Online Multi-Coloring with Advice [*]

Marie G. Christ Lene M. Favrholdt Kim S. Larsen

University of Southern Denmark
Odense, Denmark

{christm,lenem,kslarsen}@imada.sdu.dk

October 15, 2018

Abstract

We consider the problem of online graph multi-coloring with advice. Multi-coloring is often used to model frequency allocation in cellular networks. We give several nearly tight upper and lower bounds for the most standard topologies of cellular networks, paths and hexagonal graphs. For the path, negative results trivially carry over to bipartite graphs, and our positive results are also valid for bipartite graphs. The advice given represents information that is likely to be available, studying for instance the data from earlier similar periods of time.

1 Introduction

We consider the problem of graph multi-coloring, where each node may receive multiple requests. Whenever a node is requested, a color must be assigned to the node, and this color must be different from any color previously assigned to that node or to any of its neighbors. The goal is to use as few colors as possible. In the online version, the requests arrive one by one, and each request must be colored without any information about possible future requests. The underlying graph is known to the online algorithm in advance.

[*]Supported in part by the Danish Council for Independent Research and the Villum Foundation. An extended abstract will appear in the Twelfth Workshop on Approximation and Online Algorithms (WAOA), Lecture Notes in Computer Science. Springer, 2014.

The problem is motivated by frequency allocation in cellular networks. These networks are formed by a number of base transceiver stations, each of which covers what is referred to as a cell. Due to possible interference, neighboring cells cannot use the same frequencies. In this paper, we use classic terminology and refer to these cells as nodes in a graph where nodes are connected by an edge if they correspond to neighboring cells in the network. Frequencies can then be modeled as colors. Multiple requests for frequencies can occur in one cell and overall bandwidth is a critical resource.

Two basic models dominate in the discussion of cellular networks, the highway and the city model. The former is modeled by linear cellular networks, corresponding to paths, and the latter by hexagonal graphs. We consider the problem of multi-coloring such graphs.

1.1 Analyzing online algorithms

If A is a multi-coloring algorithm, we let $\mathrm{A}(I)$ denote the number of colors used by A on the input sequence I. When I is clear from the context, we simply write A instead of $\mathrm{A}(I)$. The quality of an online algorithm is often given in terms of the competitive ratio [35, 26]. An online multi-coloring algorithm is *c-competitive* if there exists a constant α such that for all input sequences I, $\mathrm{A}(I) \leq c\,\textsc{Opt}(I) + \alpha$. The (asymptotic) *competitive ratio* of A is the infimum over all such c. Results that can be established using $\alpha = 0$ are referred to as *strict* (or absolute). Often, it is a little unclear when one refers to an *optimal* online algorithm, whether this means that the solution produced is as good as the one produced offline or that no better online algorithm can exist. For that reason, we may use the term *strictly* 1-*competitive* to emphasize that an algorithm is as good as an optimal offline algorithm, and *optimal* to mean that no better online algorithm exists under the given conditions. Throughout, we let n denote the number of requests in a given input sequence.

1.1.1 Relaxing the concept of online

A way of relaxing the very strict and unnatural assumption that the algorithm has no information about the input sequence is to give the algorithm some *advice*. The possibly most famous online problem of paging, where no deterministic online algorithm is better than k-competitive on a cache size of k, can be solved optimally with just one bit of advice per request, saying whether to keep the requested page in cache until its next request [18, 5].

A recent trend in the analysis of online algorithms has been to consider advice, formalized under the notion of *advice complexity*, starting in [18]. Theoretically, results along these lines give some information in the direction of the hardness stemming from the problem being online, relaying information concerning how much we need to know about the future to perform better. For practical applications, the assumption that absolutely nothing is known about the future is often unrealistic, and though many problems must be addressed without knowing in which order requests arrive, quite often something is known about the sequence of requests as a whole.

This realization that input is not arbitrary (uniformly random, for instance) is not new, and work focused on locality of reference in input data has tried to capture this. Early work includes access graph results, starting in [7], and with references to additional related work in [9], but also more distributional models, such as [1], have been developed. An entirely different approach was initiated in [11] and further developed in [12, 8]. The idea behind the concept of accommodating sequences is that for many problems requiring resources, there is a close connection between the resources available and the resources required for an optimal offline algorithm, as when capacity of transportation systems are matched with expected demand. This leans itself very closely up against many of the results that we report here, where the advice needed to do better is often some information regarding the resources required by an optimal offline algorithm.

Thus, the results in this paper could have practical applications. The results establish which type of information is useful, how algorithms should be designed to exploit this information, and what the limits are for what can be obtained.

1.1.2 Modeling advice complexity

Returning to the advice complexity modeling, some problems need very little advice. On the other hand, complete information about the input or the desired output is a trivial upper bound on the amount of advice needed to be optimal. The first approach to formalizing the concept of advice measured the number of bits per request [18]. This model is well suited for some problems where information is tightly coupled with requests and the number of bits needed per request is constant. However, for most problems, we prefer the model where we simply measure the total advice needed throughout the execution of the algorithm. As also discussed in [5, 23], this model avoids some modeling issues present in the “per request” modeling, and at the same time makes it possible to derive sublinear advice requirements. Thus, we use the advice model from [23], where the online algorithm has access to an infinite advice tape, written by an offline oracle with infinite computation power. In other words, the online algorithm can ask for the answer to any question and read the answer from the tape. Competitiveness is defined and measured as usual, and the advice complexity is simply the number of bits read from the tape, i.e., the maximum index of the bits read from the advice tape.

As the advice tape is infinite, we need to specify how many bits of advice the algorithm should read and if this knowledge is not implicitly available, it has to be given explicitly in the advice string. For instance, if we want OPT as advice (the number of colors an optimal offline algorithm uses on a given sequence, for instance), then we cannot merely read $\lceil \log(\textsc{Opt}+1)\rceil$ (all $\log$s in this paper are base 2) bits, since this would require knowing something about the value of OPT. One can use a self-delimiting encoding as introduced in [20]. We use the variant from [10], defined as follows: The value of a non-negative integer X is encoded by a bit sequence, partitioned into three consecutive parts. The last part is X written in binary. The middle part gives the number of bits in the last part, written in binary. The first part gives the number of bits in the middle part, written in unary and terminated with a zero. These three parts require $\lceil \log(\lceil \log(X+1)\rceil+1)\rceil+1$, $\lceil \log(\lceil \log(X+1)\rceil+1)\rceil$, and $\lceil \log(X+1)\rceil$ bits, respectively, adding a lower-order term to the number of bits of information required by an algorithm. We define *enc*(x) to be the minimum number of bits necessary to encode a number x, and note that the encoding above is a (good) upper bound on *enc*(x).

1.2 Previous and new results

We now discuss previous work related to multi-coloring and advice complexity and then state our results. When working with online algorithms, decisions are generally irrevocable, i.e., once a color is assigned to a node, this decision is final. However, in some applications, local changes of colors may be allowed (reassignment of frequencies). This is called *recoloring*. An algorithm is *d-recoloring* if, in the process of treating a request, it may recolor up to a distance d away from the node of the request.

1.2.1 Previous results

For multi-coloring a path, the algorithm 4-BUCKET is $\frac{4}{3}$-competitive [17], and this is optimal [14]. Even with 0-recoloring allowed (that is, colors at the requested node may be changed), 4-BUCKET is optimal [15]. Furthermore, if 1-recoloring is allowed, the algorithm GREEDYOPT is strictly 1-competitive [15].

For multi-coloring bipartite graphs, the optimal asymptotic competitive ratio lies between $\frac{10}{7} \approx 1.428$ and $\frac{18-\sqrt{5}}{11} \approx 1.433$ [16].

