ of degree δ , $|\gamma(v)|/k \geq \rho(\delta)/(\delta + 1)$.

We call $\rho(\delta)$ the *approximation ratio* of a $(\rho(\delta), k)$ -multicoloring. Because in a one-shot algorithm (cf. the next section for a formal definition), a node of degree δ cannot distinguish G from $K_{\delta+1}$, the approximation ratio of every one-shot algorithm needs to be at most 1.

The multicoloring problem is related to the fractional coloring problem in the following way. Assume that every node is assigned the same number c of colors and that the total number of colors is k . Taking every color with fraction $1/c$ then leads to a fractional (k/c) -coloring of G . Hence, in this case, k/c is lower bounded by the fractional chromatic number $\chi_f(G)$ of G .

3.3. One-Shot Algorithms

As outlined in the introduction, we are interested in local algorithms to compute multicolorings of an n -node graph $G = (V, E)$. For a parameter $N \geq n$, we assume that every node v has a unique ID $x_v \in [N]$. In deterministic algorithms, every node has to compute a color set based on its own ID as well as the IDs of its neighbors. For randomized algorithms, we assume that nodes also know the random bits of their neighbors. Formally, a one-shot algorithm can be defined as follows.

Definition 3.2 (One-Shot Algorithm). We call a distributed algorithm a one-shot algorithm if every node v performs (a subset of) the following three steps:

1. Generate sequence R_v of random bits (deterministic algorithms: $R_v = \emptyset$)
2. Send x_v, R_v to all neighbors
3. Compute solution based on x_v, R_v , and the received information

Assume that G is a network graph such that two nodes u and v can directly communicate with each other iff they are connected by an edge in G . In the standard *synchronous message passing* model, time is divided into rounds and in every round, every node of G can send a message to each of its neighbors. One-shot algorithms then exactly correspond to computations that can be carried out in a single communication round.

For deterministic one-shot algorithms, the output of every node v is a function of v 's ID x_v and the IDs of v 's neighbors. We call this information on which v bases its decisions, the *one-hop view* of v .

Definition 3.3 (One-Hop View). Consider a node v with ID x_v and let Γ_v be the set of IDs of the neighbors of v . We call the pair (x_v, Γ_v) the one-hop view of v .

Let (x_u, Γ_u) and (x_v, Γ_v) be the one-hop views of two adjacent nodes. Because u and v are neighbors, we have $x_u \in \Gamma_v$ and that $x_v \in \Gamma_u$. It is also not hard to see that

$$\forall x_u, x_v \in [N] \text{ and } \forall \Gamma_u, \Gamma_v \in 2^{[N]} \text{ such that } x_u \neq x_v, x_u \in \Gamma_v \setminus \Gamma_u, x_v \in \Gamma_u \setminus \Gamma_v, \quad (3.1)$$

there is a labeled graph that has two adjacent nodes u and v with one-hop views (x_u, Γ_u) and (x_v, Γ_v) , respectively. Assume that we are given a graph with maximum degree Δ (i.e., for all

one-hop views (x_v, Γ_v) , we have $|\Gamma_v| \leq \Delta$. A one-shot vertex coloring algorithm maps every possible one-hop view to a color. A correct coloring algorithm must assign different colors to two one-hop views (x_u, Γ_u) and (x_v, Γ_v) iff they satisfy Condition (3.1). This leads to the definition of the *neighborhood graph* $\mathcal{N}_1(N, \Delta)$ [15] (the general notion of neighborhood graphs has been introduced in [16]). The nodes of $\mathcal{N}_1(N, \Delta)$ are all one-hop views (x_v, Γ_v) with $|\Gamma_v| \leq \Delta$. There is an edge between (x_u, Γ_u) and (x_v, Γ_v) iff the one-hop views satisfy Condition (3.1). Hence, a one-shot coloring algorithm must assign different colors to two one-hop views iff they are neighbors in $\mathcal{N}_1(N, \Delta)$. The number of colors that are needed to properly color graphs with maximum degree Δ by a one-shot algorithm therefore exactly equals the chromatic number $\chi(\mathcal{N}_1(N, \Delta))$ of the neighborhood graph (see [15, 16] for more details). Similarly, a one-shot $(\rho(\delta), k)$ -multicoloring algorithm corresponds to a $(\rho(\delta), k)$ -multicoloring of the neighborhood graph.

4. Upper Bounds

In this section, we prove all the upper bounds claimed in Section 1. We first prove that an efficient deterministic one-shot multicoloring algorithm exists in Section 4.1. Based on similar ideas, we derive an almost optimal randomized algorithm in Section 4.2. Finally, in Section 4.3, we introduce constructive methods to obtain one-shot multicoloring algorithms. For all algorithms, we assume that the nodes know the size of the ID space N as well as Δ , an upper bound on the largest degree in the network. It certainly makes sense that nodes are aware of the used ID space. Note that it is straight-forward to see that there cannot be a non-trivial solution to the one-shot multicoloring problem if the nodes do not have an upper bound on the maximum degree in the network.

