Online Algorithms for Disk Graphs*

Ioannis Caragiannis[1], Aleksei Fishkin[2], and 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] Institut für Informatik und Praktische Mathematik
Universität zu Kiel, Christian-Albrechts Platz 4, 24098, Germany

Abstract. We study the on-line versions of two fundamental graph problems, maximum independent set and minimum coloring, for the case of *disk graphs* which are graphs resulting from intersections of disks on the plane. In particular, we investigate whether randomization can be used to break known lower bounds for deterministic on-line independent set algorithms and present new upper and lower bounds; we also present an improved upper bound for on-line coloring.

1 Introduction

We study two fundamental graph problems, maximum independent set and minimum coloring. Given a graph G, the maximum independent set problem is to find an independent set (i.e., a set of nodes without edges between them) of maximum size, while the minimum coloring problem is to find an assignment of colors (i.e, 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 graphs modelling intersections of disks in 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 strictly smaller than 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, maximum independent set and minimum coloring are important since they can model resource allocation problems in radio communication networks [11]. Consider a set of transmitters located in fixed positions within a

* This work was partially supported by the European Union under IST FET Project CRESCCO, IST FET Project DELIS and RTN ARACNE.

J. Fiala et al. (Eds.): MFCS 2004, LNCS 3153, pp. 215–226, 2004.

geographical region. Each transmitter uses a specific frequency 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, important engineering problems that have to be solved are the frequency assignment problem where the objective is to minimize the number of frequencies used all over the network, and the call admission problem where the objective is to find a maximum-sized set of transmitters which can use the same frequency. Assuming that all transmitters have circular range, the graph reflecting possible interference between pairs of transmitters is a disk graph. The frequency assignment and call admission problems are equivalent to minimum coloring and maximum independent set problems, respectively.

An instance of the maximum independent set or 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 [13].

The maximum independent set in disk graphs has been proved to be NP-hard even for unit disk graphs and even if the disk representation is given as part of the input [3]. A naive independent set algorithm is the algorithm First-Fit: starting from an empty set, it incrementally constructs an independent set by examining the nodes of the graph in an arbitrary order and including a node in the independent set only if none of its neighbors has been previously included. When applied to unit disk graphs, First-Fit has approximation ratio at most 5 and does not use the disk representation [18] (also implicit in [14]). In [18], a 3-approximation algorithm is obtained by computing a specific ordering of the nodes of a unit disk graph and running First-Fit according to this ordering. A similar idea leads to a 5-approximation algorithm in general disk graphs [18]. Furthermore, as it has been observed in [5], a $(2.5+\epsilon)$-approximation algorithm for unit disk graphs follows by a more general result presented in [12]. None of the algorithms above use the disk representation. Polynomial-time approximation schemes have been presented for both unit disk graphs [15,19] and general disk graphs [6,2] when the disk representation is given.

The minimum coloring problem has also been proved to be NP-hard in [3,9] even for unit disk graphs. Again, First-Fit algorithm can be used. 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. Algorithm First-Fit computes 5-approximate solutions in unit disk graphs [8,18]. By processing the nodes of the graph in a specific order, First-Fit computes 3-approximate solutions in unit disk graphs [9,18,20]. In general disk graphs, a smallest-degree-last version of First-Fit achieves an approximation ratio of 5 [8,17,18].

In the on-line versions of the problems, 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 on-line independent set algorithm decides either to accept the node by including it in the independent set or to reject it, while an on-line coloring algorithm decides which color to assign to the node. In each case, the decisions of the algorithm cannot change in the future. The performance of an on-line algorithm is measured in terms of its competitive ratio (or competitiveness). For on-line independent set algorithms, the competitive ratio is defined as the maximum over all possible sequences of disks of the ratio of the size of the maximum independent set over the size of the independent set computed by the algorithm. For on-line coloring algorithms, the competitive ratio 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.

First-Fit is essentially an on-line algorithm. For the independent set problem, it has competitive ratio 5 in unit disk graphs [14,18] and $O(\min\{n, \sigma^2\})$ in σ-bounded disk graphs with n nodes [5]. As it is observed in [5], First-Fit is optimal within the class of deterministic on-line algorithms.

The First-Fit coloring algorithm 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 [16,10]. 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 [8,18] 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 on-line coloring algorithm is presented in [7]. 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]. Better on-line coloring algorithms exist for σ-bounded disk graphs in the case where the disk representation is given. Most of them use First-Fit as a subroutine. The best competitiveness upper bound in this case is $O(\min\{\log n, \log \sigma\})$ [4].

