On-line coloring of geometric intersection graphs

Thomas Erlebach*[1] Jiří Fiala†[2]

[1] ETH Zürich, Computer Engineering and Networks Lab,
CH-8092 Zürich, Switzerland.
E-mail: erlebach@tik.ee.ethz.ch
Corresponding author.

[2] Charles University, Faculty of Mathematics and Physics,
Institute for Theoretical Computer Science (ITI)†,
Malostranské nám. 2/25, 118 00, Prague, Czech Republic.
E-mail: fiala@kam.mff.cuni.cz

Abstract

This paper studies on-line coloring of geometric intersection graphs. It is shown that no deterministic on-line algorithm can achieve competitive ratio better than $\Omega(\log n)$ for disk graphs and for square graphs with n vertices, even if the geometric representation is given as part of the input. Furthermore, it is proved that the standard First-fit heuristic achieves competitive ratio $O(\log n)$ for disk graphs and for square graphs and is thus best possible.

Keywords: coloring, competitive analysis, intersection graphs, disk graphs, first-fit

1 Introduction

The *intersection graph* of a set of geometric objects is the graph with a vertex for each object and an edge between two vertices if and only if the

*Partially supported by EU Thematic Network APPOL II (IST-2001-32007), with funding provided by the Swiss Federal Office for Education and Science (BBW).

†This author was partially supported by EU APPOL project IST-1999-14084, during his postdoctoral fellowship at Christian-Albrechts University of Kiel.

†Supported by the Ministry of Education of the Czech Republic as project LN00A056.

corresponding objects intersect. A graph G is a *disk graph* if there exists a set of disks in the Euclidean plane whose intersection graph is G. Such a set of disks is then called a *disk representation* of G. The class of disk graphs has been studied for many years for its theoretical aspects as well as for its applications. As an example of a classical result we mention that Koebe proved in 1936 that every planar graph can be represented as a coin graph, i.e. a disk graph where disks are not allowed to overlap [13] (see also the more accessible discussion of Koebe's result by Sachs [19]). Concerning applications, Hale pointed out in 1980 that the frequency assignment problem can be modeled as a graph problem [9]. If we assume that all transmitters have circular range and transmitters with intersecting ranges are to use different frequencies, the underlying graph is a disk graph and the frequency assignment problem (without additional constraints) is equivalent to the graph coloring problem [7, 15]. Observe also that in this application the disk representation can be obtained from the placement of transmitters and their ranges.

The question of determining for a given graph whether it is a disk graph has been studied by several authors. In contrast to the case of planar graphs, no efficient algorithms are known for the recognition of disk graphs. Breu and Kirkpatrick have shown that the recognition problem is NP-hard for unit disk graphs (intersection graphs of disks with equal diameter) [2] and for disk graphs with bounded diameter ratio (intersection graphs of disks where the ratio of the largest diameter to the smallest diameter is bounded by an arbitrary constant) [1]. Hliněný and Kratochvíl proved NP-hardness for the recognition problem of arbitrary disk graphs [10].

The hardness of the recognition problem implies that a disk representation cannot be derived from the graph in polynomial time unless $P = NP$. Therefore, an important factor in the design of algorithms for disk graphs is whether the disk graph is given only as a set of edges and vertices, or whether the centers and diameters of the disks (i.e. the disk representation of the graph) form the input to the algorithm. Some problems can be solved efficiently no matter whether the disk representation is given or not. The problem of computing a maximum clique in a unit disk graph is an example: While the first polynomial-time algorithm for this problem, which was due to Clark, Colbourn, and Johnson [4], had required the disk representation, Raghavan and Spinrad recently presented an efficient algorithm that does not require the representation [18]. The situation seems to be different with respect to the approximability of the maximum independent set problem in disk graphs. Here, a polynomial-time approximation scheme has been found for the case of given disk representation [5], but only a 5-approximation algorithm is known for the case that the representation is not available [16].

The problem of coloring unit disks with three colors has been shown *NP*-complete in [4]. On the other hand, Peeters showed that the class of unit disk graphs admits a 3-approximate coloring algorithm [17], and this approach was generalized to disk graphs by Marathe et al. [16], who obtained a 5-approximation algorithm (see also Gräf [7] and Malesińska [15]). The proofs of these results also show that the number of colors required to color a unit disk graph or general disk graph with maximum clique size ω can be bounded by $3\omega - 2$ and $6\omega - 6$, respectively.

A class of graphs that are similar to disk graphs in certain aspects is the class of intersection graphs of squares (whose sides are parallel to the coordinate axes) in the plane, called *square graphs*. They have not been studied as intensively as disk graphs so far, except that it has often been noted that results for disk graphs can be adapted to intersection graphs of arbitrary regular polygons (including squares), e.g. [5, 11].

We focus our research on the on-line coloring problem for disk graphs and square graphs. The vertices of the graph are presented to the on-line coloring algorithm one by one. When a vertex v is revealed, all edges joining v to previously presented vertices are revealed as well. The algorithm must assign a color to v immediately without knowledge of future vertices and edges. The color assigned to v must be different from the colors of all previously colored vertices that are adjacent to v. We say that an on-line algorithm is *ρ-competitive* or achieves *competitive ratio ρ* if it always uses at most ρ times as many colors as an optimal coloring.