4.1. Existence of an Efficient Deterministic Algorithm

The existence of an efficient, deterministic one-shot multicoloring algorithm is established by the following theorem.

Theorem 4.1. *Assume that we are given a graph with maximum degree Δ and node IDs in $[N]$. Then, for all $0 < \varepsilon \leq 1$, there is a deterministic, one-shot $(1 - \varepsilon, \mathcal{O}(\Delta^2 \log(N)/\varepsilon^2))$ -multicoloring algorithm.*

Proof. We use permutations to construct colors as described in [15]. For $i = 1, \dots, k$, let \prec_i be a global order on the ID set $[N]$. A node v with 1-hop view (x_v, Γ_v) includes color i in its color set iff $\forall y \in \Gamma_v : x_v \prec_i y$. It is clear that with this approach the color sets of adjacent nodes are disjoint. In order to show that nodes of degree δ obtain a $\rho/(\delta + 1)$ -fraction of all colors, we need to show that for all $\delta \in [\Delta]$, all $x \in [N]$, and all $\Gamma \in \binom{[N] \setminus \{x\}}{\delta}$, for all $y \in \Gamma$, $x \prec_i y$ for at least $k\rho/(\delta + 1)$ global orders \prec_i . We use the probabilistic method to show that a set of size $k = 2(\Delta + 1)^2 \ln(N)/\varepsilon^2$ of global orders \prec_i exists such that every node of degree $\delta \in [\Delta]$ gets at least an $(1 - \varepsilon)/(\delta + 1)$ -fraction of the k colors. Such a set implies that there exists an algorithm that satisfies the claimed bounds for all graphs with maximum degree Δ and IDs in $[N]$.

Let \prec_1, \dots, \prec_k be k global orders chosen independently and uniformly at random. The probability that a node v with degree δ and 1-hop view (x_v, Γ_v) gets color i is $1/(\delta + 1)$ (note that $|\Gamma_v| = \delta$). Let X_v be the number of colors that v gets. We have $\mathbb{E}[X_v] = k/(\delta + 1) \geq k/(\Delta + 1)$. Using a Chernoff bound, we then obtain

$$\mathbb{P} \left[X_v < (1 - \varepsilon) \cdot \frac{k}{\delta + 1} \right] = \mathbb{P} [X_v < (1 - \varepsilon) \cdot \mathbb{E}[X_v]] < e^{-\varepsilon^2 \mathbb{E}[X_v]/2} \leq \frac{1}{N^{\Delta+1}}. \quad (4.1)$$

Algorithm 1 Explicit Deterministic Multicoloring Algorithm: Basic Construction

Input: one-hop view (x, Γ) , parameter $\ell \geq 0$

Output: set S of colors, initially $S = \emptyset$

- 1: **for all** $(\alpha_0, \alpha_1, \dots, \alpha_\ell) \in \mathbb{F}_{q_0} \times \mathbb{F}_{q_1} \times \dots \times \mathbb{F}_{q_\ell}$ **do**
 - 2: $\beta_{0,x} := \varphi_{0,x}(\alpha_0); \forall y \in \Gamma : \beta_{0,y} := \varphi_{0,y}(\alpha_0)$
 - 3: **for** $i := 1$ **to** ℓ **do**
 - 4: $\beta_{i,x} := \varphi_{i,\beta_{i-1,x}}(\alpha_i); \forall y \in \Gamma : \beta_{i,y} := \varphi_{i,\beta_{i-1,y}}(\alpha_i)$
 - 5: **if** $\forall y \in \Gamma : \beta_{\ell,x} \neq \beta_{\ell,y}$ **then**
 - 6: $S := S \cup (\alpha_0, \alpha_1, \dots, \alpha_\ell, \beta_{\ell,x})$
-

The total number of different possible one-hop views can be bounded as $|\mathcal{N}_1(N, \Delta)| = N \cdot \sum_{\delta=1}^{\Delta} \binom{N-1}{\delta} < N^{\Delta+1}$. By a union bound argument, we therefore get that with positive probability, for all $\delta \in [\Delta]$, all possible one-hop views (x_v, Γ_v) with $|\Gamma_v| = \delta$ get at least $(1 - \varepsilon) \cdot k/(\delta + 1)$ colors. Hence, there exists a set of k global orders on the ID set $[N]$ such that all one-hop views obtain at least the required number of colors. ■

Remark: Note that if we increase the number of permutations (i.e., the number of colors) by a constant factor, all possible one-hop views (x, Γ) with $|\Gamma| = \delta$ get a $(1 - \varepsilon)/(\delta + 1)$ -fraction of all colors w.h.p.