In [13], it was shown that, for hexagonal graphs, no online algorithm can be better than $\frac{3}{2}$-competitive or have a better strict competitive ratio than 2. They also gave an algorithm, HYBRID, with an asymptotic competitive ratio of approximately 1.9 on hexagonal graphs. On k-colorable graphs, it is strictly $\frac{k+1}{2}$-competitive, and hence, it has an optimal strict competitive ratio on hexagonal graphs. Recoloring was studied in [25]: No d-recoloring algorithm for hexagonal graphs has an asymptotic competitive ratio better than $1 + \frac{1}{4(d+1)}$. For $d = 0$, the lower bound was improved to $\frac{9}{7}$. In [36], a $\frac{4}{3}$-competitive 2-recoloring algorithm is given. The best known 1-recoloring algorithm for hexagonal graphs is $\frac{33}{24}$-competitive [37]. For the offline problem of multi-coloring hexagonal graphs, no polynomial time algorithm can obtain an absolute approximation ratio better than $\frac{4}{3}$ [30, 32, 33], unless P $=$ NP.

Many other problems have been considered in the advice models, including paging [5], disjoint path allocation [2], and job shop scheduling [5], as well as k-server [4], knapsack [6], set cover [28], metrical task systems [21], and buffer management [19]. Also graph coloring has been considered, but in a very different online setting, where the graph itself is not available from the beginning. Instead, the nodes are revealed one by one and results have been obtained for paths [22], bipartite graphs [3], and 3-colorable graphs [34]. In [29], a coloring problem with restrictions going beyond the immediate neighbors is considered. Furthermore, there are interesting connections between advice and randomization and sometimes results on advice complexity can be used to obtain efficient random algorithms [5, 27, 6].

1.2.2 Our results

An overview of our results is given in Table 1. For the path, these results are nearly tight, even with upper bounds that also apply to bipartite graphs. For hexagonal graphs, note that with a linear number of advice bits, it is possible to be $\frac{4}{3}$-competitive, and the lower bound for being better than $\frac{5}{4}$-competitive is close to linear. The advice given to the algorithms is essentially (an approximation of) OPT or the maximum number of requests given to any clique in the graph. For the underlying problem of frequency allocation, guessing these values based on previous data may not be unrealistic.

	Ratio	Lower	Type	Thm	Upper	Type	Thm
Paths	1	$\log n - 2$	s	1	$\log n + O(\log\log n)$	s	4
	$1+\frac{1}{2^b}$	$b-2$	a	2	$b+1+O(\log\log n)$	s	5
	$< \frac{4}{3}$	$\omega(1)$	a	3			
Hexagonal	1				$(n+1)\lceil\log n\rceil$	s	8
	$< \frac{5}{4}$	$\Omega(n)$	a	7			
	$\frac{4}{3}$				$n+2\|V\|$	a	10
	$\frac{3}{2}$	$\left\lfloor\frac{n-1}{3}\right\rfloor$	s	6	$\log n + O(\log\log n)$	a	9

Table 1: Overview of our results. Recall that n denotes the number of requests in the input sequence. We mark the ratios that are strict by “s” and the ones that are asymptotic by “a”. Note that a strict lower bound can be larger than an asymptotic upper bound. For each bound, we indicate the number of the theorem proving the result. For readability, many of the bounds stated are weaker than those proven in the paper. Moreover, the upper bounds for the path hold for any bipartite graph. The result of Theorem 3 in the third row of the table is valid only for *neighborhood-based* algorithms, as defined just before Theorem 3 in Section 2.

2 The Path

As explained earlier, we establish all lower bounds for paths, and since a path is bipartite, all these negative results carry over to bipartite graphs. Similarly, all our (constructive) upper bounds are given for bipartite graphs and therefore also apply to paths. We start with three lower bound results.

2.1 Lower bounds

Theorem 1 Any strictly 1-competitive online algorithm for multi-coloring paths of at least 10 nodes has advice complexity at least $\left\lceil\log(\left\lfloor\frac{n}{4}\right\rfloor+1)\right\rceil$.

Proof We let $m = \left\lfloor\frac{n}{4}\right\rfloor$ and define a set S of $m+1$ sequences, all having the same prefix of length $2m$. The set S will have the following property: for no two sequences in S can their prefixes be colored in the same way while ending up using the optimal number of colors on the complete sequence. Starting from one end of the path, we denote the nodes $v_1, v_2, \ldots$.

We define the set S to consist of the sequences $I_0, I_1, \ldots, I_m$, where I_i is defined in the following way. First m requests are given to each of the nodes v_1 and v_4. Then

i requests to each of v_2 and v_3. To give all sequences the same length, the sequence ends with $\lceil n-2m-2i \rceil$ requests distributed as evenly as possible among v_6, v_8, and v_{10}. Since $\lceil \lceil n-2m-2i \rceil /3 \rceil \leq m$, the optimal number of colors will not be influenced by this part of the sequence.

Note that $\textsc{Opt}(I_i) = m+i$. In order not to use more than $\textsc{Opt}(I_i)$ colors for I_i, exactly i of the colors assigned to v_4 have to be different from the colors assigned to v_1. The prefixes of length $2m$ in S are identical, so all information to distinguish between the different sequences must be given as advice. The cardinality of S is $m+1$. To specify one out of $m+1$ possible actions, $\lceil \log(m+1) \rceil$ bits are necessary. $\square$

For algorithms that are $\frac{9}{8}$-competitive or better, we give the following lower bound.

Theorem 2 Consider multi-coloring paths of at least 10 nodes. For any $b \geq 3$ and any $(1+\frac{1}{2^b})$-competitive algorithm, A, there exists an $N \in \mathbb{N}$ such that A has advice complexity at least $b-2$ on sequences of length at least N.

Proof For any $(1+\frac{1}{2^b})$-competitive algorithm, A, there exists an $\alpha \geq 1$ such that $\mathrm{A}(I) \leq (1+\frac{1}{2^b})\,\textsc{Opt}(I)+\alpha$, for any input sequence I. We consider sequences of length $n \geq 2^{2b+2}\alpha + 3$.

Let $m = \left\lfloor \frac{n}{4} \right\rfloor$ and consider the same set of sequences as in the proof of Theorem 1. Recall that

$$\textsc{Opt}(I_i) = m+i\,.$$

For the sequence I_i, let x_i denote the number of colors that A uses on v_4 but not on v_1. Then, A uses $m+x_i$ colors in total for v_1 and v_4. On v_3, it can use at most x_i of the colors used at v_1, so the total number of colors used at v_1, v_2, and v_3 is at least $m+2i-x_i$. Thus,

$$\mathrm{A}(I_i) \geq \max\left\{m+x_i, m+2i-x_i\right\}\,.$$

We will prove that there are $p \geq 2^{b-2}$ sequences $I_{i_1}, I_{i_2}, \ldots, I_{i_p}$ such that, for any pair $i_j \neq i_k$, we have $x_{i_j} \neq x_{j_k}$, or otherwise A would not be $(1+\frac{1}{2^b})$-competitive on sequences of at least 2^{b+2} requests. This will immediately imply that A must use at least $b-2$ advice bits.