We distinguish on-line algorithms that are given the disk or square representation and on-line algorithms that are given only the intersection graph. In the former case, when a vertex is presented to the algorithm, the corresponding disk (given by the coordinates of the center and the diameter) or the corresponding square (given by the coordinates of the corners) is given to the algorithm as well. In the latter case, no geometric information is given to the algorithm.

As a motivation we shall mention the result of Gyárfás and Lehel from 1988 showing that there exists a tree T on n vertices such that for every on-line coloring algorithm there exists a specific ordering of the vertices of T, such that the algorithm is forced to use $\Omega(\log n)$ distinct colors [8]. Every tree is planar, so this result together with Koebe's theorem immediately shows that this lower bound is valid also for disk graphs *without given representation*. This leads naturally to the question of whether the

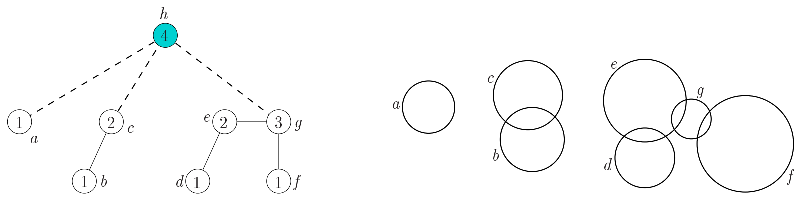

Figure 1: The disk representation constrains the adversary.

knowledge of the disk representation can help an on-line algorithm to get a better competitive ratio.

Note that in the setting with given representation, the choices of an adversary who wants to force the algorithm to use many colors are constrained by the geometry. To illustrate this, we consider an adversary who has already presented a forest consisting of several trees to the algorithm. See Figure 1 for an example. If the disk representation is not given to the algorithm (left-hand side of Figure 1), the adversary can next present a vertex that is adjacent to an arbitrary vertex from each of the trees, e.g. a new vertex h that is adjacent to a, c, and g, which forces the algorithm to use a fourth color. The resulting graph is a tree and, therefore, a disk graph. On the other hand, if the disk representation is given to the algorithm and if the situation is as shown on the right-hand side of Figure 1, the only possibilities for presenting a disk intersecting one disk from each of the trees are a disk intersecting a, c, and e or a disk intersecting a, b, and d. It is impossible to present a disk that intersects a, c, and g, but does not intersect e. If the adversary presents the vertices a to g with a different disk representation, the algorithm might assign different colors to the disks, again preventing the adversary from forcing a fourth color on a tree with 8 vertices. Furthermore, if the situation is as shown in the figure, the on-line algorithm *knows* that the adversary cannot present a disk intersecting only a, c, and g in the future.

Since the adversary is constrained by the geometry and since the algorithm is aware of these constraints, one might expect that the disk representation can help an on-line algorithm to get a better competitive ratio. The main result of this paper is that this is not the case. We prove that for every on-line disk coloring algorithm there exists a sequence of n disks such that the algorithm is forced to use $\Omega(\log n)$ distinct colors, while an optimal coloring uses only two colors. We also adapt a result of Irani [12] and show that a competitive ratio of $O(\log n)$ is achieved by the First-fit coloring algorithm. This shows that the First-fit algorithm is optimal for on-line coloring of disk graphs up to a constant factor.

For simplicity, we first prove the lower bound result for intersection graphs of squares (Section 2) and then provide arguments how the proof can be adapted to disks (Section 3). In Section 4, we prove that the First-fit heuristic gives competitive ratio $O(\log n)$ for disk graphs and square graphs. In Section 5, we consider disk graphs in which the ratio of the diameter of the largest disk to the diameter of the smallest disk, the so-called *diameter ratio*, is bounded by σ. We present an algorithm that makes use of the disk representation (in fact, it requires only knowledge of the diameters of the disks) to achieve an improved competitive ratio of $O(\min\{\log n, \log \sigma\})$, while the competitive ratio of First-fit can be bounded by $O(\min\{\log n, \sigma^2\})$ in this case. For square graphs, we call the ratio of the side length of the largest square to the side length of the smallest square the *side-length ratio* and we show that the results for disk graphs with bounded diameter ratio can be adapted to square graphs with bounded side-length ratio. We give our conclusions in Section 6. A summary of our results for on-line coloring of disk graphs is given in Table 1.

Table 1: Known results for on-line coloring of disk graphs with n vertices in the setting with (+) and without (−) given representation. UDG stands for unit disk graphs, DG_σ for disk graphs with diameter ratio bounded by σ, and DG for general disk graphs. New results obtained in this paper are shown in bold.

Graph class	Disk rep.	Competitive ratio lower bound	Competitive ratio upper bound
UDG	+	2 [6]	5 [16, 7]
DG_σ	+	$\boldsymbol{\Omega(\log\log\sigma)}$	$\boldsymbol{O(\min\{\log n, \log\sigma\})}$
DG	+	$\boldsymbol{\Omega(\log n)}$	$\boldsymbol{O(\log n)}$
UDG	−	2 [6]	5 [16, 7]
DG_σ	−	$\boldsymbol{\Omega(\log\log\sigma)}$	$\boldsymbol{O(\min\{\log n, \sigma^2\})}$
DG	−	$\Omega(\log n)$ [8]	$\boldsymbol{O(\log n)}$

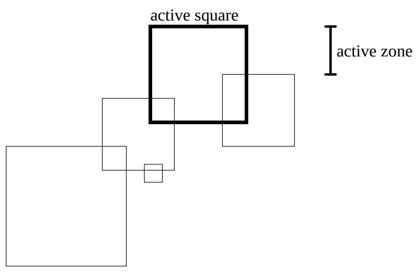

Figure 2: Example of an active square and its active zone.