4.2. Randomized Algorithms

We will now show that with the use of randomization, the upper bound of Section 4.1 can be significantly improved if the algorithm only needs to be correct w.h.p. We will again use random permutations. The problem of the deterministic algorithm is that the algorithm needs to assign a large set of colors to all roughly N^Δ possible one-hop views. With the use of randomization, we essentially only have to assign colors to n randomly chosen one-hop views.

For simplicity, we assume that every node knows the number of nodes n (knowing an upper bound on n is sufficient). For an integer parameter $k > 0$, every $v \in V$ chooses k independent random numbers $x_{v,1}, \dots, x_{v,k} \in [kn^4]$ and sends these random numbers to all neighbors. We use these random numbers to induce k random permutations on the nodes. Let $\Gamma(v)$ be the set of neighbors of a node v . A node v selects all colors i for which $x_{v,i} < x_{u,i}$ for all $u \in \Gamma(v)$.

Theorem 4.2. *Choosing $k = 6(\Delta + 1) \ln(n)/\varepsilon^2$ leads to a randomized one-shot algorithm that computes a $(1 - \varepsilon, k)$ -multicoloring w.h.p.*

Remark: In the above algorithm, every node has to generate $\mathcal{O}(\Delta \log^2(n)/\varepsilon^2)$ random bits and send these bits to the neighbors. Using a (non-trivial) probabilistic argument, it is possible to show that the same result can be achieved using only $\mathcal{O}(\log n)$ random bits per node.

4.3. Explicit Algorithms

We have shown in Section 4.1 that there is a deterministic one-shot algorithm that almost matches the lower bound (cf. Theorem 5.2). Unfortunately, the techniques of Section 4.1 do not yield an explicit algorithm. In this section, we will present constructive methods to obtain a one-shot multicoloring algorithm.

We develop the algorithm in two steps. First, we construct a multicoloring where in the worst case, every node v obtains the same fraction of colors independent of v 's degree. We then show how to increase the fraction of colors assigned to low-degree nodes. For an integer parameter $\ell \geq 0$,

let q_0, \dots, q_ℓ be prime powers and let d_0, \dots, d_ℓ be positive integers such that $q_0^{d_0+1} \geq N$ and $q_i^{d_i+1} \geq q_{i-1}$ for $i \geq 1$. For a prime power q and a positive integer d , let $\mathcal{P}(q, d)$ be the set of all q^{d+1} polynomials of degree at most d in $\mathbb{F}_q[z]$, where \mathbb{F}_q is the finite field of order q . We assume that we are given an injection φ_0 from the ID set $[N]$ to the polynomials in $\mathcal{P}(q_0, d_0)$ and injections φ_i from $\mathbb{F}_{q_{i-1}}$ to $\mathcal{P}(q_i, d_i)$ for $i \geq 1$. For a value x in the respective domain, let $\varphi_{i,x}$ be the polynomial assigned to x by injection φ_i . The first part of the algorithm is an adaptation of a technique used in a coloring algorithm described in [16] that is based on an algebraic construction of [7]. There, a node v with one-hop view (x, Γ) selects a color $(\alpha, \varphi_{0,x}(\alpha))$, where $\alpha \in \mathbb{F}_{q_0}$ is a value for which $\varphi_{0,x}(\alpha) \neq \varphi_{0,y}(\alpha)$ for all $y \in \Gamma$ (we have to set q_0 and d_0 such that this is always possible). We make two modifications to this basic algorithm. Instead of only selecting one value $\alpha \in \mathbb{F}_{q_0}$ such that $\forall y \in \Gamma : \varphi_{0,x}(\alpha) \neq \varphi_{0,y}(\alpha)$, we select all values α for which this is true. We then use these values recursively (as if $\varphi_{i,x}(\alpha_i)$ was the ID of v) ℓ times to reduce the dependence of the approximation ratio of the coloring on N . The details of the first step of the algorithm are given by Algorithm 1.

Lemma 4.3. *Assume that for $0 \leq i \leq \ell$, $q_i \geq f_i \Delta d_i$ where $f_i > 1$. Then, Algorithm 1 constructs a multicoloring with $q_\ell \cdot \prod_{i=0}^{\ell} q_i$ colors where every node at least receives a λ/q_ℓ -fraction of all colors where $\lambda = \prod_{i=0}^{\ell} (1 - 1/f_i)$.*

Proof. All colors that are added to the color set in line 6 are from $\mathbb{F}_{q_0} \times \mathbb{F}_{q_1} \times \dots \times \mathbb{F}_{q_\ell} \times \mathbb{F}_{q_\ell}$. It is therefore clear that the number of different colors is $q_\ell \cdot \prod_{i=0}^{\ell} q_i$ as claimed. From the condition in line 5, it also follows that the color sets of adjacent nodes are disjoint.

