A Tight Bound for Online Coloring of Disk Graphs*

Ioannis Caragiannis[1], Aleksei V. Fishkin[2], Christos Kaklamanis[1], and Evi Papaioannou[1]

[1] Research Academic Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece
[2] Max Planck Institut für Informatik, Stuhlsatzenhausweg 85, Geb. 46.1, 66123 Saarbrücken, Germany

Abstract. We present an improved upper bound on the competitiveness of the online coloring algorithm First-Fit in disk graphs which are graphs representing overlaps of disks on the plane. We also show that this bound is best possible for deterministic online coloring algorithms that do not use the disk representation of the input graph. We also present a related new lower bound for unit disk graphs.

1 Introduction

We study minimum coloring, a fundamental combinatorial optimization problem in graphs. Given a graph G, the minimum coloring problem is to find an assignment of colors (denoted by positive integers) to the nodes of the graph so that no two nodes connected by an edge are assigned the same color and the number of colors used is minimized. We consider intersection graphs modeling overlaps of disks on the plane.

The intersection graph of a set of disks in the Euclidean plane is the graph having a node for each disk and an edge between two nodes if and only if the corresponding disks overlap. Each disk is defined by its radius and the coordinates of its center. Two disks overlap if the distance between their centers is at most equal to the sum of their radii. A graph G is called a *disk graph* if there exists a set of disks in the Euclidean plane whose intersection graph is G. The set of disks is called the disk representation of G. A disk graph is called *unit disk graph* if all disks in its disk representation have the same radius. A disk graph is *σ-bounded* if the ratio between the maximum and the minimum radius among all the disks in its disk representation is at most σ.

In disk graphs, minimum coloring is important since it can model frequency assignment problems in radio communication networks utilizing the Frequency Division Multiplexing technology [10]. Consider a set of transmitters located in

* This work was partially supported by the European Union under IST FET Project CRESCCO.

A. Pelc and M. Raynal (Eds.): SIROCCO 2005, LNCS 3499, pp. 78–88, 2005.

fixed positions within a geographical region. Each transmitter may select to use a specific frequency from an available spectrum in order to transmit its messages. Two transmitters can successfully (i.e., without signal interference) transmit messages simultaneously either if they use different frequencies or if they use the same frequency and their ranges do not overlap. Given a set of transmitters in a radio network, in order to guarantee successful transmissions simultaneously, the important engineering problem to be solved is the frequency assignment problem where the objective is to minimize the number of frequencies used all over the network. Assuming that all transmitters have circular range, the graph reflecting possible interference between pairs of transmitters is a disk graph. The frequency assignment problem is equivalent to minimum coloring.

An instance of the minimum coloring problem may or may not include the disk representation (i.e., disk center coordinates and/or radii) of the disk graph as part of the input. Clearly, the latter case is more difficult. Information about the disk representation of a disk graph is not easy to extract. Actually, determining whether a graph is a disk graph is an NP-complete problem [11].

The minimum coloring problem has been proved to be NP-hard in [3, 8] even for unit disk graphs. A naive algorithm is algorithm First-Fit: it examines the nodes of the graph in an arbitrary order and assigns to each node the smallest color not assigned to its already examined neighbors. Clearly, First-Fit does not use the disk representation. It computes 5-approximate solutions in unit disk graphs [7, 14]. By processing the nodes of the graph in a specific order, First-Fit computes 3-approximate solutions in unit disk graphs [8, 14, 15]. In general disk graphs, a smallest-degree-last version of First-Fit achieves an approximation ratio of 5 [7, 13, 14].

In the online versions of the problem, the disk graph is not given in advance but is revealed in steps. In each step, a node of the graph appears together with its edges incident to nodes appeared in previous steps (and possibly, together with the center coordinates and/or the radius of the corresponding disk). When a node appears, an online coloring algorithm decides which color to assign to the node. The decisions of the algorithm at a step cannot change in the future.