2 A lower bound for on-line coloring of squares

Let $\mathcal{A}$ be an arbitrary deterministic on-line square coloring algorithm. We prove that for any k there exists a sequence $(S_1, S_2, \dots, S_n)$ of $n \leq 2^k$ squares on which $\mathcal{A}$ uses at least k colors, while the intersection graph of the squares is a tree and can be colored optimally with only two colors.

We deal only with squares whose sides are parallel to the axes of the coordinate system. The minimum and maximum x-coordinate of the square S_i are denoted by $\underline{x_i}$ and $\overline{x_i}$, respectively. Similarly, $\underline{y_i}$ and $\overline{y_i}$ denote the minimum and maximum y-coordinate, respectively. Thus, we have

$$S_i = \{(x, y) \mid \underline{x_i} \leq x \leq \overline{x_i},\ \underline{y_i} \leq y \leq \overline{y_i}\}.$$

We will denote a square S_i also by the tuple $(\underline{x_i}, \underline{y_i}, \overline{x_i}, \overline{y_i})$.

If $\mathcal{S}$ and $\mathcal{S}'$ are sequences of squares, then their concatenation is denoted by $\mathcal{S} \circ \mathcal{S}'$. If $\mathcal{S}$ is a nonempty sequence of squares, then $\mathcal{S}^-$ denotes the same sequence without the last element.

We say that squares of a sequence are *in general position* if every pair of squares differs in the maximum y-coordinate.

Now assume that the intersection graph of an arbitrary sequence of squares $(S_1, ..., S_n)$ in general position is a forest F. (Only such sequences will occur in our lower bound construction.) In each connected component of F we define the *active square* to be the one with the highest $\overline{y}$ coordinate. The *active zone* of an active square S_i is delimited by the interval $[a_1, a_2]$, where $a_2 = \overline{y_i}$ and $a_1 = \max_j\{\overline{y_j} : S_j \neq S_i$ and S_j belongs to the same connected component as $S_i\}$. If the connected component of S_i contains only S_i, we let $a_1 = \underline{y_i}$. See Figure 2 for an example.

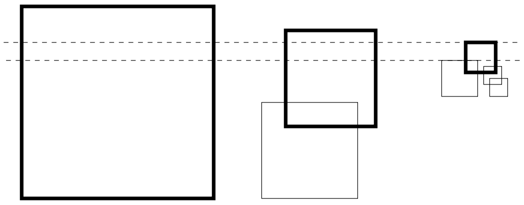

Figure 3: The active zone of a sequence of squares.

The active zone of a sequence of squares $\mathcal{S} = (S_1, \ldots, S_n)$ is defined as the intersection of the active zones of all active squares in $\mathcal{S}$. The width of the active zone is equal to the length of the corresponding interval, or is equal to 0 if the zone is empty. An example is shown in Figure 3, where the active squares are drawn in bold and the active zone of the sequence is indicated by two dashed lines.

The motivation for introducing active zones is as follows: If the sequence of squares presented to the algorithm so far has an active zone of positive width, then the adversary can present a new square that intersects all active squares but no other squares, and the resulting graph is still a tree. Therefore, the goal of the adversary is to construct sequences of squares with an active zone of positive width for which the algorithm assigns many different colors to active squares.

If all squares in a sequence $(S_1, S_2, \ldots, S_n)$ are contained in another square B, then we call B a *bounding square* for the sequence.

Lemma 1 *For each on-line square coloring algorithm $\mathcal{A}$, each $k \geq 2$, and each bounding square B with side length ℓ, there exists a sequence $\mathcal{S}$ of at most 2^k squares such that*

- *every $S \in \mathcal{S}$ is contained in B,*
- *the squares in $\mathcal{S}$ are in general position,*
- *the intersection graph of $\mathcal{S}$ is a tree,*
- *the smallest square in $\mathcal{S}$ has side length at least $4^{-2^{k-1}+1}\ell$,*
- *the active zones of $\mathcal{S}$ and $\mathcal{S}^-$ have width at least $4^{-2^{k-1}+1}\ell$,*
- *the algorithm $\mathcal{A}$ uses at least $k-1$ distinct colors on the active squares of $\mathcal{S}^-$, and*

- *the last square in $\mathcal{S}$ intersects all active squares of $\mathcal{S}^-$ and is the active square of $\mathcal{S}$.*

Proof: In the proof, we assume without loss of generality that B is the unit square $\{(x,y) \mid 0 \le x \le 1,\ 0 \le y \le 1\}$ and thus we have $\ell = 1$; if B is a different bounding square, all squares in the constructed sequence are scaled and shifted accordingly.

We prove the statement by induction over k. The statement is clearly true for $k = 2$: we can select $\mathcal{S}$ as two intersecting squares $(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}, \frac{1}{2})$ and $(\frac{1}{4}, \frac{1}{4}, \frac{3}{4}, \frac{3}{4})$. Clearly, the width of the active zones of $\mathcal{S}^-$ and $\mathcal{S}$ is at least $\frac{1}{4}$ and the smallest square has side length at least $\frac{1}{4}$.