To determine the approximation ratio, we count the number of colors, a node v with one-hop view (x, Γ) gets. First note that the condition in line 5 of the algorithm implies that (and is therefore equivalent to demand that) $\beta_{i,x} \neq \beta_{i,y}$ for all $y \in \Gamma$ and for all $i \in \{0, \dots, \ell\}$ because $\beta_{i,x} = \beta_{i,y}$ implies $\beta_{j,x} = \beta_{j,y}$ for all $j \geq i$. We therefore need to count the number of $(\alpha_0, \dots, \alpha_\ell) \in \mathbb{F}_{q_0} \times \dots \times \mathbb{F}_{q_\ell}$ for which $\beta_{i,x} \neq \beta_{i,y}$ for all $i \in \{0, \dots, \ell\}$ and all $y \in \Gamma$. We prove by induction on i that for $i < \ell$, there are at least $\prod_{j=0}^i q_j \cdot (1 - 1/f_j)$ tuples $(\alpha_0, \dots, \alpha_i) \in \mathbb{F}_{q_0} \times \dots \times \mathbb{F}_{q_i}$ with $\beta_{j,x} \neq \beta_{j,y}$ for all $j \leq i$. Let us first prove the statement for $i = 0$. Because the IDs of adjacent nodes are different, we know that $\varphi_{0,x} \neq \varphi_{0,y}$ for all $y \in \Gamma$. Two different degree d_0 polynomials can be equal at at most d_0 values. Hence, for every $y \in \Gamma$, $\varphi_{0,x}(\alpha) = \varphi_{0,y}(\alpha)$ for at most d_0 values α . Thus, since $|\Gamma| \leq \Delta$, there are at least $q_0 - \Delta d_0 \geq q_0 \cdot (1 - 1/f_0)$ values α for which $\varphi_{0,x} \neq \varphi_{0,y}$ for all $y \in \Gamma$. This establishes the statement for $i = 0$. For $i > 0$, the argument is analogous. Let $(\alpha_0, \dots, \alpha_{i-1}) \in \mathbb{F}_{q_0} \times \dots \times \mathbb{F}_{q_{i-1}}$ be such that $\beta_{j,x} \neq \beta_{j,y}$ for all $y \in \Gamma$ and all $j < i$. Because $\beta_{i-1,x} \neq \beta_{i-1,y}$, we have $\varphi_{i,x} \neq \varphi_{i,y}$. Thus, with the same argument as for $i = 0$, there are at least $q_i \cdot (1 - 1/f_i)$ values α_i such that $\beta_{i,x} \neq \beta_{i,y}$ for all $y \in \Gamma$. Therefore, the number of colors in the color set of every node is at least $\prod_{i=0}^{\ell} q_i \cdot (1 - 1/f_i) = \lambda \cdot \prod_{i=0}^{\ell} q_i$. This is a (λ/q_ℓ) -fraction of all colors. ■

The next lemma specifies how the values of q_i , d_i , and f_i can be chosen to obtain an efficient algorithm.

Lemma 4.4. *Let ℓ be such that $\ln^{(\ell)} N > \max\{e, \Delta\}$. For $0 \leq i \leq \ell$, we can then choose q_i , d_i , and f_i such that Algorithm 1 computes a multicoloring with $\mathcal{O}(\ell \Delta)^{\ell+2} \cdot \log_\Delta N \cdot \log_\Delta \ln^{(\ell)} N$ colors and such that every node gets at least a $1/(4e^{9/4} \Delta \lceil \log_\Delta \ln^{(\ell)} N \rceil)$ -fraction of all colors.*

The number of colors that Algorithm 1 assigns to nodes with degree almost Δ is close to optimal even for small values of ℓ . If we choose $\ell = \Theta(\log^* N - \log^* \Delta)$, nodes of degree $\Theta(\Delta)$ even receive at least a (d/Δ) -fraction of all colors for some constant d . Because the number of colors assigned to a node v is independent of v 's degree, however, the coloring of Algorithm 1 is far

Algorithm 2 Explicit Deterministic Multicoloring Algorithm: Small Number of Colors

Input: one-hop view (x, Γ) , instances $\mathcal{A}_{2^i, N}$ for $i \in [\lceil \log \Delta \rceil]$ of Algorithm 1, parameter $\varepsilon \in [0, 1]$

Output: set S of colors, initially $S = \emptyset$

- 1: **for all** $i \in [\lceil \log \Delta \rceil]$ **do**
- 2: $\omega_i := \left\lceil (\Delta/2^{i-1})^\varepsilon \cdot |\mathcal{C}_{2^{\lceil \log \Delta \rceil}, N}| / |\mathcal{C}_{2^i, N}| \right\rceil$
- 3: **for all** $i \in \{\lceil \log |\Gamma| \rceil, \dots, \lceil \log \Delta \rceil\}$ **do**
- 4: **for all** $c \in \mathcal{C}_{2^i, N}[x, \Gamma]$ **do**
- 5: **for all** $j \in [\omega_i]$ **do** $S := S \cup (c, i, j)$

from optimal for low-degree nodes. In the following, we show how to improve the algorithm in this respect.