The performance of an online algorithm is measured in terms of its competitive ratio (or competitiveness, [1]) which is defined as the maximum over all possible sequences of disks of the ratio of the number of colors used by the algorithm over the minimum number of colors sufficient for coloring the graph (i.e., its chromatic number).

First-Fit is essentially an online algorithm. It has been widely studied in a more general context and has been proved to be $\Theta(\log n)$-competitive in inductive graphs with n nodes [12, 9]. Disk graphs are inductive [4, 5] so the upper bound holds for disk graphs as well. The lower bound holds also for trees (which are disk graphs) so the $\Theta(\log n)$ bound holds for general disk graphs. In unit disk graphs, First-Fit is at most 5-competitive [7, 14] while for σ-bounded disk graphs with n nodes, it is at most $O(\min\{\log n, \sigma^2\})$-competitive [4]. For unit disk graphs, a lower bound of 2 on the competitiveness of any deterministic online coloring algorithm is presented in [6]. The best known lower bound on the competitiveness of deterministic coloring algorithms in σ-bounded disk graphs is $\Omega(\min\{\log n, \log\log\sigma\})$ [4]. A competitive ratio of $O(\min\{\log n, \log\sigma\})$ is achieved by two algorithms presented in [4] and [2]. The former uses the disk representation while the latter does not but it is quite impractival. Both algorithms use First-Fit as a subroutine.

In this paper we show that algorithm First-Fit itself is $O(\log\sigma)$-competitive when applied to σ-bounded disk graphs. This significantly improves the previously known upper bound of $O(\sigma^2)$ on the competitiveness of First-Fit. Furthermore, it matches the best known upper bound for online deterministic coloring algorithms, previously achieved either by algorithms that use the disk representation [4] or by quite impractical algorithms that do not use the disk representation [2]. Our second result indicates that First-Fit has optimal competitiveness (within constant factors) among all deterministic online algorithms for disk graphs that do not use the disk representation. In particular, we show that any deterministic online coloring algorithm that does not use the disk representation has competitive ratio $\Omega(\log\sigma)$ on σ-bounded disk graphs. Combined with previous results, our lower bound establishes a tight bound of $\Theta(\min\{\log n, \log\sigma\})$ on the optimal competitiveness of deterministic online coloring algorithms in σ-bounded disk graphs with n nodes that do not use the disk representation. We also prove a new lower bound of 2.5 on the competitiveness of deterministic online coloring algorithms for unit disk graphs that do not use the disk representation. This result improves a previous lower bound of 2 [6].

The rest of the paper is structured as follows. A discussion on previous upper bounds and the proof of the upper bound on the competitiveness of First-Fit are presented in Section 2. The lower bounds are presented in Section 3. We conclude with open problems in Section 4.

2 The Upper Bound

In this section, we prove the upper bound for algorithm First-Fit. Although this upper bound can also be achieved by two other known algorithms presented in [4] and [2], respectively, our result is important because of the simplicity of algorithm First-Fit.

The algorithm of Erlebach and Fiala [4] classifies the disks into a logarithmic number of classes so that the disks belonging to the same class form a 2-bounded disk graph and runs algorithm First-Fit in each class using disjoint sets of colors for coloring the disks of different classes. The classification is performed according to the radii of the disks; hence, the algorithm uses the disk representation. The proof of the $O(\log\sigma)$ upper bound follows by the fact that algorithm First-Fit has constant competitive ratio on 2-bounded disk graphs.

The algorithm Layered classifies the disks into layers and applies algorithm First-Fit to each layer separately, using a different set of colors in each layer. Layers are numbered with integers $1, 2, ...$ and a disk is classified into the smallest layer possible under the constraint that it cannot be classified into a layer if it overlaps with at least 16 mutually non-overlapping disks belonging to this layer.