Now assume that the statement is correct for k and we want to prove the statement for $k+1$. We apply the inductive hypothesis for k to the algorithm $\mathcal{A}$ and the bounding square $B' = (0, 0, \frac{1}{4}, \frac{1}{4})$. From this we obtain a sequence $\mathcal{S}'$ of at most 2^k squares, all contained in B', such that $\mathcal{A}$ uses at least $k-1$ distinct colors on the active squares of $\mathcal{S}'^-$. Furthermore, the active zone of $\mathcal{S}'^-$ has positive width.

Assume that the active zone of $\mathcal{S}'^-$ is an interval $[a_1, a_2]$ and let $\varepsilon = a_2 - a_1$. Note that $4^{-2^{k-1}+1} \cdot \frac{1}{4} \le \varepsilon \le \frac{1}{4}$ and that the smallest square in $\mathcal{S}'^-$ has side length at least $4^{-2^{k-1}+1} \cdot \frac{1}{4}$.

Now we apply the inductive hypothesis for k a second time to the algorithm $\mathcal{A}$ (after it has colored $\mathcal{S}'^-$) and the bounding square $B'' = (\frac{5}{12}, a_1, \frac{5}{12} + \varepsilon, a_2)$. This gives us a sequence $\mathcal{S}''$ of at most 2^k squares, all contained in B'', such that $\mathcal{A}$ uses at least $k-1$ colors on the active squares of $\mathcal{S}''^-$. Note that the squares in $\mathcal{S}''$ lie in the active zone of $\mathcal{S}'^-$, but are strictly to the right of all squares in $\mathcal{S}'^-$. We know that the active zones of $\mathcal{S}''$ and $\mathcal{S}''^-$ have width at least $4^{-2^{k-1}+1}\varepsilon \ge 4^{-2^k+1}$, and that the smallest square in $\mathcal{S}''$ has side length at least $4^{-2^{k-1}+1}\varepsilon \ge 4^{-2^k+1}$.

Now we have to consider two cases (illustrated in Figure 4):

(1) The sets of colors used on active squares of $\mathcal{S}'^-$ and $\mathcal{S}''^-$ are the same. Then the color of the last square S of $\mathcal{S}''$ is a new k-th color. Denote the active zone of S by $[b_1, b_2]$ and let $b = (b_1 + b_2)/2$. Then we can take
$$\mathcal{S} = \mathcal{S}'^- \circ \mathcal{S}'' \circ \left\{ \left(0, b, \frac{2}{3}, b + \frac{2}{3}\right) \right\}$$
and it is clear that all the conditions of the statement are satisfied for $k+1$.

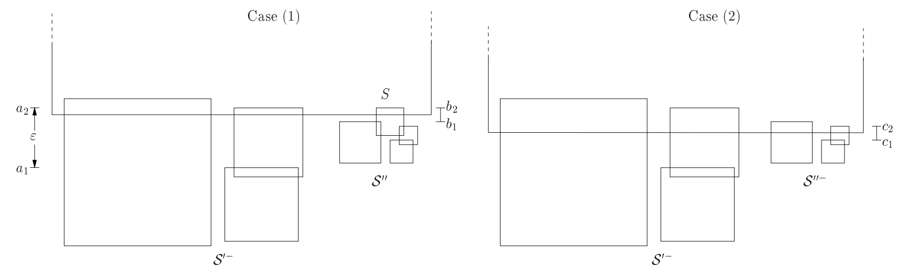

Figure 4: Illustration of lower bound construction.

(2) Active squares of $\mathcal{S}'^-$ and $\mathcal{S}''^-$ have different colors. This means that at least k different colors appear on active squares of the sequence $\mathcal{S}'^- \circ \mathcal{S}''^-$. Let the active zone of this sequence (which must have positive width) be $[c_1, c_2]$ and let $c = (c_1 + c_2)/2$. Then we can take

$$\mathcal{S} = \mathcal{S}'^- \circ \mathcal{S}''^- \circ \left\{ \left(0, c, \frac{2}{3}, c + \frac{2}{3}\right) \right\}$$

and it is again easy to see that the statement holds for $k+1$.

This completes the inductive step and thus the proof of the lemma. □

Substituting $n = 2^k$ in the lemma we get the following theorem.

Theorem 2 *For every on-line square coloring algorithm $\mathcal{A}$ and arbitrarily large values of n there exists a sequence of n squares $\mathcal{S}$ such that $\mathcal{A}$ uses $\Omega(\log n)$ distinct colors on $\mathcal{S}$ while $\mathcal{S}$ can be colored optimally with two colors.*

Furthermore, the proof of Lemma 1 shows that the ratio of the largest side length to the smallest side length of a square used in the lower bound construction is at most $4^{2^{k-1}-1}$. This leads to the following corollary.

Corollary 3 *There cannot be an on-line square coloring algorithm with competitive ratio $o(\log\log\sigma)$ for square graphs with side-length ratio bounded by σ, even if the square representation is given as part of the input.*

3 A lower bound for on-line coloring of disks

Instead of giving a detailed proof of the lower bound for disk graphs, we show how to adapt the approach of the previous section for square graphs and only point out which aspects require a different treatment.