In this paper, we study the on-line version of both problems. For the independent set problem, we investigate whether randomization helps in improving the competitiveness of on-line algorithms. For randomized on-line independent set algorithms, the competitive ratio is defined as the maximum over all possible sequences of disks of the ratio of the size of the maximum independent set over the expected size of the independent set computed by the algorithm. We assume that the sequences of disks are selected by oblivious adversaries, i.e., adversaries that have no knowledge of the random choices of the algorithms (but may know the probability distribution used by the algorithm for making random choices). This is a typical assumption usually made in the study of randomized on-line algorithms [1]. Somewhat surprisingly, we show that, in general, randomization does not help against oblivious adversaries even if the disk representation is given, i.e., we construct sequences of disks for which no (possibly randomized) on-line algorithm can be better than $\Omega(n)$-competitive. In the case that the disk representation is not given, we prove a lower bound of $\Omega(\min\{n, \sigma^2\})$ on the

Table 1. Summary of results for the on-line independent set problem. (*) indicates results in this paper.

		Deterministic alg.	Randomized algorithms	
	Disk repr.	Lower/upper bound	Lower bound	Upper bound
σ-DG	+	$\Theta(\min\{n, \sigma^2\})$	* $\Omega(\min\{n, \log\sigma\})$	* $O(\min\{n, \log\sigma\})$ * $O\left(\min\{n, \frac{1}{\epsilon}\log\sigma\log^{1+\epsilon}\log\sigma\}\right)$
σ-DG	–	$\Theta(\min\{n, \sigma^2\})$	* $\Omega(\min\{n, \sigma^2\})$	$O(\min\{n, \sigma^2\})$
UDG	+	5	* 2.5	* $\frac{8\sqrt{3}}{\pi} \approx 4.41$
UDG	–	5	* 3	5

competitiveness of on-line algorithms on σ-bounded disk graphs with n nodes meaning that algorithm First-Fit is optimal within a small constant factor. For the case of σ-bounded disk graphs with given representation, we present randomized algorithms with competitive ratio almost logarithmic in σ and show that they are optimal. For unit disk graphs, we present a randomized algorithm with competitive ratio 4.41. We also show lower bounds of 2.5 and 3 for randomized algorithms in unit disk graphs. Our results for the on-line independent set problem together with the previously known results on deterministic on-line algorithms are summarized in Table 1. For the coloring problem, we show how to achieve the best known upper bound of $O(\min\{\log n, \log\sigma\})$ for σ-bounded sequences of n disks even if the disk representation is not given.

The rest of the paper is structured as follows. Section 2 is devoted to the on-line independent set problem in σ-bounded disk graphs. Our results for unit disk graphs are presented in Section 3 while our coloring algorithm is presented in Section 4. We conclude with extensions and open problems in Section 5. Due to lack of space, most of the proofs have been omitted. They will appear in the final version of the paper.

2 Independent Sets in σ-Bounded Disk Graphs

2.1 Upper Bounds

In this section we present the randomized on-line algorithm Classify for computing independent sets in disk graphs. It has a competitive ratio $O(\min\{n, \log\sigma\})$ against oblivious adversaries on σ-bounded disk graphs with n nodes. The algorithm uses the value of σ which is supposed to be known in advance and makes its random choices based on the disk representation. Despite these limitations, this is the first algorithm achieving a competitive ratio logarithmic in σ and (as we will prove in Section 2.2) is optimal among the on-line algorithms that use the disk representation.

Algorithm Classify works as follows. When the first disk is presented, the algorithm flips a coin. On heads, it accepts the disk and executes algorithm First-Fit for disks having radii in the interval $[R, 2R)$, where R is the radius of the first disk presented, ignoring (i.e., rejecting) all other disks. On tails, the algorithm selects equiprobably a number i from the set $\{-\lceil \log \sigma \rceil, -\lceil \log \sigma \rceil + 1, ..., -1, 1, ..., \lceil \log \sigma \rceil\}$ and executes algorithm First-Fit for disks of radius in the interval $[R2^i, R2^{i+1})$, ignoring (i.e., rejecting) all other disks.

We prove the following theorem.