The proof that algorithm Layered is $O(\log \sigma)$-competitive is based on the following arguments. First, if a disk of radius R belongs to some layer $i > 1$, then there is a disk of radius at most $R/2$ belonging to layer $i-1$. Hence, if the disk graph given as input to algorithm Layered is σ-bounded, the number of layers is at most $1 + \log \sigma$. The logarithmic competitive ratio follows since the maximum independent set in the neighborhood of each node within each layer has size at most 15, and, hence, algorithm First-Fit is proved to have constant competitive ratio within each layer. Clearly, algorithm Layered does not use the disk representation.

For checking whether a new node presented has 16 or more non-overlapping disks of some layer in its neighborhood may require time $\Omega(n^{16})$. This could be decreased to $\Omega(n^8)$ by changing the constraint so that a disk cannot be classified into a layer if it overlaps with at least 8 mutually non-overlapping disks belonging to this layer. Still, it can be proved that there exists a constant $\alpha > 1$ such that for each disk of radius R belonging to some layer $i > 1$, there exists a disk at layer $i-1$ of radius smaller than R/α. This is the best possible improvement in the idea of algorithm Layered since, for any $\alpha > 1$ arbitrarily close to 1 (e.g., $\alpha = 1 + 1/\sigma$), a disk of radius R can overlap with 7 mutually non-overlapping disks of radius R/α, and, hence, the logarithmic upper bound on the number of layers cannot be established.

Surprisingly, we show that algorithm First-Fit itself is at most $O(\log \sigma)$-competitive, improving the previously known $O(\sigma^2)$ upper bound. Combining this result with the $O(\log n)$ upper bound which is known for the competitive ratio of First-Fit we obtain that First-Fit is $O(\min\{\log n, \log \sigma\})$-competitive. Algorithm First-Fit runs in time proportional to the number of edges of the disk graph, i.e., $O(n^2)$, and does not use the disk representation. Hence, it is much simpler than the previously known algorithms that achieve the same bounds.

Theorem 1. First-Fit *is* $O(\log \sigma)$*-competitive for* σ*-bounded disk graphs.*

Proof. Let G be a σ-bounded disk graph with chromatic number κ. Assume that the nodes of G appear online and are colored by algorithm First-Fit. Consider a representation of G by overlapping disks on the plane of radii between r (the radius of the smallest disk) and $R_{\max}$ (the radius of the largest disk) so that $R_{\max}/r \leq \sigma$. We classify the nodes into levels $0, 1, ..., \lfloor \log (R_{\max}/r) \rfloor$ as follows: a node corresponding to a disk of radius R belongs to level $\lfloor \log (R/r) \rfloor$. Since $R_{\max}/r \leq \sigma$, the index of the last level is at most $\lfloor \log \sigma \rfloor$.

We will first show that a node of G belonging to level $i \geq 0$ is adjacent to at most $15(\kappa - 1)$ other nodes of level at least i.

Assume otherwise that there exists a node u of G at level i which is adjacent to at least $15\kappa - 14$ other nodes of level at least i. Let R be the radius of the disk d corresponding to node u in the disk representation. Then $i = \lfloor \log (R/r) \rfloor$. Also, let S_d be the set of disks corresponding to nodes adjacent to d which belong to levels at least i. Clearly, all the disks of S_d have radii at least $r2^i$.

We apply the following shrinking procedure on the disks of S_d. We shrink each disk d' in S_d into a disk of radius $r2^{\lfloor \log R/r \rfloor}$ as follows: If the center $c_{d'}$ of d' is inside d, we shrink d' into a disk having the same center $c_{d'}$. Otherwise, let

$p_{d'}$ be the point in the periphery of d' which is closest to the center of d. We shrink d' so that $p_{d'}$ is again the point in the periphery of d' which is closest to the center of d. Denote by S'_d the set of shrunk disks. Clearly, each of the disks in S'_d overlaps with d since either its center is contained in d or a point in its periphery is contained in d. This means that all disks in S'_d are completely contained into the disk of radius $R + 2^{i+1}$ centered at the center c_d of disk d. An example of the shrinking procedure is depicted in Figure 1.