Let $\varepsilon = \frac{1}{2^b} + \frac{1}{2^{2b}}$. From $\mathrm{A}(I_i) \leq (1+\frac{1}{2^b})\,\textsc{Opt}(I_i)+\alpha$ and $m \geq 2^{2b}\alpha$, we obtain the inequalities

$$m+x_i \leq (1+\varepsilon)(m+i)$$

and

$$m+2i-x_i \leq (1+\varepsilon)(m+i)$$

which reduce to

$$x_i \leq \varepsilon m + (1+\varepsilon)i \tag{1}$$

and

$$i \leq \frac{x_i + \varepsilon m}{1-\varepsilon} \tag{2}$$

Hence, by (1), $x_0 \leq \varepsilon m$. Thus, by (2), we can have $x_i = x_0$, only if $i \leq \frac{2\varepsilon m}{1-\varepsilon}$. Therefore, we let $i_1 = 0$ and $i_2 = \lfloor \frac{2\varepsilon m}{1-\varepsilon} + 1 \rfloor$. In general, we ensure $x_{i_j} \neq x_{i_{j+1}}$ by letting $i_{j+1} = \lfloor \frac{x_{i_j}+\varepsilon m}{1-\varepsilon} + 1 \rfloor$. Thus,

$$\begin{aligned} i_{j+1} &\leq \frac{x_{i_j} + \varepsilon m}{1-\varepsilon} + 1 \\ &\leq \frac{\varepsilon m + (1+\varepsilon)i_j + \varepsilon m}{1-\varepsilon} + 1, \text{ by } (1) \\ &= \frac{1+\varepsilon}{1-\varepsilon} \cdot i_j + \frac{2\varepsilon m}{1-\varepsilon} + 1 \end{aligned}$$

Solving this recurrence relation, we get

$$\begin{aligned} i_{j+1} &\leq \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^j \cdot i_1 + \sum_{k=0}^{j-1} \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^k \left(\frac{2\varepsilon m}{1-\varepsilon} + 1\right) \\ &= \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^j \cdot 0 + \frac{\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^j - 1}{\frac{1+\varepsilon}{1-\varepsilon} - 1} \left(\frac{2\varepsilon m}{1-\varepsilon} + 1\right) \\ &= \frac{\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^j - 1}{1+\varepsilon-1+\varepsilon} (2\varepsilon m + 1 - \varepsilon) \\ &= \frac{\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^j - 1}{2\varepsilon} (2\varepsilon m + 1 - \varepsilon) \end{aligned}$$

We let p equal the largest j for which $i_j \leq m$:

$$
\begin{aligned}
& m < i_{p+1} \leq \frac{\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^p - 1}{2\varepsilon}\,(2\varepsilon m + 1 - \varepsilon) \\
\Rightarrow\; & 2\varepsilon m < \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^p (2\varepsilon m + 1 - \varepsilon) - (2\varepsilon m + 1 - \varepsilon) \\
\Leftrightarrow\; & \frac{4m\varepsilon + 1 - \varepsilon}{2m\varepsilon + 1 - \varepsilon} < \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^p \\
\Leftrightarrow\; & \ln\left(2 - \frac{1-\varepsilon}{2m\varepsilon + 1 - \varepsilon}\right) < p \cdot \ln\left(1 + \frac{2\varepsilon}{1-\varepsilon}\right) \\
\Rightarrow\; & \ln\left(2 - \frac{1-\varepsilon}{2m\varepsilon + 1 - \varepsilon}\right) < p \cdot \frac{2\varepsilon}{1-\varepsilon}, \text{ since } \ln(1+x) \leq x, \text{ for } x > -1 \\
\Rightarrow\; & \ln\left(2 - \frac{1}{3}\right) < p \cdot \frac{2\varepsilon}{1-\varepsilon}, \text{ since } m\varepsilon > 2^b > 1 - \varepsilon, \\
\Rightarrow\; & \ln\left(\sqrt{e}\right) < p \cdot \frac{2\varepsilon}{1-\varepsilon} \\
\Leftrightarrow\; & \frac{1}{2} < p \cdot \frac{2\varepsilon}{1-\varepsilon} \\
\Leftrightarrow\; & p > \frac{1-\varepsilon}{4\varepsilon} \\
\Leftrightarrow\; & p > \frac{1 - \frac{1}{2^b} - \frac{1}{2^{2b}}}{\frac{4}{2^b} + \frac{4}{2^{2b}}} = \frac{2^{2b} - 2^b - 1}{2^{b+2} + 4} > 2^{b-2} - 1 \\
\Rightarrow\; & p \geq 2^{b-2}, \text{ since } p \text{ is an integer}
\end{aligned}
$$

This completes the proof. □

For the following theorem, we define the class of *neighborhood-based* algorithms: A multi-coloring algorithm, A, is called neighborhood-based, if there exists a constant d such that, when assigning a color to a request to a node v, A bases its decision only on requests to nodes a distance of at most d away from v. Note that, in particular, a neighborhood-based algorithm cannot base its decision on the current value of OPT.

Theorem 3 No neighborhood-based online algorithm for multi-coloring paths with advice complexity $O(1)$ can be better than $\frac{4}{3}$-competitive.

Proof Having an online algorithm with advice complexity $O(1)$ gives an algorithm a constant number of possible algorithmic behaviors; it is equivalent to having $O(1)$ online algorithms without advice and choosing one of these according to the given advice.

As shown in [15], the family of sequences used in the proofs of Theorems 1 and 2, can be used to prove that any online algorithm without advice has a competitive ratio of at least $\frac{4}{3}$. The result is asymptotic, since the construction works with any scaling of the number of requests to each node. This means that for each algorithm, there exists an infinite family of sequences indexed by n, the length of the sequences, establishing the lower bound for each algorithm.

For any neighborhood-based algorithm A, there is a constant d such that, when assigning a color to a request, A ignores all requests given to nodes a distance of more than d away from the requested node. For any n and any family, there is a smallest and a largest node on the path which is requested, and the part of the path from this smallest to the largest node defines a subpath. We now rename nodes in these infinite families so that the subpaths used by the different families are separated by d unused nodes. We then form one request sequence by concatenating all these renamed subsequences. We scale the number of requests in each sequence such that the value of OPT is the same for each sequence.

Clearly, no matter which of the $O(1)$ algorithms are run on this constructed family, its performance tends to at least $\frac{4}{3}$ OPT. $\square$

2.2 Upper bounds

For multi-coloring of a path, there exists a strictly 1-competitive 1-recoloring algorithm, GREEDYOPT [15]. GREEDYOPT divides the nodes into two sets, *upper* and *lower*, such that every second node belongs to *upper* and the remaining nodes belong to *lower*. The following invariant is maintained: After each request, each node in *lower* uses consecutive colors starting with the color 1 and each node in *upper* uses consecutive colors ending with a color no larger than the optimal number of colors for the sequence of requests seen so far.

The algorithm for paths from [15] is easily generalized to work on bipartite graphs, letting the nodes of one partition, L, belong to *lower* and the nodes of the other partition, U, belong to *upper*. Recoloring is only needed if the number of colors used by an optimal offline algorithm is not known. Hence, using *enc*(OPT) advice bits, an online algorithm can be strictly 1-competitive, even if recoloring is not allowed. We call the resulting algorithm GREEDYOPTADVICE.

To describe the algorithm GREEDYOPTADVICE in detail, we need some notation: Let $f_i(v)$ denote the set of colors assigned to node v after the first i requests, starting with request 1. Also, for notational convenience, we define $f_0(v) = \emptyset$ for all v. To smoothly handle initially empty sets of colors in the algorithm, we define that if $f_i(v)$ is the empty set, then $\min f_i(v) = \max f_i(v) = 0$. This notation will be used throughout the appendix. GREEDYOPTADVICE is listed as Algorithm 1.

Algorithm 1 The multi-coloring algorithm GREEDYOPTADVICE.

```
1:  Assume that a bipartite graph is given by the partition into $L$ and $U$.
2:  **Advice:** $m = \text{OPT}$
3:  **for** $i = 1$ **to** $n$ **do**
4:      Assume that the $i$th request, $r$, is to node $v$
5:      **if** $v \in U$ **then**
6:          /* using the upper colors top-down */
7:          **if** $f_{i-1}(v) = \emptyset$ **then**
8:              give $r$ color $m$
9:          **else**
10:             give $r$ color $\min f_{i-1}(v) - 1$
11:     **else**
12:         /* $v \in L$; using the lower colors bottom-up */
13:         give $r$ color $\max f_{i-1}(v) + 1$
```

Theorem 4 Algorithm GREEDYOPTADVICE is correct, strictly 1-competitive, and has advice complexity *enc*(OPT).