Again we consider active disks and we define the active zone of an active disk analogously, i.e. as an interval $[a_1, a_2]$, where a_2 is the maximum y-coordinate of the active disk and a_1 is the maximum y-coordinate of any other disk in the same connected component. If the active disk does not intersect any other disk, we take a_1 to be its minimum y-coordinate. We use the notions of the active zone of a sequence, of a bounding square, and of being in general position as they were defined in the previous section.

We adapt the construction of the previous section in order to establish the following lemma.

Lemma 4 *For each on-line disk coloring algorithm $\mathcal{A}$, each $k \geq 2$, and each bounding square B with side length ℓ, there exists a sequence $\mathcal{D}$ of at most 2^k disks such that*

- *every $D \in \mathcal{D}$ is contained in B,*
- *the disks in $\mathcal{D}$ are in general position,*
- *the intersection graph of $\mathcal{D}$ is a tree,*
- *the smallest disk in $\mathcal{D}$ has diameter at least $4 \cdot 10^{-4^{k-2}} \ell$,*
- *the active zones of $\mathcal{D}$ and $\mathcal{D}^-$ have width at least $2 \cdot 10^{-4^{k-2}} \ell$,*
- *the algorithm $\mathcal{A}$ uses at least $k-1$ distinct colors on the active disks of $\mathcal{D}^-$, and*
- *the last disk in $\mathcal{D}$ intersects the active zones of all active disks of $\mathcal{D}^-$ and is the active disk of $\mathcal{D}$.*

The only problem that might arise in the new construction is that the diameter of the last disk in $\mathcal{D}$ might have to be very large in order to ensure that this disk intersects the tiny active zones of all active disks in $\mathcal{D}^-$, but no other disks. In this case, the first condition of the lemma might be violated. However, we can handle this issue as follows: We select the bounding squares B' and B'' for the application of the inductive hypothesis scaled by a factor depending on k such that the last active disk of the sequence $\mathcal{D}$ remains inside the bounding square B in any case that may arise after the recursive construction of the sets $\mathcal{D}'$ and $\mathcal{D}''$ (corresponding to $\mathcal{S}'$ and $\mathcal{S}''$ in the proof of Lemma 1).

Proof: We mimic the proof of Lemma 1 and assume without loss of generality that B is the unit square. For $k = 2$ the lemma holds if we choose the sequence consisting of two disks with centers at $(\frac{2}{5}, \frac{2}{5})$ and $(\frac{3}{5}, \frac{3}{5})$, both of diameter $\frac{2}{5}$.

Consider the inductive step from k to $k + 1$. We select B' of side length $\frac{1}{2}10^{-2\cdot 4^{k-2}}$ with lower left corner at $(\frac{1}{2} - 10^{-2\cdot 4^{k-2}}, \frac{1}{12})$. By the inductive hypothesis, there is a sequence $\mathcal{D}'$ for $\mathcal{A}$ consisting of disks contained in B' and having the properties claimed in the lemma. Similarly as in the case of squares we place the bounding square B'' into the active zone of $\mathcal{D}'^-$ to the right of B', say, with its left side at $x = \frac{1}{2}$. By the inductive hypothesis, there is a sequence $\mathcal{D}''$ of disks contained in B'' that satisfies the properties claimed in the lemma for the algorithm $\mathcal{A}$ after it has colored $\mathcal{D}'^-$.

The minimum diameter of a disk in $\mathcal{D}'^- \circ \mathcal{D}''$ is at least $4 \cdot 10^{-4^{k-1}}$ and the active zones of $\mathcal{D}''$ and $\mathcal{D}''^-$ have width at least $2 \cdot 10^{-4^{k-1}}$. Depending on the way the algorithm $\mathcal{A}$ colors the sequence $\mathcal{D}'^- \circ \mathcal{D}''^-$, we either extend the sequence $\mathcal{D}'^- \circ \mathcal{D}''$ or the sequence $\mathcal{D}'^- \circ \mathcal{D}''^-$ by adding the last disk D_n to get the desired sequence $\mathcal{D}$, analogous to the proof of Lemma 1. In the sequences $\mathcal{D}'^- \circ \mathcal{D}''$ and $\mathcal{D}'^- \circ \mathcal{D}''^-$, the active zone intersects the active disks in a strip of (horizontal) length less than $2 \cdot 10^{-2\cdot 4^{k-2}}$ and (vertical) width at least $2 \cdot 10^{-4^{k-1}}$. As is depicted in Figure 5, a last disk D_n intersecting all active disks can be found with diameter at most $\frac{1}{2} + 2 \cdot 10^{-4^{k-1}} < \frac{2}{3}$. Therefore, the entire sequence $\mathcal{D}$ is contained inside the unit square. (To place $\mathcal{D}$ into a prespecified square B different from the unit square, we shift and scale the bounding squares B' and B'' accordingly.) $\square$

Theorem 5 *For every on-line disk coloring algorithm $\mathcal{A}$ and arbitrarily large integers n there exists a sequence of n disks $\mathcal{D}$ such that $\mathcal{A}$ uses $\Omega(\log n)$ distinct colors on $\mathcal{D}$ while $\mathcal{D}$ can be colored optimally with two colors.*

Furthermore, we see that in the proof of Lemma 4 the diameter ratio σ of the disks in $\mathcal{D}$ is bounded from above by $\frac{1}{4} \cdot 10^{4^{k-2}}$. Since $\mathcal{A}$ uses at least k colors on this instance, the competitive ratio of $\mathcal{A}$ must be $\Omega(\log\log\sigma)$.