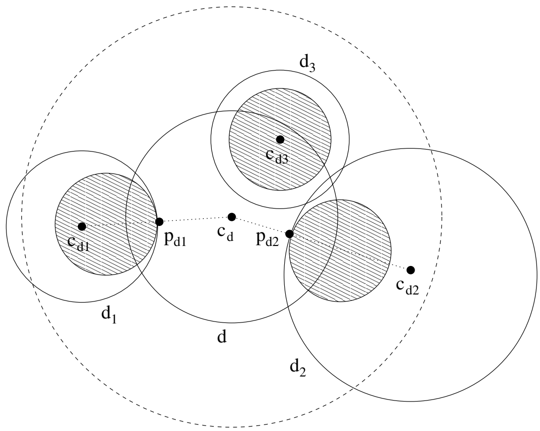

Fig. 1. The shrinking procedure. The disk d overlaps with three disks. Grey disks are the three corresponding shrunk disks

The node-induced subgraph H of G defined by the nodes of G corresponding to the disks of S_d is $(\kappa - 1)$-colorable since the graph G is κ-colorable and the nodes of H are all adjacent to u in G. Consequently, since by our assumption H contains the neighbors of u in G with levels at least i and since there are at least $15\kappa - 14$ such nodes, the maximum independent set in H has size at least 16. Consider such an independent set of 16 nodes in H and the 16 non-overlapping disks of S_d corresponding to these nodes. Clearly, the 16 corresponding shrunk disks of S'_d are also non-overlapping. Each of these disks has radius $r2^i$ and, hence, their total area is

$$16\pi \left(r2^i\right)^2 = \pi \left(r2^{i+1} + r2^{i+1}\right)^2$$
$$> \pi \left(R + r2^{i+1}\right)^2 .$$

Since these disks are non-overlapping, this contradicts the fact that all disks of S'_d are completely contained in the disk of radius $R + r2^{i+1}$ centered at c_d. Consequently, our assumption is incorrect and u is adjacent to at most $15(\kappa - 1)$ other nodes of G of level at least i.

We will now show that each node of G at level $i \geq 0$ is colored by algorithm First-Fit with a color in the range $[1, (15\kappa - 14)(i + 1)]$. Hence, the maximum color that can be assigned to a node of G by First-Fit is at most

$(15\kappa - 14)(\lfloor \log \sigma \rfloor + 1)$. Since G has chromatic number κ, this implies that algorithm First-Fit is $O(\log \sigma)$-competitive.

We use induction on the level of nodes. The statement is true for nodes of level 0, since any such node is adjacent to at most $15(\kappa - 1)$ nodes of G and, hence, it will be assigned a color in $[1, 15\kappa - 14]$.

Now assume that the statement is true for nodes of level $i = 0, ..., k$ (for $k < \lfloor \log \sigma \rfloor$). We will show that it also holds for nodes of level $k + 1$. Consider a node u at level $k+1$. This node will be adjacent to nodes of smaller levels which may use up to color $(15\kappa-14)(k+1)$ and to at most $15(\kappa-1)$ additional nodes of levels at least $k+1$. Hence, the maximum color that can be assigned by algorithm First-Fit to node u is $(15\kappa - 14)(k + 1) + 15(\kappa - 1) + 1 = (15\kappa - 14)(k + 2)$. This completes the proof of the theorem. □

3 The Lower Bound

The result proved in the following establishes that algorithm First-Fit achieves the best possible competitive ratio (within constant factors) among all deterministic online coloring algorithms that do not use the disk representation.