Proof We consider correctness first. Clearly, at time i, the maximum color assigned to a node $v \in L$ is $\max f_i(v) = |f_i(v)|$ and the minimum color assigned to a node $v \in U$ is $\min f_i(v) = \text{OPT} + 1 - |f_i(v)|$ (assuming v has received at least one request).

Assume for the sake of contradiction that, at some time i, a request to a node $l \in L$ gets assigned the same color c as a request to a neighboring node u, which must belong to U. This means that $c = |f_i(l)|$ and $c = \text{OPT} + 1 - |f_i(u)|$, and, as l and u are neighbors, $\text{OPT} \geq |f_i(l)| + |f_i(u)|$, but then $\text{OPT} \geq |f_i(l)| + |f_i(u)| = c + \text{OPT} + 1 - c = \text{OPT} + 1$. This is a contradiction, so GREEDYOPTADVICE is correct.

It follows directly that the maximum color that GREEDYOPTADVICE assigns is OPT, implying that GREEDYOPTADVICE is strictly 1-competitive.

The maximum color that an optimal offline algorithm uses, given a sequence of length n, is n. Therefore, $\text{OPT} \leq n$ and *enc*(OPT) advice bits are sufficient. □

We turn to nonoptimal variants of GREEDYOPTADVICE using fewer than *enc*(OPT) advice bits. We show how to obtain a particular competitive ratio of $1 + \frac{1}{2^b}$, using $b + 1 + O(\log \log \text{OPT})$ bits of advice. Thus, essentially, we are approaching optimality exponentially fast in the number of bits of advice.

Theorem 5 For any integer $b \geq 1$, there exists a strictly $(1 + \frac{1}{2^{b-1}})$-competitive online algorithm for multi-coloring bipartite graphs with advice complexity $b + \mathit{enc}(a)$, where $a + b$ is the total number of bits in the value $\textsc{Opt}$.

Proof As advice, the algorithm asks for the b high order bits of the value $\textsc{Opt}$, as well as the number $a = \lceil \log(\textsc{Opt} + 1) \rceil - b$ of low order bits, but not the value of these bits. The algorithm knows b and can therefore just read the first b bits, while a needs to be encoded. Thus, $b + \mathit{enc}(a)$ bits are sufficient to encode the advice.

First, if $\textsc{Opt}$ contains fewer than b bits, this is detected by a being zero. In this case, some of the b bits may be leading zeros. By Theorem 4, we can then be strictly 1-competitive.

Now, assume this is not the case. Let $\textsc{Opt}_b = \left\lfloor \frac{\textsc{Opt}}{2^a} \right\rfloor$ denote the value represented by the b high order bits. Then the algorithm computes $m = 2^a \,\textsc{Opt}_b + 2^a - 1$ and runs $\textsc{GreedyOptAdvice}$ with this m. Since $\textsc{Opt} \leq m \leq \textsc{Opt} + 2^a - 1$, the algorithm is correct and uses at most $\textsc{Opt} + 2^a - 1$ colors.

For any number $x \geq 1$, consisting of c bits, with the most significant bit being one, $2^c \leq 2x$. Thus, $2^{b+a} \leq 2\,\textsc{Opt}$, so $2^a \leq \frac{2\,\textsc{Opt}}{2^b}$. This means that the number of colors used by $\textsc{GreedyOptAdvice}$ is less than $\textsc{Opt} + \frac{2\,\textsc{Opt}}{2^b} = (1 + \frac{1}{2^{b-1}})\,\textsc{Opt}$, so the algorithm is strictly $(1 + \frac{1}{2^{b-1}})$-competitive. □

Considering the lower bound of Theorem 1 versus the upper bound of Theorem 4, as well as the lower bound of Theorem 2 versus the upper bound of Theorem 5, in both cases there is a small discrepancy of a few bits, in addition to a low order term. The lower bound proof of Theorem 1 demonstrates the need of advice to distinguish between $\left\lfloor \frac{n}{4} \right\rfloor + 1$ different scenarios to be optimal. It will vary with n whether or not the division by four saves one or two bits compared with $\log n$, and similar reasoning applies to Theorem 2. Thus, when stating the lower bound, we have to subtract two bits (refer to Table 1). Using encoding tricks, to for instance identify cases where $\textsc{Opt}$ has a very small value, we can also sometimes get down to a bit less than $\log n$ for the upper bound. Thus, our results are nearly tight, up to low order terms, but because of rounding, it seems difficult to squeeze the missing few bits out of the bounds in every case. Note that for upper bounds, one could perform better by distinguishing between different cases, but finding out which case to use requires extra bits, by which we lose the advantage again.

Corollary 1 For any $\varepsilon > 0$, there exists a strictly $(1+\varepsilon)$-competitive deterministic online algorithm for multi-coloring bipartite graphs with advice complexity

$$O(\log\log \textsc{Opt}).$$

Proof Except for the term b, the advice stated in Theorem 5 is $O(\log\log \textsc{Opt})$ and $\textsc{Opt} \leq n$. Thus, we just need to bound the term b. For a given ε, choose b large enough such that $\frac{1}{2^{b-1}} \leq \varepsilon$. Using this value for b in Theorem 5, we obtain an algorithm with a strict competitive ratio of at most $1 + \frac{1}{2^{b-1}} \leq 1 + \varepsilon$. Since, for any given ε, b is a constant, the total amount of advice is $O(\log\log \textsc{Opt})$. □

2.3 Cancellations

Algorithm 2 The multi-coloring algorithm GreedyOptAdviceCancel.

```
1: Assume that a bipartite graph is given by the partition into $L$ and $U$.
2: Advice: $m = \textsc{Opt}$
3: for $i = 1$ to $n$ do
4:     Assume that the $i$th request, $r$, is to node $v$
5:     if $r$ is a color request then
6:         if $v \in U$ then
7:             /* using the upper colors top-down */
8:             if $f_{i-1}(v) = \emptyset$ then
9:                 give $r$ color $m$
10:            else
11:                give $r$ color $\min f_{i-1}(v) - 1$
12:        else
13:            /* $v \in L$; using the lower colors bottom-up */
14:            give $r$ color $\max f_{i-1}(v) + 1$
15:    else
16:        /* $r$ is a cancellation */
17:        if $v \in U$ then
18:            if the color of $r$ is different from $\min f_{i-1}(v)$ then
19:                recolor the request that has color $\min f_{i-1}(v)$, giving it the
                   color of $r$
20:        else
21:            if the color of $r$ is different from $\max f_{i-1}(v)$ then
22:                recolor the request that has color $\max f_{i-1}(v)$, giving it
                   the color of $r$
```

The Multi-Coloring problem is sometimes considered in the context of request cancellations, i.e., a color already given to a node disappears again. We observe that even using the weakest form of recoloring, namely 0-recoloring, where only requests at the node where the cancellation takes place may be recolored, we can extend the algorithm GREEDYOPTADVICE, using the same advice, to a strictly 1-competitive algorithm. This is simply done by recoloring at most one request per cancellation to ensure that the invariants regarding lower and upper nodes are maintained, i.e., ensuring that the colors used at any node form a consecutive sequence starting from one and increasing and starting from OPT and decreasing for lower and upper nodes, respectively. This algorithm, GREEDYOPTADVICECANCEL, is listed as Algorithm 2. Note that the difference to Algorithm 1 is the check in line 5 as to whether the current request is a color request and the addition of lines 15–22 handling cancellations.

3 Hexagonal Graphs

A hexagonal graph is a graph that can be obtained by placing (at most) one node in each cell of a hexagonal grid (such as the one sketched in Figure 1) and adding an edge between any pair of nodes placed in neighboring cells. Note that any hexagonal graph can be 3-colored. This is easily seen, since it is possible to use the three colors cyclically on the cells of each row of the underlying hexagonal grid, such that no two neighboring cells receive the same color.