Theorem 1. *Algorithm* Classify *is* $O(\min\{n, \log \sigma\})$*-competitive against oblivious adversaries on* σ*-bounded disk graphs with* n *nodes.*

Proof. Since the first disk is accepted with probability 1/2, the algorithm has competitive ratio $O(n)$. In what follows, we show that the algorithm is $O(\log \sigma)$-competitive as well. Denote by OPT the optimal independent set of the sequence. For $i = -\lceil \log \sigma \rceil, -\lceil \log \sigma \rceil - 1, ..., \lceil \log \sigma \rceil$, denote by S_i the set of disks with radius in the interval $[R2^i, R2^{i+1})$ and by OPT_i the maximum independent set among the disks belonging to set S_i. Clearly, $|OPT_i| \geq |OPT \cap S_i|$ since $OPT \cap S_i$ is an independent set for S_i. Assume that the algorithm selects set S_i and executes algorithm First-Fit on the disks of that set. Observe that disks in S_i form a 2-bounded disk graph. In such graphs, the following lemma gives a guarantee on the performance of algorithm First-Fit for computing independent sets.

Lemma 1. *Algorithm* First-Fit *is at most* 15*-competitive on* 2*-bounded disk graphs.*

Using the lemma, we obtain that the algorithm accepts at least

$$B_i \geq \frac{1}{15}|OPT_i| \geq \frac{1}{15}|OPT \cap S_i|$$

disks of S_i. Now, the expected size of the independent set computed by algorithm Classify is

$$\begin{aligned} E[B] &= \sum_{i=-\lceil \log \sigma \rceil}^{\lceil \log \sigma \rceil} (\Pr[S_i \text{ is selected}] \cdot B_i) \\ &\geq \frac{1}{15} \sum_{i=-\lceil \log \sigma \rceil}^{\lceil \log \sigma \rceil} (\Pr[S_i \text{ is selected}] \cdot |OPT \cap S_i|) \\ &\geq \frac{1}{15} \min_i \{\Pr[S_i \text{ is selected}]\} \cdot |OPT| \\ &\geq \frac{1}{60\lceil \log \sigma \rceil} \cdot |OPT|. \end{aligned}$$

Hence, the competitive ratio of the algorithm is $O(\log \sigma)$. □

We now present algorithm Guess which achieves a slightly weaker competitive ratio but does not need to known neither n nor σ in advance. Consider a sequence of n disks and let R be the radius of the first disk of the sequence. Then, for any $i = 0, 1, ..., 2\lceil \log \sigma \rceil - 1$, define the set of disks S_i with radii at least $R/2^{\lceil \log \sigma \rceil - i}$ and smaller than $R/2^{\lceil \log \sigma \rceil - i - 1}$. When the first disk of each set is presented, the algorithm probabilistically determines whether it will consider disks from that specific set and ignore all disks from all other sets. Sets are divided into epochs. Epoch 0 consists of the set containing the first disk of the sequence. For $j = 1, 2, ...,$ epoch j consists of the 2^j sets presented after the sets of epoch $j-1$.

When the first disk appears, the algorithm tosses a coin with $\Pr[HEADS] = 1/2$. On heads, the algorithm decides to run First-Fit on the disks belonging to the first set of epoch 0 and ignores (i.e., rejects) all other disks; on tails, it decides to reject all disks belonging to the first set of epoch 0. When the first disk of the first set of the j-th epoch (for $j > 0$) is presented, if the algorithm has rejected all sets in epochs $0, ..., j-1$, it tosses a coin with $\Pr[HEADS] = \frac{\epsilon}{i+2}$. On HEADS, it equiprobably selects one of the 2^j sets of the epoch and decides to run First-Fit on the disks belonging to that set and ignores (i.e., rejects) disks from all other sets; on tails, it rejects all disks from all sets of the epoch.

Theorem 2. *Algorithm Guess is at most $O\left(\min\left\{n, \frac{1}{\epsilon}\log\sigma\log^{1+\epsilon}\log\sigma\right\}\right)$-competitive, for any $\epsilon \in (0,1]$, against oblivious adversaries in σ-bounded disk graphs with n nodes.*

2.2 Lower Bounds

The lower bounds presented in this section show that, in general, randomization does not help, i.e., there are sequences of n disks for which any on-line algorithm is $\Omega(n)$-competitive even if the disk representation is given. For σ-bounded disk graphs, the next lower bound states that when the disk representation is not given, on-line algorithms with competitive ratio logarithmic in σ do not exist.