We present an adversary ADV which, on input an integer k and a deterministic online coloring algorithm A, outputs a tree T with at most 2^{k-1} nodes such that A colors T with at least k colors. We describe the adversary ADV in the following. This is a non-recursive description of the adversary very similar to that used in [4] for proving lower bounds on disk graphs and (in a more general form) in [12] for proving lower bounds on inductive graphs. We use the notation $\langle s, t\rangle$ to represent the nodes in the tree but the root, where $t \geq 1$ is an integer representing the level of the node and s a binary string of length $k - t - 1$. We use the function $str()$ which, on input integers $i \geq 0$ and j with $0 \leq j \leq 2^i$, returns a string of length i (possibly empty) which is the binary representation of j with i binary digits. We use the symbol $\odot$ to denote the concatenation of two strings. The adversary ADV can be described as follows.

Create 2^{k-2} nodes labeled as $\langle str(i, k-2), 1\rangle$, for each $i = 0, ..., 2^{k-2} - 1$, and introduce them to algorithm A.
For $i = 2$ to $k - 1$
  For $j = 0$ to $2^{k-i-1} - 1$
    Let S_ℓ be the set of colors assigned by A to nodes

$$\langle str(j, k-i-1) \odot \overbrace{000...00}^{i-1 \text{ times}}, 1\rangle, \langle str(j, k-i-1) \odot \overbrace{00...00}^{i-2}, 2\rangle,$$

$$..., \langle str(j, k-i-1) \odot 0, i-1\rangle.$$

    Let S_r be the set of colors assigned by A to nodes

$$\langle str(j, k-i-1) \odot 1\overbrace{00...00}^{i-2}, 1\rangle, \langle str(j, k-i-1) \odot 1\overbrace{0..00}^{i-3}, 2\rangle,$$

$$..., \langle str(j, k-i-1) \odot 1, i-1\rangle.$$

If $S_\ell = S_r$ then
Create a new node $\langle str(j,k-i-1),i\rangle$ connected to nodes

$$\langle str(j,k-i-1)\odot 1\overbrace{00..00}^{i-2},1\rangle, \langle str(j,k-i-1)\odot 1\overbrace{0..00}^{i-3},2\rangle,$$

$$..., \langle str(j,k-i-1)\odot 1, i-1\rangle$$

and introduce it to algorithm A.
else
Let $\langle s,t\rangle$ be the node to which A assigns a color not in S_ℓ.
Rename it as $\langle str(j,k-i-1),i\rangle$
Endif
Endfor
Endfor
Create a new node r connected to nodes labeled

$$\langle\overbrace{000...00}^{k-2},1\rangle, \langle\overbrace{000...00}^{k-3},2\rangle, ..., \langle 0,k-2\rangle, \langle\emptyset,k-1\rangle,$$

and introduce it to algorithm A.

The adversary forces the algorithm to use at least k colors. In each iteration, it can be shown by induction on i that all the $i-1$ nodes examined for defining the set S_ℓ (and similarly for S_r) are colored with $i-1$ different colors. Hence, after the if-then-else statement, the adversary will have forced algorithm A to use i different colors, i.e., $k-1$ colors at the end of all iterations. This clearly holds if the sets S_ℓ an S_r are not the same (else statement). Otherwise, it is guaranteed by the introduction of a new node which is connected to nodes colored with the $i-1$ different colors of S_r (if-then statement). Then, the last node r is connected to nodes with $k-1$ different colors and will be assigned a k-th color by algorithm A.

Also, it can be seen that when a new node is introduced in an iteration, the nodes to which it is connected will not be connected to other nodes in subsequent iterations. Hence, in general, the resulting graph is a forest. The number of nodes is at most 2^{k-1} since there are 2^{k-2} nodes at level 1, at most one new node in each iteration and one more node at the end. Actually, when the adversary runs against algorithm First-Fit, then the constructed graph is a tree with exactly 2^{k-1} nodes (in each iteration, a new node is introduced). We denote this tree by $T_{FF}(k)$ and we will first show that this is an α^{k-1}-bounded disk graph, for every $\alpha > 2$. Then, we will show how to adapt this construction for forests produced by the adversary against other deterministic online coloring algorithms.