3.1 Lower bounds

Theorem 6 Any online algorithm for multi-coloring hexagonal graphs with a strict competitive ratio strictly smaller than $\frac{3}{2}$ has advice complexity at least $\left\lfloor \frac{n-1}{3} \right\rfloor$.

Proof First, we explain a small part of the construction that we will use in many copies. We consider two sequences with the same prefix of length 2. Both sequences can be colored with two colors, but this requires coloring the two prefixes of length two differently. Consider the left-most part of Figure 1 (surrounded by thick lines) consisting of the "double" nodes D_1 and D_2, the "outer" nodes O_0 and O_1 and the "single" nodes S_1, and S_2. These nodes form the same type of configuration as the nodes D_3, D_4, O_1, O_2, S_3, and S_4. If a pair of outer nodes are given some requests, they can later be "connected" by follow-up requests to either the two double nodes or the single node between them.

First the nodes O_0 and O_1 get one request each. Then, either D_1 and D_2 or S_1 and S_2 receive one request each. The node S_2 is used to get up to the same sequence

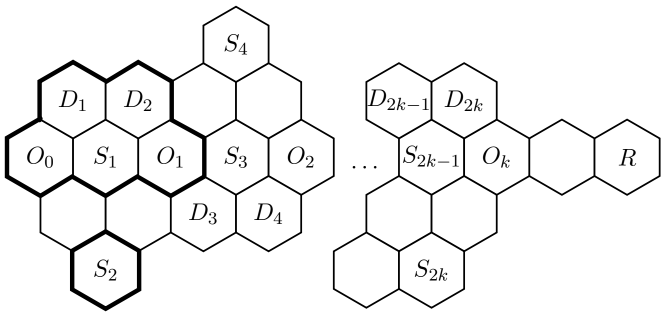

Figure 1: Hexagonal lower bound construction.

length in all cases. In order not to use more than two colors, the outer nodes have to use different colors if we later give requests to the two D-nodes. Similarly, the O-nodes should have the same color if we later give a request to the S-node in between them. Since the prefix of length two is $\langle O_0, O_1 \rangle$ for both sequences, all information for an algorithm to distinguish between the two sequences must be given as advice.

We can repeat this graph pattern $\left\lfloor \frac{n-1}{3} \right\rfloor$ times, as illustrated in Figure 1 with $k = \left\lfloor \frac{n-1}{3} \right\rfloor$, giving the requests to all O-nodes first.

We now define the set of sequences S of cardinality $2^{\left\lfloor \frac{n-1}{3} \right\rfloor}$ formally, i.e., we define a sequence for each possible combination of requests to either D_{2j-1} and D_{2j} or S_{2j-1} and S_{2j} for $j = 1, 2, \dots, \left\lfloor \frac{n-1}{3} \right\rfloor$. A sequence is defined for any chosen combination of the following i-values, i.e., by choosing a tuple $(i_1, i_2, \dots, i_{\left\lfloor \frac{n-1}{3} \right\rfloor}) \in \{0,1\}^{\left\lfloor \frac{n-1}{3} \right\rfloor}$. For any such choice, we define the sequence as a concatenation of the subsequences given below. In the description of the subsequences, we use the notation $\mathrm{Req}(v, m)$ to denote a sequence of m requests to a node v, and also use this notation for $m = 0$, denoting the empty request sequence, and $m = 1$, denoting one request.

- $\mathrm{Req}(O_j, 1)$ for $j = 0, 1, \dots, \left\lfloor \frac{n-1}{3} \right\rfloor$
- $\mathrm{Req}(D_{2j-1}, i_j)$, $\mathrm{Req}(D_{2j}, i_j)$ for $j = 1, 2, \dots, \left\lfloor \frac{n-1}{3} \right\rfloor$
- $\mathrm{Req}(S_{2j-1}, 1 - i_j)$, $\mathrm{Req}(S_{2j}, 1 - i_j)$ for $j = 1, 2, \dots, \left\lfloor \frac{n-1}{3} \right\rfloor$
- $\mathrm{Req}(R, n - (3 \left\lfloor \frac{n-1}{3} \right\rfloor + 1))$

Note that for any i_j, $\{i_j, 1 - i_j\} = \{0, 1\}$ and we either give requests to the D-

nodes or the S-nodes. The possible requests to R simply gets all sequences up to a length of n.

The node O_0 is given some color. After that, we have $\left\lfloor \frac{n-1}{3} \right\rfloor$ independent choices of coloring each node O_i in the prefix of any sequence identically to O_{i-1} or not. Since the prefixes are the same, all information for an algorithm to distinguish between the different sequences must be given as advice. To specify one out of $2^{\left\lfloor \frac{n-1}{3} \right\rfloor}$ possible actions, $\left\lceil \log 2^{\left\lfloor \frac{n-1}{3} \right\rfloor} \right\rceil = \left\lfloor \frac{n-1}{3} \right\rfloor$ bits are necessary. □

Theorem 7 Any online algorithm for multi-coloring hexagonal graphs with competitive ratio strictly smaller than $\frac{5}{4}$ has advice complexity $\Omega(n)$.

Proof We use the basic construction from Theorem 6. Assume p requests are given to one of the components like this:

First, we give $\frac{p}{4}$ requests to each of O_0 and O_1. Let q, $0 \leq q \leq \frac{p}{4}$, denote the number of colors used at both nodes. Then following up by giving $\frac{p}{4}$ requests to each S-node results in a minimum of $\frac{3p}{4} - q$ colors used, while giving the requests to the D-nodes instead results in a minimum of $\frac{p}{2} + q$ colors.

Note that OPT $= \frac{p}{2}$, independent of in which of the two ways the sequence is continued. Thus, for any $\varepsilon > 0$, any $(\frac{5}{4} - \varepsilon)$-competitive algorithm must choose q such that, for some constant α, $\frac{3p}{4} - q \leq \left(\frac{5}{4} - \varepsilon\right) \frac{p}{2} + \alpha$ and $\frac{p}{2} + q \leq \left(\frac{5}{4} - \varepsilon\right) \frac{p}{2} + \alpha$. Adding these two inequalities, we obtain $\frac{5p}{4} \leq (\frac{5}{4} - \varepsilon)p + 2\alpha$ which is equivalent to $\varepsilon p \leq 2\alpha$. Thus, if p is non-constant, no $(\frac{5}{4} - \varepsilon)$-competitive algorithm can use the same value of q for both sequences.

Now assume for the sake of contradiction that for some advice of $g(n) \in o(n)$ bits, we can obtain a ratio of $\frac{5}{4} - \varepsilon$. Let $f(n) = \frac{1}{2}\frac{n}{g(n)}$. Since $g(n) \in o(n)$, $f(n) \in \omega(1)$. The idea is now to repeat the construction as in the proof of Theorem 6 and give $f(n)$ requests to each construction ($f(n)$ has the role of p in the above). Since a pair of neighboring constructions share $f(n)/4$ requests, this results in $\frac{n - f(n)/4}{3f(n)/4} = \frac{4n - f(n)}{3f(n)} \geq \frac{n}{f(n)}$ constructions. We assume without loss of generality that all our divisions result in integers.

In order to be $(\frac{5}{4} - \varepsilon)$-competitive, an online algorithm must, for each two neighboring O-nodes, choose between at least two different values of q. These are independent decisions, and the ratio only ends up strictly better than $\frac{5}{4}$ if the algorithm decides correctly in every subconstruction. Thus, it needs at least $\frac{n}{f(n)}$ bits of advice. However, $\frac{n}{f(n)} = \frac{n}{\frac{1}{2}\frac{n}{g(n)}} = 2g(n) > g(n)$, which is a contradiction. □

3.2 Upper bounds

We have the following trivial upper bound on the advice necessary to be optimal, independent of the graph topology:

Theorem 8 There is a strictly 1-competitive online multi-coloring algorithm with advice complexity $(n+1)\lceil \log \textsc{Opt} \rceil$.