Theorem 3. *Any randomized on-line algorithm for computing independent sets in σ-bounded disk graphs with n nodes is $\Omega(\min\{n, \sigma^2\})$-competitive against oblivious adversaries, if the disk representation is not given.*

Proof. Let κ be a positive integer. We will construct an adversary which generates a graph G_κ with an independent set of size $\kappa+1$ such that the expectation of the size of the independent set of G_κ that any randomized on-line algorithm can find is at most 2.

The graph G_κ generated by the adversary is defined as follows. The nodes of G_κ are partitioned into κ levels $0, 1, ..., \kappa-1$. Each level i has two nodes: a left node v_l^i and a right node v_r^i. The two nodes of a level are non-adjacent. First, the adversary generates the two nodes of level 0. For $i = 1, ..., \kappa-1$, the nodes of level i are generated after the nodes of level $i-1$. The adversary tosses a coin in order to connect the nodes of level i with nodes of smaller levels. On heads, it connects both nodes of level i to node v_l^{i-1} and to all nodes of levels $i-2, i-3, ..., 0$ to which node v_l^{i-1} is connected; on tails, it connects both nodes of level i to node v_r^{i-1} and to all nodes of levels $i-2, ..., 0$ to which node v_r^{i-1} is connected.

Consider the set of nodes consisting of the two nodes of level $\kappa-1$ and, for $i = 0, ..., \kappa-2$, of the node of level i which is not connected to nodes of higher

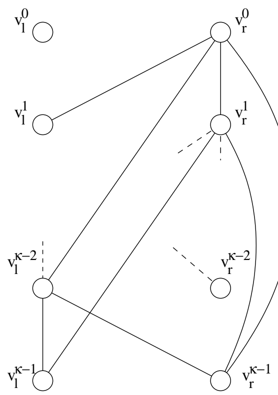

Fig. 1. An example of graph G_κ.

levels. This is an independent set of G_κ. Hence, the optimal independent set of G_κ has size at least $\kappa + 1$.

In what follows we will show that the size of the independent set of G_κ any (possibly randomized) on-line algorithm can compute is at most 2. Consider the application of an algorithm A on t sequences of disks produced by the adversary. Denote by l_i the number of executions in which the algorithm accepts the left node of level i, by r_i the number of executions in which the algorithm accepts the right node of level i, and by b_i the number of executions in which the algorithm accepts both nodes of level i.

For $i = 0, 1, ..., \kappa - 1$, let X_i be the random variable denoting the number of executions in which the nodes presented at level i are unconstrained by nodes of smaller levels (i.e., they are not connected to nodes of smaller levels that have been accepted by the algorithm). Then, nodes of level $i+1$ are constrained only if: (i) the left node of level i is rejected and the nodes of level $i+1$ are connected to the left node of level i or (ii) the right node of level i is rejected and the nodes of level $i+1$ are connected to the right node of level i. Hence,

$$E[X_{i+1}|X_i] \leq X_i - \frac{r_i + l_i}{2} - b_i$$

and

$$E[X_{\kappa-1}] \leq t - \sum_{i=0}^{\kappa-2} \frac{r_i + l_i + 2b_i}{2} \Rightarrow$$

$$\sum_{i=0}^{\kappa-2} (r_i + l_i + 2b_i) \leq 2t - 2E[X_{\kappa-1}] \tag{1}$$

Since the number of executions in which the algorithm accepts at least one node from level $\kappa - 1$ is at most the number of executions in which the nodes of level $\kappa - 1$ are unconstrained, it is $r_{\kappa-1} + l_{\kappa-1} + b_{\kappa-1} \leq E[X_{\kappa-1}]$. Now, using (1), we obtain that the expectation of the size of the independent set $B(G_\kappa)$ computed by the algorithm is

$$\begin{aligned} E[|B(G_\kappa)|] &= \frac{1}{t}\sum_{i=0}^{\kappa-1}(r_i + l_i + 2b_i) \\ &= \frac{1}{t}\sum_{i=0}^{\kappa-2}(r_i + l_i + 2b_i) + \frac{1}{t}(r_{\kappa-1} + l_{\kappa-1} + 2b_{\kappa-1}) \\ &\le 2 - \frac{1}{t}(2E[X_{\kappa-1}] - (r_{\kappa-1} + l_{\kappa-1} + 2b_{\kappa-1})) \\ &\le 2. \end{aligned}$$

We conclude that the competitive ratio of the algorithm is at least $\frac{\kappa+1}{2}$.