Given a disk d of radius R corresponding to some node of the tree $T_{FF}(k)$, we define the vertical stripe of d to be the vertical stripe of width $2R$ which completely contains d. In our construction, the disk representation of $T_{FF}(k)$ is such that the disks corresponding to nodes in the subtree of a node u do not cross the boundaries of the vertical stripe of the disk d corresponding to u. Furthermore, the vertical stripes of any two disks corresponding to children of the same node are disjoint. These two invariants guarantee that the disks corresponding to nodes belonging to different subtrees do not overlap.

We first locate a disk of radius α^{k-1} corresponding to the root r of the tree. Disks corresponding to nodes of level i (for $i = 1, ..., k-1$) will have radius α^{i-1}.

A node u at level i with $i = 2, ..., k$, has $i-1$ children u_1, ..., u_{i-1} in $T_{FF}(k)$ with levels $1, ..., i-1$, respectively. Let d be the disk corresponding to node u and let $d_1, ..., d_{i-1}$ be the disks corresponding to its children u_1, ..., u_{i-1}, respectively. Assuming that the center of the disk d has horizontal coordinate h, the center of the disk d_j has horizontal coordinate $h_j = h - \alpha^{i-1} + 3\alpha^{i-1}2^{-i+j}$. The horizontal stripes of the disks $d_1, ..., d_{i-1}$ are disjoint since for two disks d_j and $d_{j'}$ with $j > j'$, their centers differ in the horizontal coordinate by

$$\begin{aligned}
h_j - h_{j'} &= (h - \alpha^{i-1} + 3\alpha^{i-1}2^{-i+j}) - (h - \alpha^{i-1} + 3\alpha^{i-1}2^{-i+j'}) \\
&= \frac{3}{2}\alpha^{j-1}(\alpha/2)^{i-j}(2 - 2^{-j+j'+1}) \\
&> \frac{3}{2}\alpha^{j-1} \\
&> \alpha^{j-1} + \alpha^{j'-1}
\end{aligned}$$

which is the sum of their radii.

Furthermore, neither the leftmost disk d_1 nor the rightmost disk d_{i-1} cross the boundary of the horizontal stripe of d. Indeed, the leftmost point of d_1 has horizontal coordinate

$$\begin{aligned}
h_1 - 1 &= h - \alpha^{i-1} + 3\alpha^{i-1}2^{-i+1} - 1 \\
&= h - \alpha^{i-1} + 3(\alpha/2)^{i-1} - 1 \\
&> h - \alpha^{i-1} + 2 \\
&> h - \alpha^{i-1}
\end{aligned}$$

which is the horizontal coordinate of the left boundary of the horizontal stripe of d. Also, the rightmost point of d_{i-1} has horizontal coordinate

$$\begin{aligned}
h_{i-1} + \alpha^{i-2} &= h - \alpha^{i-1} + 3\alpha^{i-1}/2 + \alpha^{i-2} \\
&= h + \alpha^{i-1}/2 + \alpha^{i-2} \\
&< h + \alpha^{i-1}
\end{aligned}$$

which is the horizontal coordinate of the right boundary of the horizontal stripe of d.

The vertical coordinate of the center of disk d_j is defined so that it is smaller than the vertical coordinate of the lowest point in the intersection of disk d_j with disk d. This guarantees that, among all disks corresponding to nodes in the subtree of u, the disks that d overlaps with are those corresponding to its children in $T_{FF}(k)$.

In the disk representation of the tree $T_{FF}(k)$, we use disks of radii between 1 and α^{k-1}. So, $T_{FF}(k)$ is an α^{k-1}-bounded disk graph. An example of the construction is depicted in Figure 2.