Proof Start by asking for the number of bits necessary to represent values up to $\textsc{Opt}$. Then for each request, read $\lceil \log(\textsc{Opt}+1) \rceil$ bits telling, which color to use. This gives *enc*$(\lceil \log \textsc{Opt} \rceil) + n \lceil \log \textsc{Opt} \rceil < (n+1) \lceil \log \textsc{Opt} \rceil$. □

In the following, we will show how two known approximation algorithms can be converted to online algorithms with advice. In the description of the algorithms, we let the *weight* of a clique denote the total number of requests to the nodes of the clique. Note that the only maximal cliques in a hexagonal graph are isolated nodes, edges, or triangles. We let ω denote the maximum weight of any clique in the graph.[1]

A $\frac{3}{2}$-competitive algorithm called the Fixed Preference Allocation algorithm, FPA, was proposed in [24]. In [31], the strategy was simplified and it was noted that the algorithm can be converted to a 1-recoloring online algorithm. We describe the simplified offline algorithm below.

The algorithm uses three color classes, R, G, and B. The color classes represent a partitioning of the nodes in the graph so that no two neighbors are in the same partition. Each of the three color classes has its own set of $\lceil \frac{\omega}{2} \rceil$ colors, and each node in a given color class uses the colors of its color class, starting with the smallest. This set of colors is also referred to as the node's *private* colors. If more than $\lceil \frac{\omega}{2} \rceil$ requests are given to a node, then it borrows colors from the private colors of one of its neighbors, taking the highest available color. R nodes can borrow colors from G nodes, G from B, and B from R.

For completeness, we give the arguments that FPA is correct and obtains an approximation ratio of $\frac{3}{2}$. Assume for the purpose of contradiction that the coloring produced by the algorithm causes a conflict between an R node and a G node. This means that their combined number of requests must be greater than ω, which is a contradiction. The same argument holds for the other color combinations. Thus, the coloring is legal. Any optimal algorithm needs at least ω colors, so $\textsc{Opt} \geq \omega$ and the algorithm is a $\frac{3}{2}$-approximation algorithm.

[1]The Greek letter ω is traditionally used here, so we will also do that. Since there is no argument, this should not give rise to confusion with the $\omega(f)$, stemming from asymptotic notation.

Algorithm 3 The $\frac{3}{2}$-competitive algorithm, FPA, with advice.

```
1:  **Advice:** $\lceil \frac{\omega}{2} \rceil$
2:  $\textsc{Red} = \{1, 2, \ldots, \lceil \frac{\omega}{2} \rceil\}$,
3:  $\textsc{Green} = \{\lceil \frac{\omega}{2} \rceil + 1, \lceil \frac{\omega}{2} \rceil + 2, \ldots, 2\lceil \frac{\omega}{2} \rceil\}$,
4:  $\textsc{Blue} = \{2\lceil \frac{\omega}{2} \rceil + 1, 2\lceil \frac{\omega}{2} \rceil + 2, \ldots, 3\lceil \frac{\omega}{2} \rceil\}$
5:  **Function** $\mathrm{Class}(v)$
6:      **return** $v$'s color class: R, G, or B
7:  **Function** $\mathrm{Borrow}(c)$
8:      **return** the next class in the wrap-around sequence R, G, or B
9:  **Function** $\mathrm{Colors}(c)$
10:     **return** the set of private colors of class $c$
11: **for** $i = 1$ **to** $n$ **do**
12:     Assume that the $i$th request, $r$, is to node $v$
13:     **if** $|f_{i-1}(v)| < \lceil \frac{\omega}{2} \rceil$ **then**
14:         give $r$ color $\min(\mathrm{Colors}(\mathrm{Class}(v)) \setminus f_{i-1}(v))$
15:     **else**
16:         give $r$ color $\max(\mathrm{Colors}(\mathrm{Borrow}(\mathrm{Class}(v))) \setminus f_{i-1}(v))$
```

Since $\lceil \frac{\omega}{2} \rceil \leq \lceil \frac{\textsc{Opt}}{2} \rceil$, we can give $\lceil \frac{\omega}{2} \rceil$ as advice, resulting in Algorithm 3. Note that the f-notation used in the pseudo-code was defined in connection with Algorithm 1.

Theorem 9 There is a $\frac{3}{2}$-competitive online algorithm for multi-coloring hexagonal graphs with advice complexity *enc*$(\lceil \frac{\textsc{Opt}}{2} \rceil)$.

Proof Given $\lceil \frac{\omega}{2} \rceil \leq \lceil \frac{\textsc{Opt}}{2} \rceil$ as advice, FPA can be used as an online algorithm (Algorithm 3). □

In [30], an algorithm with an improved approximation ratio of $\frac{4}{3}$ was introduced. We now describe this algorithm. For completeness, we also give the arguments that the algorithm is correct and is a $\frac{4}{3}$-approximation algorithm:

The algorithm uses color classes in the same way as FPA, except that the private color sets contain only $\lfloor \frac{\omega+1}{3} \rfloor$ colors each. We use the following notation. For any node v, we let n_v denote the number of requests to v. Furthermore, b_v denotes the maximum number of colors that v can borrow, i.e., $b_v = \max\{0, \lfloor \frac{\omega+1}{3} \rfloor - n'_v\}$, where n'_v is the maximum number of requests to any of the neighboring nodes in the color class that v can borrow from.

The algorithm can be seen as working in up to three phases:

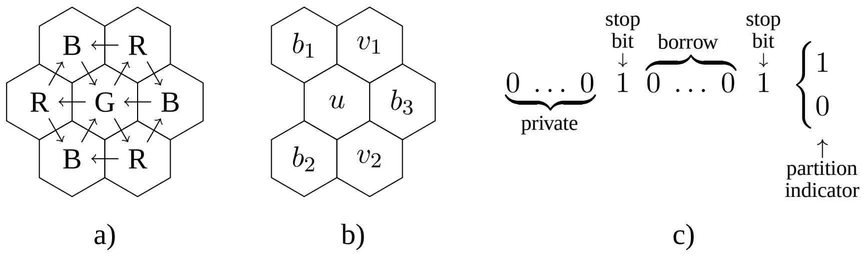
B
R
R
G
B
B
R
a)
b_1
v_1
u
b_3
b_2
v_2
b)
stop bit
borrow
stop bit
$\underbrace{0 \ldots 0}_{\text{private}}$ 1 $\overbrace{0 \ldots 0}^{\text{borrow}}$ 1
1
0
partition indicator
c)

Figure 2: Illustration of the $\frac{4}{3}$-approximation algorithm. a) The borrow pattern. Arrows show the direction of the flow of colors in Phase 2. b) Part of a graph induced by nodes still having unprocessed requests after Phase 2. c) The subsequence of advice bits connected to one node. The sequence of advice bits is a merge of such sequences.

In the *first phase*, the algorithm colors $\min\{n_v, \left\lfloor \frac{\omega+1}{3} \right\rfloor\}$ requests to each node, v, using the node's private colors. Let G_1 be the graph induced by the nodes that still have uncolored requests after Phase 1.

For any node, v, in G_1, $\left\lfloor \frac{\omega+1}{3} \right\rfloor$ requests to v are colored with v's private colors in Phase 1. By the definition of ω, this immediately implies that any pair of neighboring nodes have a total of at most $\omega - 2\left\lfloor \frac{\omega+1}{3} \right\rfloor$ uncolored requests already after Phase 1.