It remains to show that graph G_κ for $\kappa = \Omega\left(\min\{n, \sigma^2\}\right)$ is a σ-bounded disk graph. This is stated in the following lemma.

Lemma 2. *For any $\sigma \ge 2$, graph G_{4d^2} for $d = \lfloor \frac{\sigma+2}{4} \rfloor$ is a σ-bounded disk graph.*

The proof of the lemma will appear in the final version of the paper. This completes the proof of the theorem. □

The next lower bound states that algorithm Classify is optimal.

Theorem 4. *Any randomized on-line algorithm for computing independent sets in σ-bounded disk graphs with n nodes is $\Omega(\min\{n, \log\sigma\})$-competitive against oblivious adversaries.*

3 Independent Sets in Unit Disk Graphs

In this section, we present new upper and lower bounds on the competitiveness of on-line randomized independent set algorithms for unit disk graphs.

We first present algorithm Filter, an on-line randomized algorithm for computing independent sets in unit disk graphs. We show that the algorithm is $\frac{8\sqrt{3}}{\pi} \approx 4.41$-competitive against oblivious adversaries.

At the beginning, algorithm Filter selects α and β uniformly at random from the intervals $[0, 4)$ and $[0, 2\sqrt{3})$, respectively. When a new disk centered at point (x, y) appears, the algorithm does the following: If there are integers κ, λ such that the point $(x + \alpha, y + \beta)$ has distance less than 1 from the point with coordinates $(4\kappa + 2(\lambda \bmod 2), 2\lambda\sqrt{3})$, then Filter executes algorithm First-Fit, else it ignores the disk.

Theorem 5. *Algorithm* Filter *is $\frac{8\sqrt{3}}{\pi}$-competitive against oblivious adversaries.*

Proof. Consider the application of algorithm Filter on a sequence $\mathcal{D}$ of disks of unit radius. Let $\mathcal{D}'$ denote the (random) subsequence of $\mathcal{D}$ consisting of the disks not ignored by the algorithm. We denote by $A(\mathcal{D})$ the maximum independent set of a sequence $\mathcal{D}$ and by $B(\mathcal{D})$ the set of disks accepted by the algorithm.

We first show that the probability that a disk is not ignored by the algorithm is $\frac{\pi}{8\sqrt{3}}$. Consider a disk D with center at point (x, y) and the rectangle defined by the diagonal points (x, y) and $(x + 4, y + 2\sqrt{3})$. Also, consider the unit disks containing the points at distance less than 1 from points with coordinates $(4\kappa + 2(\lambda \bmod 2), 2\lambda\sqrt{3})$ for integer κ and λ, and observe that the total area of the intersection of these disks with the rectangle equals the area of a disk with radius 1 (see Figure 2. Since point $(x + \alpha, y + \beta)$ is uniformly distributed within the rectangle, the probability that the disk D is not ignored by algorithm Filter is equal to the area of a disk of radius 1 over the area of the rectangle, i.e., $\frac{\pi}{8\sqrt{3}}$.

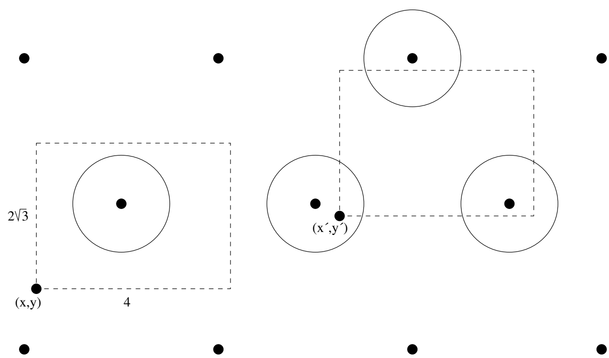

Fig. 2. Disk centers (x, y) and (x', y') and the rectangles where points $(x + \alpha, y + \beta)$ and $(x' + \alpha, y' + \beta)$ are uniformly distributed.