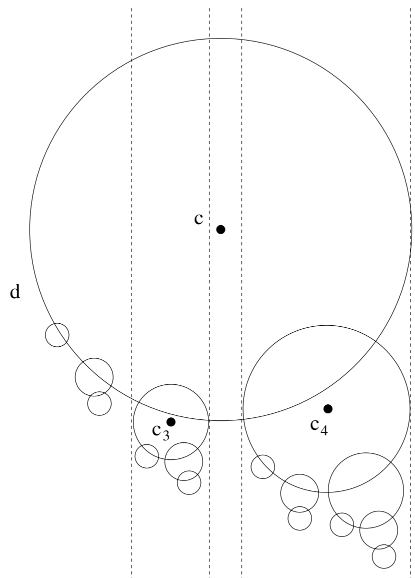

Fig. 2. The disk representation of the tree produced by algorithm ADV on input $k = 5$ and algorithm First-Fit. The dashed lines indicate the boundaries of the horizontal stripes of two disks

Now, consider the forest created by the adversary on input an integer k and some other algorithm A. In this case, some iterations may have renamed some nodes instead of introducing new ones. We construct the disk representation of such a forest by starting with the disk representation of $T_{FF}(k)$. We follow the execution of the adversary on input k and some algorithm A. When, the adversary executes the else statement in an iteration, we remove the disk corresponding to the node $\langle str(j, k-i-1), i\rangle$ (since this node is not introduced) and move horizontally all the disks corresponding to nodes of the subtree of node $\langle s, t\rangle$ so that the disk d corresponding to node $\langle s, t\rangle$ (and no other disk in its subtree) overlaps with the disk corresponding to the parent node of node $\langle str(j, k-i-1), i\rangle$ in $T_{FF}(k)$. From now on, the node to which disk d corresponds has been renamed as $\langle str(j, k-i-1), i\rangle$.

Clearly, this also yields an α^{k-1}-bounded disk graph. The above discussion leads to the following lemma.

Lemma 1. *For any $\alpha > 2$, the forest constructed by the adversary ADV on input an integer $k \geq 3$ and any deterministic online coloring algorithm A is an α^{k-1}-bounded graph.*

Now given a sufficiently large σ, the graph produced by the adversary on input $k = 1 + \lfloor \log_\alpha \sigma \rfloor$ and any deterministic algorithm is a 2-colorable σ-bounded disk graph which the algorithm colors with at least $1 + \lfloor \log_\alpha \sigma \rfloor$ colors. We obtain the following.

Theorem 2. *Any deterministic online algorithm for coloring σ-bounded disk graphs that does not use the disk representation has competitive ratio $\Omega(\log \sigma)$.*

For unit disk graphs, the best known lower bound on the competitiveness of any deterministic algorithm is 2 [6] and holds also for algorithms that use the disk representation. On input a deterministic online algorithm A, the adversary in the proof of [6] constructs a κ-colorable unit disk graph with $\kappa \in \{1, 2, 3\}$, which algorithm A colors with at least 2κ colors. In the following, we improve this lower bound for deterministic online coloring algorithm in unit disk graphs that do not use the disk representation.

Consider a deterministic online coloring algorithm A and the forest produced by adversary ADV on input 5 and algorithm A. Each connected component of the forest produced by ADV is a subtree of the tree $T_{FF}(5)$ produced by ADV on input 5 and algorithm First-Fit. The tree $T_{FF}(5)$ is a unit disk graph as shown in Figure 3.

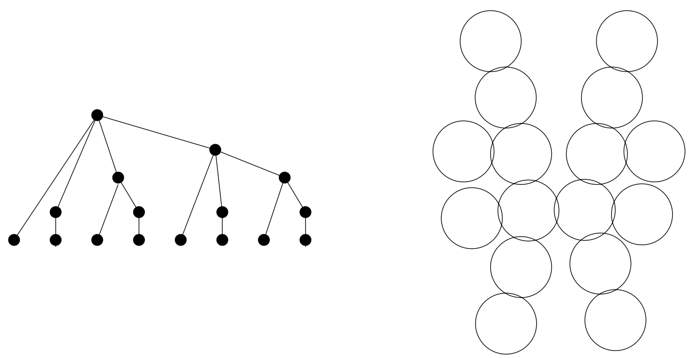

Fig. 3. The tree $T_{FF}(5)$ and its disk representation with unit disks

Hence, the output of the adversary ADV on input 5 and any deterministic algorithm A is a unit disk graph. Since this output is a forest, it is 2-colorable, while the adversary forces A to use at least 5 colors. We obtain the following.