Corollary 6 *No deterministic on-line disk coloring algorithm can have competitive ratio $o(\log\log\sigma)$ on disk graphs with diameter ratio bounded by σ, even if the disk representation is given as part of the input.*

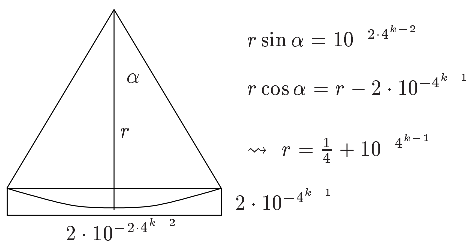

Figure 5: Disk D_n intersects the strip of horizontal length $2 \cdot 10^{-2\cdot 4^{k-2}}$ and of width $2 \cdot 10^{-4^{k-1}}$.

4 Upper bound for the First-fit algorithm

First-fit is one of the most well-known heuristics for on-line graph coloring. It simply assigns to each node the lowest-numbered color that has not yet been assigned to any of its neighbors. In order to analyze the competitive ratio of the First-fit on-line coloring algorithm for disk and square graphs, we make use of a result due to Irani [12]. A graph G is called *d-inductive* if the nodes of G can be ordered in such a way that each node has at most d edges to higher-numbered nodes. Irani's result is the following.

Theorem 7 (Irani, 1994) *If G is a d-inductive graph on n nodes, then First-fit uses $O(d \log n)$ colors to color G.*

We use $\omega(G)$ to denote the size of a maximum clique in G and obtain bounds on the inductiveness of disk graphs and square graphs in terms of $\omega(G)$.

Lemma 8 *Every disk graph G is $6\omega(G)$-inductive.*

Proof: Let $\mathcal{D}$ be a set of disks that is a disk representation of G. Order the disks in $\mathcal{D}$ according to non-decreasing diameters. Consider some disk D_i. We partition the higher-numbered disks that intersect D_i into six groups such that the disks in each group form a clique: A higher-numbered disk D_j that intersects D_i is assigned to group $\lfloor 3\alpha/\pi \rfloor$, where α is the angle that the line from the center of D_i to the center of D_j forms with the positive x axis.

See Figure 6 for an example with two disks in group 1. Since all higher-numbered disks are at least as big as D_i, simple geometric arguments show that any two disks assigned to the same group must intersect. $\square$

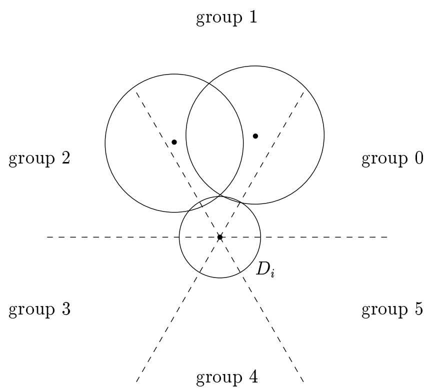

Figure 6: The larger disks intersecting D_i can be partitioned into six cliques.

Lemma 9 *Every square graph G is $4\omega(G)$-inductive.*

Proof: Order the squares in order of non-decreasing size. Consider some square S_i. All higher-numbered squares that intersect S_i contain at least one corner of S_i. Therefore, the higher-numbered squares that intersect S_i can be partitioned into four cliques. $\square$

From Theorem 7 and Lemmas 8 and 9 we get the following result.

Theorem 10 *First-fit uses $O(\omega(G)\log n)$ colors to color a disk graph or square graph G with n nodes.*

Since an optimal coloring requires at least $\omega(G)$ colors, we get the following corollary.

Corollary 11 *First-fit is an $O(\log n)$-competitive on-line coloring algorithm for disk graphs and square graphs with n nodes.*

Note that First-fit does not require the geometric representation of the disk graph or square graph.

5 Coloring disks with bounded diameter ratio or squares with bounded side-length ratio

We now focus our attention on the case that the ratio of the diameter of the largest disk and the smallest one is bounded by some value σ and that the disk representation is given as part of the input. (In fact, it would suffice that the diameters of the disks are given as part of the input.)

We prove that there exists an on-line coloring algorithm with competitive ratio $O(\min\{\log n, \log \sigma\})$. The algorithm is a composition of two algorithms: The first algorithm $\mathcal{A}$ is the First-fit algorithm for disk graphs with arbitrary diameter. It provides the bound $O(\log n)$. The second algorithm $\mathcal{B}$ is the First-fit technique applied separately on $\log \sigma$ layers of disks, where the diameters of the disks on each layer are within a factor of two so that First-fit has constant competitive ratio on each layer.