In the *second phase*, each node v with more than $\left\lfloor \frac{\omega+1}{3} \right\rfloor$ requests borrows $\min\{n_v - \left\lfloor \frac{\omega+1}{3} \right\rfloor, b_v\}$ colors. Let G_2 be the graph induced by nodes that still have uncolored requests after Phase 2.

In [30] it is proven that G_2 is bipartite and that any pair of neighbors in G_2 has a total of at most $\omega - 2\left\lfloor \frac{\omega+1}{3} \right\rfloor \leq \left\lfloor \frac{\omega+1}{3} \right\rfloor + 1$ uncolored requests after Phase 2. Thus, in the third phase, the remaining requests can be colored with GREEDYOPT (see the path section) using $\left\lfloor \frac{\omega+1}{3} \right\rfloor + 1$ additional colors.

To see that G_2 is bipartite, first note that G_1 (and hence G_2) cannot contain triangles. Each node in such a triangle would have received at least $\left\lfloor \frac{\omega+1}{3} \right\rfloor + 1$ requests, contradicting the definition of ω.

Using the fact that G_2 does not contain triangles, we can now argue that G_2 is acyclic and hence bipartite. Assume to the contrary that G_2 does contain a cycle, C. Assume without loss of generality that the R, G, B coloring of the underlying hexagonal grid is as shown in Figure 2 a) and let u be a leftmost node of C. Then, referring to Figure 2 b), two of the nodes v_1, v_2, and b_3 must also be part of C. Note that b_3 cannot be part of C, since then there would be a triangle after Phase 1.

Thus, u, v_1, and v_2 are part of the cycle and hence receive at least $\left\lfloor \frac{\omega+1}{3} \right\rfloor + 1$ requests each.

Since u could not borrow enough colors from the nodes in the color class it is allowed to borrow from, one of the b-nodes, say b_j, together with u must have a total of at least $2\left\lfloor \frac{\omega+1}{3} \right\rfloor + 1$ requests. So, b_j and u must form a triangle together with either v_1 or v_2 so that the three nodes together have received a total of at least $(2\left\lfloor \frac{\omega+1}{3} \right\rfloor + 1) + (\left\lfloor \frac{\omega+1}{3} \right\rfloor + 1)$ requests. This quantity is strictly larger than ω, contradicting the definition of ω.

This ends the argument that the algorithm is correct.

Since the total number of colors used is at most $3\left\lfloor \frac{\omega+1}{3} \right\rfloor + (\omega - 2\left\lfloor \frac{\omega+1}{3} \right\rfloor) \leq \frac{4\omega+1}{3}$, the algorithm is a $\frac{4}{3}$-approximation algorithm.

We now show how an online algorithm, given the right advice, can behave as the offline $\frac{4}{3}$-approximation algorithm. Note that the three phases of the offline $\frac{4}{3}$-approximation algorithm are characterized by the coloring strategy (using the node's own private colors, borrowing private colors from neighbors, or coloring a bipartite graph). However, when requests arrive online, the nodes may not go from one phase to the next simultaneously.

Theorem 10 There is a $\frac{4}{3}$-competitive online algorithm for multi-coloring hexagonal graphs with advice complexity at most $n + 2|V|$.

Proof We describe the algorithm and advice resulting in a coloring with at most $\frac{4}{3}$ OPT colors (see Algorithm 4, where we use the f-notation defined in connection with Algorithm 1).

Initially, each node is in Phase 1. On a request, the algorithm reads an advice bit and if it is zero, the next color from its private colors is used. If, instead, a one is read, this is treated as a stop bit for Phase 1, and this particular node enters Phase 2.

The algorithm starts with empty private color sets, and adds one color to each set whenever necessary, i.e., whenever a Phase 1 node that has already used all its private colors receives an additional request (this includes the first request to the node). As soon as a node leaves Phase 1, the algorithm knows that this node received $\left\lfloor \frac{\omega+1}{3} \right\rfloor$ requests, which is then the final size of each private color set. Knowing the size of the private color sets, the algorithm can calculate the maximum color for the complete coloring of the graph as $m = 4\left\lfloor \frac{\omega+1}{3} \right\rfloor + 1$.

In Phase 2, every zero indicates that the algorithm should borrow a color. When another stop bit is received (which could be after no zeros at all if the borrowing phase is empty), it moves to Phase 3. In Phase 3, it reads one bit to decide which partition, upper or lower, of the bipartite graph it is in, and does not need more information after that, since it simply uses the colors $3\left\lfloor\frac{\omega+1}{3}\right\rfloor+1,\ldots,m$, either top-down or bottom-up.

If we allow the algorithm one bit per request, it needs at most two more bits per node, since the stop bits are the only bits that do not immediately tell the algorithm which action to take. Thus, $n+2|V|$ bits of advice suffice. □

This algorithm can be used in many different ways, as long as the algorithm gets the information it needs. One other simple encoding would be to give the algorithm the value $\left\lfloor\frac{\omega+1}{3}\right\rfloor$ from the beginning and only give bit-wise advice after a node has used all its private colors. Since at least one color is private, this will save a total of at least $|V|$ bits, and result in at most *enc*$(\left\lfloor\frac{\omega+1}{3}\right\rfloor)+n+|V|$ bits of advice. This variant, and others, that are incomparable to each other, depending on the values of n, ω, and $|V|$, could all be used at the same time by first asking for a few bits to decide how to proceed. Thus, one could formulate a less readable but more accurate theorem basically taking the minimum of all the expressions. We have chosen clarity over precision, since the other expressions are mostly better in less interesting cases, where n is small compared to $|V|$, for instance.

3.3 Concluding Remarks

When considering advice complexity of multi-coloring on a path, we can achieve 1-competitiveness with a small amount of advice. A recoloring algorithm needs to be 1-recoloring to achieve the same. The advice is basically the maximum number of requests to any two neighboring nodes. Thus, whether one has that global information once and for all, or can obtain and adjust according to the local variant of this information gives the same result.

For multi-coloring of hexagonal graphs, there is a similar connection between recoloring distance and advice. The 1-recoloring online version of FPA has an advice variant and again, this advice represents information about the maximum number of requests to neighboring nodes. With additional global information about the bipartite induced subgraph, we can overcome the limitations of 1-recoloring algorithms and be as good as any known polynomial-time approximation algorithm.

References

[1] S. Albers, L.M. Favrholdt, and O. Giel. On paging with locality of reference. *Journal of Computer and System Sciences*, 70(2):145–175, 2005.

[2] K. Barhum, H.-J. Böckenhauer, M. Forisek, H. Gebauer, J. Hromkovič, S. Krug, J. Smula, and B. Steffen. On the power of advice and randomization for the disjoint path allocation problem. In *SOFSEM*, volume 8327 of *LNCS*, pages 89–101. Springer, 2014.

[3] M. Paola Bianchi, H.-J. Böckenhauer, J. Hromkovic, and L. Keller. Online coloring of bipartite graphs with and without advice. In *COCOON*, volume 7434 of *LNCS*, pages 519–530, 2012.

[4] H.-J. Böckenhauer, D. Komm, R. Královič, and R. Královič. On the advice complexity of the k-server problem. In *ICALP*, volume 6755 of *LNCS*, pages 207–218, 2011.

[5] H.-J. Böckenhauer, D. Komm, R. Královič, R. Královič, and T. Mömke. On the advice complexity of online problems. In *ISAAC*, volume 5878 of *LNCS*, pages 331–340, 2009.

[6] H.-J. Böckenhauer, D. Komm, R. Královič, and P. Rossmanith. On the advice complexity of the knapsack problem. In *LATIN*, volume 6139 of *LNCS*, pages 61–72, 2012.

[7] A. Borodin, S. Irani, P. Raghavan, and B. Schieber. Competitive paging with locality of reference. *Journal of Computer and System Sciences*, 50(2):244–258, 1995.