Theorem 3. *Any deterministic online algorithm for coloring unit disk graphs that does not use the disk representation has competitive ratio at least* 2.5.

4 Open Problems

Our results on σ-bounded disk graphs can be extended to other classes of geometric graphs such as intersection graphs of squares and intersection graphs of rectangles whose height to width ratio is bounded by a constant. It still remains to show whether there exist deterministic online coloring algorithms that use the disk representation and have competitive ratio $o(\log \sigma)$. Our lower bound clearly fails in this case since a very simple online algorithm could use the information for the radii of the disks produced by the adversary ADV and color disks of radius α^i with colors 1 and 2 depending on whether i is even or odd.

Also, it would be interesting to investigate whether randomization helps in improving the known upper bounds and even beating the lower bounds for deterministic algorithms. To our knowledge, randomized online coloring algorithms have not been studied except for very general classes of graphs (e.g., in [16]) where the results are much weaker than those for disk graphs.

References

1. A. Borodin and R. El-Yaniv. Online computation and competitive analysis. *Cambridge University Press*, 1998.
2. I. Caragiannis, A. V. Fishkin, C. Kaklamanis, and E. Papaioannou. Online algorithms for disk graphs. In *Proc. of the 29th International Symposium on Mathematical Foundations of Computer Science (MFCS '04)*, LNCS 3153, Springer, pp. 215-226, 2004.
3. B.N. Clark, C.J. Colbourn, and D.S. Johnson. Unit disk graphs. *Discrete Mathematics*, 86, pp. 165-177, 1990.
4. T. Erlebach and J. Fiala. On-line coloring of geometric intersection graphs. *Computational Geometry: Theory and Applications*, 9(1-2), pp. 3-24, 2002.
5. T. Erlebach and J. Fiala. Independence and coloring problems on intersection graphs of disks. *Manuscript*, 2003.
6. J. Fiala, A.V. Fishkin, and F.V. Fomin. Off-line and on-line distance constrained labeling of graphs. In *Proc. of the 9th Annual European Symposium on Algorithms (ESA '01)*, LNCS 2161, pp. 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. Gräf, M. Stumpf, and G. Weissenfels. On coloring unit disk graphs. *Algorithmica*, 20(3), pp. 277-293, 1998.
9. A. Gyárfás and J. Lehel. On-line and first fit colorings of graphs. *Journal of Graph Theory*, 12(2), pp. 217-227, 1988.
10. D.K. Hale. Frequency assignment: theory and applications. In *Proc. of the IEEE*, 68(12), pp. 1497-1514, 1980.
11. P. Hliněný and J. Kratochvíl. Representing graphs by disks and balls. *Discrete Mathematics*, 229(1-3), pp. 101-124, 2001.
12. S. Irani. Coloring inductive graphs on-line. *Algorithmica*, 11, pp. 53-72, 1994.
13. E. Malesińska. Graph theoretical models for frequency assignment problems. *PhD Thesis*, Technical University of Berlin, 1997.
14. M.V. Marathe, H. Breu, H.B. Hunt III, S.S. Ravi, and D.J. Rosenkrantz. Simple hueristics for unit disk graphs. *Networks*, 25, pp. 59-68, 1995.
15. R. Peeters. On coloring j-unit sphere graphs. Technical report, Dept. of Economics, Tilburg University, 1991.
16. S. Vishwanathan. Randomized online coloring of graphs. In *Proc. of the 31st Annual Symposium on Foundations of Computer Science (FOCS '90)*, 1990.