More precisely, the algorithm $\mathcal{B}$ colors disks $D_1, \ldots, D_n$ as follows:

First-fit on layers $\mathcal{B}$
$L_j := \emptyset$ for all $j \in \mathbb{Z}$;
for $i := 1$ **to** n **do**
  begin
    $j := \lfloor \log_2(diam(D_i)) \rfloor$;
    $L_j := L_j \cup \{D_i\}$; { *the layer containing* D_i }
    $F := \{c(D_k) : 1 \leq k < i, D_k \in L_j, D_k \cap D_i \neq \emptyset\} \cup$
      $\{c(D_k) : 1 \leq k < i, D_k \notin L_j\}$; { *the set of forbidden colors* }
    $c(D_i) := \min(\mathbb{N} \setminus F)$
  end.

Lemma 12 *The First-fit coloring algorithm is 28-competitive on disks of diameter ratio bounded by two.*

Proof: Assume that disks in the set $\mathcal{D}$ are scaled such that the smallest disk has unit diameter. The centers of all disks that intersect a particular disk D_i are at distance at most two from the center of D_i. We divide the plane around the center of D_i up to distance two into 28 segments such that inside each segment, every pair of points is at distance at most one, see Figure 7. (Similar proof techniques have been used in [6].) If the centers of two disks lie in the same segment, then these disks intersect, and hence disks with centers in the same segment induce a clique in the intersection graph G of $\mathcal{D}$.

Further we observe that the six inner segments contain also the center of D_i and therefore the vertex D_i can have at most $28\omega(G) - 6$ neighbors in G. Then the First-fit coloring algorithm uses at most $28\omega(G) - 5$ colors on G, and hence is 28-competitive. $\square$

Since algorithm $\mathcal{B}$ achieves constant competitive ratio on each layer and there are only $O(\log \sigma)$ different layers, we obtain the following lemma.

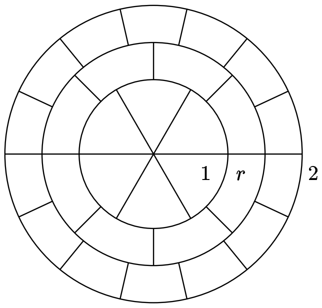

Figure 7: 28 segments of the plane around the center of D_i, $r = 1.306$.

Lemma 13 *If the disk representation is given as part of the input, First-fit on layers is an $O(\log \sigma)$-competitive coloring algorithm for disk graphs of diameter ratio bounded by σ.*

We combine First-fit and First-fit on layers as follows: We use two separate sets of colors for the algorithms $\mathcal{A}$ and $\mathcal{B}$. When a new disk D_i is presented we run $\mathcal{A}$ on D_i together with those disks colored by $\mathcal{A}$. Similarly we execute $\mathcal{B}$. Then we compare the results of these two algorithms and color D_i with the algorithm that has used fewer colors up to now (including disk D_i). The total number of colors used on the entire set $\mathcal{D}$ is the sum of the number of colors used by both algorithms. Note that at any time of the execution of the combined algorithm, the number of colors used by $\mathcal{A}$ and the number of colors used by $\mathcal{B}$ differ by at most one.

Assume that $\log n < \log \sigma$. The number of colors used by algorithm $\mathcal{A}$ is at most $O(\log n)$ times the optimal number. The number of colors used by algorithm $\mathcal{B}$ is at most one more than that of $\mathcal{A}$. So the total number of colors used by the combined algorithm is $O(\log n)$ times the optimal number of colors. A symmetric argument holds in the case that $\log \sigma \leq \log n$. Therefore, we obtain the following theorem.

Theorem 14 *If the disk representation is given as part of the input, there is an $O(\min\{\log n, \log \sigma\})$-competitive coloring algorithm for disk graphs whose diameter ratio is bounded by σ.*

An analogous result can be obtained for square intersection graphs as well. Here, we let σ denote the ratio of the largest side length of a square to the smallest side length, i.e. the side-length ratio. Lemma 12 can be adapted as follows.

Lemma 15 *The First-fit coloring algorithm is 16-competitive on squares of side-length ratio bounded by two.*

Proof: Assume that the squares are scaled such that the smallest square has side-length one. Consider a particular square S_i. Assume that the center of S_i is at the origin. Then the centers of all squares that intersect S_i are within the square with lower left endpoint $(-2,-2)$ and upper right endpoint $(2,2)$. Partition this square into 16 subsquares with side length 1 in the natural way. Any two squares with centers in the same subsquare must intersect. Furthermore, the four subsquares touching the origin contain the center of S_i. Therefore, S_i can have at most $16\omega(G)-4$ neighbors in the square graph. Then the First-fit coloring algorithm uses at most $16\omega(G)-3$ colors, and hence is 16-competitive. □

Lemma 13 and Theorem 14 can then be adapted to square graphs directly. So we obtain the analogous result for square graphs.

Theorem 16 *If the square representation is given as part of the input, there is an* $O(\min\{\log n, \log\sigma\})$*-competitive coloring algorithm for square graphs whose side-length ratio is bounded by* σ*.*

Concerning on-line coloring of disk graphs (or square graphs) with diameter ratio (side-length ratio) bounded by σ in the case without given representation, First-fit is easily seen to be $O(\sigma^2)$-competitive. This follows because the neighborhood of a disk or square can be covered by $O(\sigma^2)$ cliques. The idea of the analysis is the same as the one used in the proofs of Lemmas 12 and 15 for the case $\sigma = 2$. Combining this with the upper bound of $O(\log n)$, we get that First-fit is $O(\min\{\log n, \sigma^2\})$-competitive.