Now consider the maximum independent set $A(\mathcal{D})$ and let $A'(\mathcal{D})$ be the (random) subset of $A(\mathcal{D})$ consisting of the disks of $A(\mathcal{D})$ not ignored by algorithm Filter. Clearly, $A'(\mathcal{D})$ is an independent set for the set of disks $\mathcal{D}'$, thus, it is $|A(\mathcal{D}')| \geq |A'(\mathcal{D})|$. By linearity of expectation, we obtain that $|A'(\mathcal{D})| = \frac{\pi}{8\sqrt{3}}|A(\mathcal{D})|$ meaning that

$$E[|A(\mathcal{D}')|] \geq \frac{\pi}{8\sqrt{3}}|A(\mathcal{D})| \quad (2)$$

We now observe that each connected component of the intersection graph defined by the disks in $\mathcal{D}'$ is a clique. In particular, consider the two points O_1 with coordinates $(4\kappa_1 + 2(\lambda_1 \bmod 2), 2\lambda_1\sqrt{3})$ and O_2 with coordinates $(4\kappa_2 + 2(\lambda_2 \bmod 2), 2\lambda_2\sqrt{3})$ such that either $\kappa_1 \neq \kappa_2$ or $\lambda_1 \neq \lambda_2$ and three disks D_1, D_2, and D_3 centered at points C_1, C_2, and C_3 with coordinates (x_1, y_1), (x_2, y_2), and (x_3, y_3), respectively. Also, denote by C'_1, C'_2, and C'_3 the points with coordinates $(x_1 + \alpha, y_1 + \beta)$, $(x_2 + \alpha, y_2 + \beta)$, and $(x_3 + \alpha, y_3 + \beta)$, respectively. Assume that points C_1 and C_3 have distance smaller than 1 from point O_1, and point C_2 has distance smaller than 1 from point O_2. We will show that disks D_1 and D_3 overlap while disks D_1 and D_2 are non-overlapping. Clearly, it is $|C_1C_3| = |C'_1C'_3|$ and by triangle inequality, we obtain that $|C_1C_3| \leq |C'_1O_1| + |O_1C'_2| < 2$.

Hence, disks D_1 and D_3 overlap. Now, it can be easily verified that if either $\kappa_1 \neq \kappa_2$ or $\lambda_1 \neq \lambda_2$, it is $|O_1O_2| \geq 4$. By the triangle inequality, we have that $|O_1C_1'| + |C_1'C_2'| + |C_2'O_2| \geq 4$. Clearly, $|C_1'C_2'| = |C_1C_2|$ and since $|O_1C_1'| < 1$ and $|C_2'O_2| < 1$, it is also $|C_1C_2| > 2$ meaning that disks D_1 and D_2 do not overlap.

Now, since each connected component of the intersection graph of $\mathcal{D}'$ is a clique, the maximum independent set in the neighborhood of a disk has size at most 1. So, any disk accepted by algorithm Filter may block at most one disk in $A(\mathcal{D}')$. Hence, for the subsequence $\mathcal{D}'$ of the disks not ignored by algorithm Filter, it is $B(\mathcal{D}') \geq |A(\mathcal{D}')|$ implying that $E[B(\mathcal{D}')] \geq E[|A(\mathcal{D}')|]$. Using (2), we obtain that the competitive ratio of algorithm Filter is

$$\frac{|A(\mathcal{D})|}{E[B(\mathcal{D}')]} \leq \frac{8\sqrt{3}}{\pi}.$$

□

By adapting the lower bound construction of Section 2.2 to the case of unit disk graphs, we obtain the following statement.

Theorem 6. *No on-line (randomized) algorithm for computing independent sets in unit disk graphs can be better than* 3*-competitive against oblivious adversaries if the disk representation is not given. Even if the disk representation is given, then no on-line (randomized) algorithm can be better than* 2.5*-competitive against oblivious adversaries.*

4 An Upper Bound for Online Coloring

In this section we present an on-line coloring algorithm for disk graphs which does not require the disk representation. It achieves competitive ratio $O(\min\{\log n, \log\sigma\})$ for coloring σ-bounded sequences of n disks matching the best known upper bound for the case where the disk representation is given. The algorithm is a combination of algorithm First-Fit and algorithm Layered which is presented in the following.

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. We can show the following.