Let $\mathcal{A}_{\Delta, N}$ be an instance of Algorithm 1 for nodes with degree at most Δ and let $\mathcal{C}_{\Delta, N}$ be the color set of $\mathcal{A}_{\Delta, N}$. Further, for a one-hop view (x, Γ) , let $\mathcal{C}_{\Delta, N}[x, \Gamma]$ be the colors assigned to (x, Γ) by Algorithm $\mathcal{A}_{\Delta, N}$. We run instances $\mathcal{A}_{2^i, N}$ for all $i \in [\lceil \log \Delta \rceil]$. A node v with degree δ chooses the colors of all instances for which $2^i \geq \delta$. In order to achieve the desired trade-offs, we introduce an integer weight ω for each color c , i.e., instead of adding color c , we add colors $(1, c), \dots, (\omega, c)$. The details are given by Algorithm 2. The properties of Algorithm 2 are summarized by the next theorem. The straight-forward proof is omitted.

Theorem 4.5. *Assume that in the instances of Algorithm 1, the parameter ℓ is chosen such that for all Δ , $\mathcal{A}_{\Delta, N}$ assigns at least a $f(N)/\Delta$ -fraction of the colors to every node. Then, for a parameter $\varepsilon \in [0, 1]$, Algorithm 2 computes a $(\Omega(f(N)\varepsilon/\delta^\varepsilon), \mathcal{O}(|\mathcal{C}_{2\Delta, N}| \cdot \Delta^\varepsilon/\varepsilon))$ -multicoloring.*

Corollary 4.6. *Let $\varepsilon \in [0, 1]$ and $\ell \geq 0$ be a fixed constant in all used instances of Algorithm 1. Then, Algorithm 2 computes an $(\varepsilon/\mathcal{O}(\delta^\varepsilon \log_\Delta \ln^{(\ell)} N), \mathcal{O}(\Delta^{\ell+2} \cdot \log_\Delta N \cdot \log_\Delta \ln^{(\ell)} N))$ -multicoloring. In particular, choosing $\ell = 0$ leads to an $(\varepsilon/\mathcal{O}(\delta^\varepsilon \log_\Delta N), \mathcal{O}(\Delta^2 \log_\Delta^2 N))$ -multicoloring. Taking the maximum possible value for ℓ in all used instances of Algorithm 1 yields an $(\varepsilon/\mathcal{O}(\delta^\varepsilon), \Delta^{\mathcal{O}(\log^* N - \log^* \Delta)} \cdot \log_\Delta N)$ -multicoloring.*

5. Lower Bounds

In this section, we give lower bounds on the number of colors required for one-shot multicoloring algorithms. In fact, we even derive the lower bounds for algorithms that need to assign only one color to every node, i.e., the results even hold for standard coloring algorithms.

It has been shown in [15] that every deterministic one-shot c -coloring algorithm \mathcal{A} can be interpreted as a set of c antisymmetric relations on the ID set $[N]$. Assume that \mathcal{A} assigns a color from a set C with $|C| = c$ to every one-hop view (x, Γ) . For every color $\alpha \in C$, there is a relation \triangleleft_α such that for all $x, y \in [N]$ $x \triangleleft_\alpha y \vee y \triangleleft_\alpha x$. Algorithm \mathcal{A} can assign color $\alpha \in C$ to a one-hop view (x, Γ) iff $\forall y \in \Gamma : x \triangleleft_\alpha y$.

For $\alpha \in C$, let $\text{Bad}_\alpha(x) := \{y \in [N] : x \triangleleft_\alpha y\}$ be the set of IDs that must not be adjacent to an α -colored node with ID x . To show that there is no deterministic, one-shot c -coloring algorithm, we need to show that for every c antisymmetric relations $\triangleleft_{\alpha_1}, \dots, \triangleleft_{\alpha_c}$ on $[N]$, there is a one-hop view (x, Γ) such that $\forall i \in [c] : \Gamma \cap \text{Bad}_{\alpha_i}(x) \neq \emptyset$. The following lemma is a generalization of Lemma 4.5 in [15] and key for the deterministic and the randomized lower bounds. As the proof is along the same lines as the proof of Lemma 4.5 in [15], it is omitted here.

Lemma 5.1. *Let $X \subseteq [N]$ be a set of IDs and let t_1, \dots, t_ℓ and k_1, \dots, k_ℓ be positive integers such that*

$$t_i \cdot (\lambda(|X| - c)t_i - c) > 2c(k_i - 1) \quad \text{for } 1 \leq i \leq \ell \text{ and a parameter } \lambda \in [0, 1].$$

Then there exists an ID set $X' \subseteq X$ with $|X'| > (1 - \ell \cdot \lambda) \cdot (|X| - c)$ such that for all $i \in [\ell]$,

$$\forall x \in X', \forall \alpha_1, \dots, \alpha_{t_i} \in C : \sum_{j=1}^{t_i} |\text{Bad}_{\alpha_j}(x) \cap X| \geq k_i, \quad \forall x \in X', \forall \alpha \in C : \text{Bad}_\alpha(x) \cap X \neq \emptyset.$$

Based on several applications of Lemma 5.1 (and based on an $\Omega(\log \log N)$ lower bound in [16]), it is possible to derive an almost tight lower bound for deterministic one-shot coloring algorithms. Due to lack of space, we only state the result here.