[8] J. Boyar, L.M. Favrholdt, K.S. Larsen, and M.N. Nielsen. Extending the Accommodating Function. *Acta Informatica*, 40(1):3–35, 2003.

[9] J. Boyar, S. Gupta, and K.S. Larsen. Access graphs results for LRU versus FIFO under relative worst order analysis. In *SWAT*, volume 7357 of *LNCS*, pages 328–339. Springer, 2012.

[10] J. Boyar, S. Kamali, K.S. Larsen, and A. López-Ortiz. Online bin packing with advice. In *STACS*, volume 25 of *LIPIcs*, pages 174–186. Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, 2014.

[11] J. Boyar and K.S. Larsen. The Seat Reservation Problem. *Algorithmica*, 25(4):403–417, 1999.

[12] J. Boyar, K.S. Larsen, and M.N. Nielsen. The Accommodating Function: a generalization of the competitive ratio. *SIAM Journal on Computing*, 31(1):233–258, 2001.

[13] J.W.-T Chan, F.Y.L. Chin, D. Ye, and Y. Zhang. Absolute and asymptotic bounds for online frequency allocation in cellular networks. *Algorithmica*, 58(2):498–515, 2010.

[14] J.W.-T. Chan, F.Y.L. Chin, D. Ye, Y. Zhang, and H. Zhu. Frequency allocation problems for linear cellular networks. In *ISAAC*, volume 4288 of *LNCS*, pages 61–70. Springer, 2006.

[15] M.G. Christ, L.M. Favrholdt, and K.S. Larsen. Online multi-coloring on the path revisited. *Acta Informatica*, 50(5–6):343–357, 2013.

[16] M. Chrobak, L. Jez, and J. Sgall. Better bounds for incremental frequency allocation in bipartite graphs. *Theoretical Computer Science*, 514:75–83, 2013.

[17] M. Chrobak and J. Sgall. Three results on frequency assignment in linear cellular networks. *Theoretical Computer Science*, 411(1):131–137, 2010.

[18] S. Dobrev, R. Královič, and D. Pardubská. Measuring the problem-relevant information in input. *RAIRO Theoretical Informatics and Applications*, 43(3):585–613, 2009.

[19] R. Dorrigiv, M. He, and N. Zeh. On the advice complexity of buffer management. In *ISAAC*, volume 7676 of *LNCS*, pages 136–145, 2012.

[20] P. Elias. Universal codeword sets and representations of the integers. *IEEE Transactions on Information Theory*, 21(2):194–203, 1975.

[21] Y. Emek, P. Fraigniaud, A. Korman, and A. Rosén. Online computation with advice. *Theoretical Computer Science*, 412(24):2642–2656, 2011.

[22] M. Forisek, L. Keller, and M. Steinová. Advice complexity of online coloring for paths. In *LATA*, volume 7183 of *LNCS*, pages 228–239, 2012.

[23] J. Hromkovič, R. Královič, and R. Královič. Information complexity of online problems. In *MFCS*, volume 6281 of *LNCS*, pages 24–36, 2010.

[24] J. Janssen, K. Kilakos, and O. Marcotte. Fixed preference channel assignment for cellular telephone systems. *IEEE Transactions on Vehicular Technology*, 48(2):533–541, 1999.

[25] J. Janssen, D. Krizanc, L. Narayanan, and S.M. Shende. Distributed on-line frequency assignment in cellular networks. *Journal of Algorithms*, 36(2):119–151, 2000.

[26] A.R. Karlin, M.S. Manasse, L. Rudolph, and D.D. Sleator. Competitive snoopy caching. *Algorithmica*, 3:79–119, 1988.

[27] D. Komm and R. Královič. Advice complexity and barely random algorithms. *RAIRO Theoretical Informatics and Applications*, 45(2):249–267, 2011.

[28] D. Komm, R. Královič, and T. Mömke. On the advice complexity of the set cover problem. In *CSR*, volume 7353 of *LNCS*, pages 241–252, 2012.

[29] M. P. Bianchi and H.-J. Böckenhauer and J. Hromkovič and S. Krug and B. Steffen. On the advice complexity of the online $l(2,1)$-coloring problem on paths and cycles. In *COCOON*, volume 7936 of *LNCS*, pages 53–64. Springer, 2013.

[30] C. McDiarmid and B.A. Reed. Channel assignment and weighted coloring. *Networks*, 36(2):114–117, 2000.

[31] L. Narayanan. *Channel Assignment and Graph Multicoloring*, pages 71–94. John Wiley & Sons, Inc., 2002.

[32] L. Narayanan and S.M. Shende. Static frequency assignment in cellular networks. *Algorithmica*, 29(3):396–409, 2001.

[33] L. Narayanan and S.M. Shende. Corrigendum: Static frequency assignment in cellular networks. *Algorithmica*, 32(4):679, 2002.

[34] S. Seibert, A. Sprock, and W. Unger. Advice complexity of the online coloring problem. In *CIAC*, volume 7878 of *LNCS*, pages 345–357, 2013.

[35] D.D. Sleator and R.E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.

[36] P. Sparl and J. Zerovnik. 2-local 4/3-competitive algorithm for multicoloring hexagonal graphs. *Journal of Algorithms*, 55(1):29–41, 2005.

[37] R. Witkowski and J. Zerovnik. 1-local 33/24-competitive algorithm for multicoloring hexagonal graphs. In *WAW*, volume 6732 of *LNCS*, pages 74–84, 2011.

Algorithm 4 Combining FPA and GREEDYOPTADVICE to a $\frac{4}{3}$-competitive algorithm.

```
1:  Advice: A sequence $B$ of bits classifying each request as to whether it should
    be colored using the node's own private colors, by borrowing, or in which
    partition it falls.
2:  Function $\mathrm{Class}(v)$
3:      return $v$'s color class: R, G, or B
4:  Function $\mathrm{Borrow}(c)$
5:      return the next class in the wrap-around sequence R, G, or B
6:  Function $\mathrm{Colors}(c)$
7:      return the set of private colors of class $c$
8:  Function NextBit($B$)
9:      return the next advice bit
10: for each node $v$ do
11:     Phase$(v) = 1$
12: for $i = 1$ to $n$ do
13:     Assume that the $i$th request, $r$, is to node $v$
14:     if Phase$(v) = 1$ then
15:         if NextBit$(B) = 0$ then
16:             if $\mathrm{Colors}(\mathrm{Class}(v)) \setminus f_{i-1}(v) = \emptyset$ then
17:                 add one color to each of the three sets of private colors
18:             give $r$ color $\min(\mathrm{Colors}(\mathrm{Class}(v)) \setminus f_{i-1}(v))$
19:         else
20:             $\mathrm{Phase}(v) = 2$
21:             $\mathrm{Phase3Min} = 3\,|f_{i-1}(v)| + 1$
22:             $\mathrm{Phase3Max} = 4\,|f_{i-1}(v)| + 1$
23:     if Phase$(v) = 2$ then
24:         if NextBit$(B) = 0$ then
25:             give $r$ color $\max(\mathrm{Colors}(\mathrm{Borrow}(\mathrm{Class}(v))) \setminus f_{i-1}(v))$
26:         else
27:             $\mathrm{Phase}(v) = 3$
28:             upper$_v$ = NextBit$(B)$ /* Store the partition of $v$ */
29:     if Phase$(v) = 3$ then
30:         /* Use GREEDYOPTADVICE: */
31:         if upper$_v = 1$ then
32:             give $r$ color $\max(\{\mathrm{Phase3Min}, \dots, \mathrm{Phase3Max}\} \setminus f_{i-1}(v))$
33:         else
34:             give $r$ color $\min(\{\mathrm{Phase3Min}, \dots, \mathrm{Phase3Max}\} \setminus f_{i-1}(v))$
```