Lemma 3. *For any σ-bounded sequence of disks, the number of layers constructed by algorithm* Layered *is at most $1 + \log\sigma$.*

Theorem 7. *The algorithm* Layered *is $O(\log\sigma)$-competitive when applied to σ-bounded sequences of disks.*

Proof. Consider the application of algorithm Layered on a σ-bounded sequence of disks. Let j be the layer where the maximum number of colors has been used. Let α be the highest color used in layer j and let D_j be the disk colored with color α. Then, this disk overlaps with at least $\alpha - 1$ disks of layer j appeared prior to it. By the definition of the algorithm, the number of mutually non-overlapping disks of layer j overlapping with D_j is at most 15. This implies that the optimal algorithm should use at least $\alpha/15$ colors for coloring the disks of the sequence while, by Lemma 3, algorithm Layered uses at most $\alpha(\log \sigma + 1)$ colors, hence, it is at most $15(\log \sigma + 1)$-competitive. □

We now combine algorithms First-Fit and Layered using a technique proposed in [4,5] to obtain a better result. We use two separate sets of colors for algorithms First-Fit and Layered. When a new disk D_i is presented we run algorithm First-Fit on D_i together with those disks colored by First-Fit. Similarly, we execute Layered. Then, we compare the results of these two algorithms and color D_i with the algorithm that has used fewer colors up to that point (including the color used for disk D_i). The total number of colors used is at most the sum of the number of colors used by both methods. Note that at any time of the execution of the combined algorithm, the number of colors used by First-Fit and the number of colors used by Layered differ by at most one. Assume that $n < \sigma$. The number of colors used by First-Fit is at most $O(\log n)$ times the optimal number of colors. The number of colors used by Layered is at most one more than that of First-Fit. So, the total number of colors is at most $O(\log n)$ times the optimal number of colors. A similar argument holds in the case where $n \geq \sigma$. We obtain the following theorem.

Theorem 8. *There exists an $O(\min\{\log n, \log \sigma\}$-competitive algorithm for on-line coloring a σ-bounded disk graph with n nodes.*

5 Extensions and Open Problems

The results for the independent set extend to the more general problem where we are given $w \geq 1$ colors and the objective is to accept the maximum number of disks which can be properly colored with at most w colors (clearly, for $w = 1$, this is the independent set problem). Algorithms Classify, Guess, and Filter can be easily modified to solve this problem with the same competitiveness bounds we proved for the independent set problem.

The most interesting open problem related to the independent set problem is perhaps to close the gap on the competitiveness of (randomized) on-line algorithms in unit disk graphs. It would be very interesting even to find an algorithm with competitive ratio smaller than 5 which does not require the disk representation. For the coloring problem, there is still a large gap (in terms of σ) between the competitiveness of algorithm Layered and the known lower bounds.

References

1. A. Borodin and R. El-Yaniv. Online computation and competitive analysis. *Cambridge University Press*, 1998.
2. T.M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *Journal of Algorithms*, 46, pp. 178-189, 2003.
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. T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric graphs. In *Proc. of the 12th Annual ACM/SIAM Symposium on Discrete Algorithms (SODA '01)*, pp. 671-679, 2001.
7. 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.
8. A. Gräf. Coloring and recognizing special graph classes. Musikinformatik und Medientechnik Bericht 20/95, Johannes Gutenberg, Universität Mainz, 1995.
9. A. Gräf, M. Stumpf, and G. Weissenfels. On coloring unit disk graphs. *Algorithmica*, 20(3), pp. 277-293, 1998.
10. A. Gyárfás and J. Lehel. On-line and first fit colorings of graphs. *Journal of Graph Theory*, 12(2), pp. 217-227, 1988.
11. D.K. Hale. Frequency assignment: theory and applications. In *Proc. of the IEEE*, 68(12), pp. 1497-1514, 1980.
12. M.M. Halldórsson. Approximating discrete collections via local improvements. In *Proc. of the 6th Annual ACM/SIAM Symposium on Discrete Algorithms (SODA '95)*, pp. 160-169, 1995.
13. P. Hliněný and J. Kratochvíl. Representing graphs by disks and balls. *Discrete Mathematics*, 229(1-3), pp. 101-124, 2001.
14. D. Hochbaum. Efficient bounds for the stable set, vertex cover and set packing problems. *Discrete Applied Mathematics*, 6, pp. 243-254, 1983.
15. 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), pp. 238-274, 1998.
16. S. Irani. Coloring inductive graphs on-line. *Algorithmica*, 11, pp. 53-72, 1994.
17. E. Malesińska. Graph theoretical models for frequency assignment problems. *PhD Thesis*, Technical University of Berlin, 1997.
18. 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.
19. T. Matsui. Approximation algorithms for maximum independent set problems and fractional coloring problems on unit disk graphs. In *Discrete and Computational Geometry*, LNCS 1763, Springer, pp. 194-200, 2000.
20. R. Peeters. On coloring j-unit sphere graphs. Technical report, Dept. of Economics, Tilburg University, 1991.