Theorem 5.2. *If $N = \Omega(\Delta^2 \log \Delta)$, every deterministic one-shot coloring algorithm needs at least $\Omega(\Delta^2 + \log \log N)$ colors.*

5.1. Randomized Lower Bound

To obtain a lower bound for randomized multicoloring algorithms, we can again use the tools derived for the deterministic lower bound by applying Yao's principle. On a worst-case input, the best randomized algorithm cannot perform better than the best deterministic algorithm for a given random input distribution. Choosing the node labeling at random allows to again only consider deterministic algorithms.

We assume that the n nodes are assigned a random permutation of the labels $1, \dots, n$ (i.e., every label occurs exactly once). Note that because we want to prove a lower bound, assuming the most restricted possible ID space makes the bound stronger. For an ID $x \in [n]$, we sort all colors $\alpha \in C$ by increasing values of $|\text{Bad}_\alpha(x)|$ and let $\alpha_{x,i}$ be the i^{th} color in this sorted order. Further, for $x \in [n]$, we define $b_{x,i} := |\text{Bad}_{\alpha_{x,i}}(x)|$. In the following, we assume that

$$c = \kappa \cdot \frac{\Delta \lfloor \ln n \rfloor}{\lfloor \ln \ln n \rfloor + 2} \quad \text{and} \quad n \geq 12 \quad \text{and} \quad n \geq \Delta \cdot \ln n \quad (5.1)$$

for a constant $0 < \kappa \leq 1$ that will be determined later. By applying Lemma 5.1 in different ways, the next lemma gives lower bounds on the values of $b_{x,i}$ for $n/2$ IDs $x \in [n]$.

Lemma 5.3. *Assume that c and n are as given by Equation (5.1) and let $0 < \rho < 1/3$ be a positive constant. Further, let $\tilde{t} = \lceil \rho \ln n / \ln \ln n \rceil$ and $t_i = 2^{i-1} \cdot \lfloor \ln n \rfloor$ for $1 \leq i \leq \ell$ where $\ell = \lfloor \ln \ln n \rfloor + 2$. Then, for at least $n/2$ of all IDs $x \in [n]$, we have*

$$b_{x,1} \geq \frac{\ln \ln n}{44\kappa \cdot \ln n} \cdot \frac{n}{\Delta} - 1, \quad b_{x,\tilde{t}} \geq \frac{\rho}{48\kappa} \cdot \frac{n}{\Delta} - \frac{1}{2}, \quad b_{x,t_i} \geq 2^{i-1} \cdot \left(\frac{1}{8\kappa} \cdot \frac{n}{\Delta} - \frac{1}{2} \right) \quad \text{for } 1 \leq i \leq \ell.$$

In order to prove the lower bound, we want to show that for a randomly chosen one-hop view (x, Γ) with $|\Gamma| = \Delta$, the probability that there is a color $\alpha \in C$ for which $\Gamma \cap \text{Bad}_\alpha(x) = \emptyset$ is sufficiently small. Instead of directly looking at random one-hop views (x, Γ) with $|\Gamma| = \Delta$, we first look at one-hop views with $|\Gamma| \approx \Delta/e$ that are constructed as follows. Let $X \subseteq [n]$ be the set of IDs x of size $|X| \geq n/2$ for which the bounds of Lemma 5.3 hold. We choose x_R uniformly at random from X . The remaining $n - 1$ IDs are independently added to a set Γ_R with probability $p = \frac{\Delta}{en}$. For a color $\alpha \in C$, let \mathcal{E}_α be the event that $\Gamma_R \cap \text{Bad}_\alpha(x_R) \neq \emptyset$, i.e., \mathcal{E}_α is the event that color α cannot be assigned to the randomly chosen one-hop view (x_R, Γ_R) .