6 Conclusion

We have shown that the First-fit algorithm, which does not need the disk representation as part of the input, provides an $O(\log n)$-competitive disk coloring algorithm and that no algorithm can have competitive ratio $o(\log n)$, even if it uses the geometric representation. For instances with diameter ratio bounded by σ, we showed that the geometric representation can help to get a better ratio of $O(\min\{\log n, \log \sigma\})$. Our lower bound on the competitive ratio of any on-line disk coloring algorithm is $\Omega(\log\log \sigma)$ in this case. Analogous results hold for intersection graphs of squares.

We initiated the study of on-line coloring of disk graphs with bounded diameter ratio and square graphs with bounded side-length ratio. However, for these particular problems the gaps between the lower and upper bounds on the competitive ratio should be narrowed.

The most widely used lower bound on the chromatic number of a disk graph (i.e., the lower bound on the optimal solution) is expressed via the clique number of the graph. We hope that by use of more sophisticated arguments it could be proven that standard coloring algorithms behave even better. As a particular open problem we would ask what is the supremum of the ratio of chromatic and clique number of a unit disk graph. The only known bounds are $1.5 \leq \frac{\chi(G)}{\omega(G)} < 3$. The lower bound is derived from the coloring of the cycle C_5, and the upper bound is achieved by the algorithm due to Peeters [17].

Furthermore, it should be noted that many questions are still open for rectangle intersection graphs, a generalization of square intersection graphs. The recognition problem for rectangle intersection graphs has been proved NP-complete in [14]. For coloring a rectangle intersection graph G with clique number $\omega(G)$, it is known that $O(\omega(G)^2)$ colors suffice [3], but no example with a non-linear lower bound in terms of the clique number has been obtained. For some special cases, it is known that $O(\omega(G))$ colors suffice [15]. It would be an interesting problem for future research to investigate off-line and on-line coloring of rectangle intersection graphs.

Finally, we believe that the use of standard methods like randomized algorithms might improve the competitive ratio and we expect further results in this direction.

Acknowledgment

We would like to thank the anonymous referees for their useful suggestions.

References

[1] H. Breu and D. G. Kirkpatrick. On the complexity of recognizing intersection and touching graphs of disks. In *Proceedings of the 3rd Sym-*

posium on Graph Drawing GD'95, Lecture Notes in Computer Science 1027, pages 88–98, 1996.

[2] H. Breu and D. G. Kirkpatrick. Unit disk graph recognition is NP-hard. *Computational Geometry: Theory and Applications*, 9(1–2):3–24, 1998.

[3] J. Burling. *On coloring problems of families of prototypes.* PhD thesis, University of Colorado, 1965.

[4] B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit disk graphs. *Discrete Mathematics*, 86(1–3):165–177, 1990.

[5] T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric graphs. In *Proceedings of the 12th Annual ACM–SIAM Symposium on Discrete Algorithms SODA'01*, pages 671–679, 2001.

[6] J. Fiala, A. V. Fishkin, and F. V. Fomin. Off-line and on-line distance constrained labeling of graphs. Manuscript. An extended abstract was published in *Proceedings of the 9th European Symposium on Algorithms ESA'01*, Lecture Notes in Computer Science 2161, pages 464–475, 2001.

[7] A. Gräf. Coloring and recognizing special graph classes. Musikinformatik und Medientechnik Bericht 20/95, Johannes Gutenberg-Universität Mainz, 1995.

[8] A. Gyárfás and J. Lehel. On-line and first fit colourings of graphs. *Jornal of Graph Theory*, 12(2):217–227, 1988.

[9] W. K. Hale. Frequency assignment: Theory and applications. *Proceedings of the IEEE*, 68(12):1497–1514, 1980.

[10] P. Hliněný and J. Kratochvíl. Representing graphs by disks and balls. *Discrete Mathematics*, 229(1–3):101–124, 2001.

[11] H. B. Hunt III, M. V. Marathe, V. Radhakrishnan, S. S. Ravi, D. J. Rosenkrantz, and R. E. Stearns. NC-Approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *Journal of Algorithms*, 26(2):238–274, 1998.

[12] S. Irani. Coloring inductive graphs on-line. *Algorithmica*, 11(1):53–72, 1994.

[13] P. Koebe. Kontaktprobleme der konformen Abbildung. *Ber. Verh. Sächs. Akad. Leipzig*, 88:141–164, 1936.

[14] J. Kratochvíl. A special planar satisfiability problem and a consequence of its NP-completeness. *Discrete Applied Mathematics*, 52(3):233–252, 1994.

[15] E. Malesińska. *Graph theoretical models for frequency assignment problems.* PhD thesis, Technical University of Berlin, 1997.

[16] M. V. Marathe, H. Breu, H. B. Hunt III, S. S. Ravi, and D. J. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25(2):59–68, 1995.

[17] R. Peeters. On coloring j-unit sphere graphs. Technical report, Department of Economics, Tilburg University, 1991.

[18] V. Raghavan and J. Spinrad. Robust algorithms for restricted domains. In *Proceedings of the 12th Annual ACM–SIAM Symposium on Discrete Algorithms SODA'01*, pages 460–467, 2001.

[19] H. Sachs. Coin graphs, polyhedra, and conformal mapping. *Discrete Mathematics*, 134(1–3):133–138, 1994.