Lemma 5.4. *The probability that the randomly chosen one-hop view cannot be assigned one of the c colors in C is bounded by*

$$\mathbb{P} \left[\bigcap_{\alpha \in C} \mathcal{E}_\alpha \right] \geq \prod_{\alpha \in C} \mathbb{P}[\mathcal{E}_\alpha] \geq \prod_{\alpha \in C} \left(1 - e^{-\frac{\Delta}{en} \cdot |\text{Bad}_\alpha(x_R)|} \right) = \prod_{i=1}^c \left(1 - e^{-\frac{\Delta \cdot b_{x_R, i}}{en}} \right).$$

Proof. Note first that for $\alpha \in C$, we have

$$\mathbb{P}[\overline{\mathcal{E}_\alpha}] = \mathbb{P}[\Gamma_R \cap \text{Bad}_\alpha(x_R) = \emptyset] = (1 - p)^{|\text{Bad}_\alpha(x_R)|} \leq e^{-p|\text{Bad}_\alpha(x_R)|} = e^{-\frac{\Delta}{en} \cdot |\text{Bad}_\alpha(x_R)|}.$$

It therefore remains to prove that the probability that all events \mathcal{E}_α occur can be lower bounded by the probability that would result for independent events. Let us denote the colors in C by $\alpha_1, \dots, \alpha_c$. We then have

$$\mathbb{P} \left[\bigcap_{\alpha \in C} \mathcal{E}_\alpha \right] = \prod_{i=1}^c \mathbb{P} \left[\mathcal{E}_{\alpha_i} \mid \bigcap_{j=1}^{i-1} \mathcal{E}_{\alpha_j} \right] \geq \prod_{i=1}^c \mathbb{P}[\mathcal{E}_{\alpha_i}]. \tag{5.2}$$

The inequality holds because the events \mathcal{E}_α are positively correlated. Knowing that an element from a set $\text{Bad}_\alpha(x_R)$ is in Γ_R cannot decrease the probability that an element from a set $\text{Bad}_{\alpha'}(x_R)$ is in Γ_R . Note that this is only true because the IDs are independently added to Γ_R . More formally, Inequality (5.2) can also directly be followed from the FKG inequality [8]. ■

For space reasons, the following two lemmas are given without proof.

Lemma 5.5. *Assume that c and n are given as in (5.1) where the constant κ is chosen sufficiently small and let $\rho > 0$ be a constant as in Lemma 5.3. There is a constant $n_0 > 0$ such that for $n \geq n_0$, $\mathbb{P} \left[\bigcap_{\alpha \in C} \mathcal{E}_\alpha \right] > \frac{1}{2n^{3\rho}}$.*

Lemma 5.6. *Let (x, Γ) be a one-hop view chosen uniformly at random from all one-hop views with $|\Gamma| = \Delta$. If $\Delta \geq e(\ln n + 2)$ and n, c , and ρ are as before, the probability that none of the c colors can be assigned to (x, Γ) is at least $1/(8n^{3\rho})$.*

In the following, we call a node u together with Δ neighbors v_1, \dots, v_Δ , a Δ -star.

Theorem 5.7. *Let G be a graph with n nodes and $2n^\varepsilon$ disjoint Δ -stars for a constant $\varepsilon > 0$. On G , every randomized one-shot coloring algorithm needs at least $\Omega(\Delta \log n / \log \log n)$ colors in expectation and with high probability.*

Proof. W.l.o.g., we can certainly assume that $n \geq n_0$ for a sufficiently large constant n_0 . We choose $\rho \leq \varepsilon/4$ and consider n^ε of the $2n^\varepsilon$ disjoint Δ -stars. Let us call these n^ε Δ -stars $S_1, \dots, S_{n^\varepsilon}$. Assume that the ID assignment of the n nodes of G is chosen uniformly at random from all ID assignments with IDs $1, \dots, n$. The IDs of the star S_1 are perfectly random. We can therefore directly apply Lemma 5.6 and obtain that the probability that the center node of S_1 gets no color is at least $1/(8n^{3\rho})$. Consider star S_2 . The IDs of the nodes of S_2 are chosen at random among the $n - \Delta - 1$ IDs that are not assigned to the nodes of S_1 . Applying Lemma 5.6 we get that the probability that S_2 does not get a color is at least $1/(8(n - \Delta - 1)^{3\rho}) \geq 1/(8n^{3\rho})$ independently of whether S_1 does get a color. The probability that the stars $S_1, \dots, S_{n^\varepsilon}$ all get a color therefore is at most

$$\prod_{i=0}^{n^\varepsilon-1} \left(1 - \frac{1}{8(n - i(\Delta + 1))^{3\rho}} \right) \leq \left(1 - \frac{1}{8n^{3\rho}} \right)^{n^\varepsilon} \leq e^{-\frac{n^\varepsilon}{8n^{3\rho}}} \leq e^{-n^\rho/8}.$$

Hence, there is a constant $\eta > 0$ such that $\eta\Delta \ln n / \ln \ln n$ colors do not suffice with probability at least $1 - e^{-n^\rho/8}$ for a positive constant ρ . The lemma thus follows. ■

References

- [1] B. Awerbuch, A. V. Goldberg, M. Luby, and S. A. Plotkin. Network decomposition and locality in distributed computation. In *Proc. of 30th Symposium on Foundations of Computer Science (FOCS)*, pages 364–369, 1989.
- [2] B. Balasundaram and S. Butenko. Graph domination, coloring and cliques in telecommunications. In M. Resende and P. Pardalos, editors, *Handbook of Optimization in Telecommunications*, pages 865–890. Springer, 2006.
- [3] L. Barenboim and M. Elkin. Sublogarithmic distributed mis algorithm for sparse graphs using nash-williams decomposition. In *Proc. of 27th ACM Symposium on Principles of Distributed Computing (PODC)*, 2008.
- [4] R. Cole and U. Vishkin. Deterministic coin tossing with applications to optimal parallel list ranking. *Information and Control*, 70(1):32–53, 1986.
- [5] A. Czygrinow, M. Hańćkowiak, and M. Karoński. Distributed $O(\delta \log n)$ -edge-coloring algorithm. In *Proc. of 9th Annual European Symposium on Algorithms (ESA)*, volume 2161 of *LNCS*, pages 345–355, 2001.
- [6] G. De Marco and A. Pelc. Fast distributed graph coloring with $O(\Delta)$ colors. In *Proc. of 12th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 630–635, 2001.
- [7] P. Erdős, P. Frankl, and Z. Füredi. Families of finite sets in which no set is covered by the union of r others. *Israel Journal of Mathematics*, 51:79–89, 1985.
- [8] C. Fortuin, J. Ginibre, and P. Kasteleyn. Correlation inequalities on some partially ordered sets. *Comm. in Mathematical Physics*, 22:89–103, 1971.
- [9] P. Fraigniaud, C. Gavoille, D. Ilcinkas, and A. Pelc. Distributed computing with advice: Information sensitivity of graph coloring. In *Proc. of 34th Int. Coll. on Automata, Languages and Programming*, 2007.
- [10] S. Gandham, M. Dawande, and R. Prakash. Link scheduling in sensor networks: Distributed edge coloring revisited. In *Proc. of 24th IEEE Conference on Computer Communications (INFOCOM)*, pages 2492–2501, 2005.
- [11] A. Goldberg, S. Plotkin, and G. Shannon. Parallel symmetry-breaking in sparse graphs. *SIAM Journal on Discrete Mathematics*, 1(4):434–446, 1988.
- [12] D. A. Grable and A. Panconesi. Nearly optimal distributed edge coloring in $O(\log \log n)$ rounds. *Random Structures and Algorithms*, 10(3):385–405, 1997.
- [13] T. Herman and S. Tixeuil. A distributed TDMA slot assignment algorithm for wireless sensor networks. In *Proc. of 1st Int. Workshop on Algorithmic Aspects of Wireless Sensor Networks*, pages 45–58, 2004.
- [14] K. Kothapalli, M. Onus, C. Scheideler, and C. Schindelhauer. Distributed coloring in $O(\sqrt{\log n})$ bit rounds. In *Proc. of 20th IEEE Int. Parallel and Distributed Processing Symposium (IPDPS)*, 2006.
- [15] F. Kuhn and R. Wattenhofer. On the complexity of distributed graph coloring. In *Proc. of 25th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 7–15, 2006.
- [16] N. Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- [17] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15:1036–1053, 1986.
- [18] S. Mecke. MAC layer and coloring. In D. Wagner and R. Wattenhofer, editors, *Algorithms for Sensor and Ad Hoc Networks*, pages 63–80, 2007.
- [19] T. Moscibroda and R. Wattenhofer. Coloring unstructured radio networks. In *Proc. of 17th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 39–48, 2005.
- [20] M. Naor and L. Stockmeyer. What can be computed locally? In *Proc. of 25th ACM Symposium on Theory of Computing (STOC)*, pages 184–193, 1993.
- [21] A. Panconesi and A. Srinivasan. On the complexity of distributed network decomposition. *Journal of Algorithms*, 20(2):581–592, 1995.
- [22] A. Panconesi and A. Srinivasan. Randomized distributed edge coloring via an extension of the chernoff-hoeffding bounds. *SIAM Journal on Computing*, 26(2):350–368, 1997.
- [23] A. Pelc. Personal communication.
- [24] S. Ramanathan. A unified framework and algorithm for channel assignment in wireless networks. *Wireless Networks*, 5:81–94, 1999.
- [25] I. Rhee, A. Warrier, J. Min, and L. Xu. DRAND: Distributed randomized TDMA scheduling for wireless ad-hoc networks. In *7th ACM Symp. on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, pages 190–201, 2006.
- [26] J. Schneider and R. Wattenhofer. A log-star distributed maximal independent set algorithm for growth-bounded graphs. In *Proc. of 27th ACM Symposium on Principles of Distributed Computing (PODC)*, 2008.
- [27] X. Zhang, J. Hong, L. Zhang, X. Shan, and V. Li. CP-TDMA: Coloring- and probability-based TDMA scheduling for wireless ad hoc networks. *IEICE Transactions on Communication*, E91-B(1):322–326, 2008.