On-line and off-line distance constrained labeling of disk graphs (extended abstract)

Jiří Fiala[*1,2], Aleksei V. Fishkin [**1], and Fedor V. Fomin[***3]

[1] CAU, Institute of Computer Science and Applied Mathematics, Olshausenstr. 40, 24098, Kiel, Germany. {jfi,avf}@informatik.uni-kiel.de
[2] Charles University, Faculty of Mathematics and Physics, Institute for Theoretical Computer Science (ITI), Malostranské nám. 2/25, 118 00, Prague, Czech Republic.
[3] Faculty of Mathematics and Mechanics, St. Petersburg State University, Bibliotechnaya sq. 2, St. Petersburg, 198904, Russia. fomin@gamma.math.spbu.ru

Abstract. A disk graph is the intersection graph of a set of disks in the plane. We consider the problem of assigning labels to vertices of a disk graph satisfying a sequence of distance constrains. Our objective is to minimize the distance between the least and the largest labels. We propose an on-line labeling algorithm on disk graphs, if the maximum and minimum diameters are bounded. We give the upper and lower bounds on its competitive ratio, and show that the algorithm is asymptotically optimal. In more detail we explore the case of distance constraints $(2,1)$, and present two off-line approximation algorithms. The last one we call robust, i.e. it does not require the disks representation and either outputs a feasible labeling, or answers the input is not a unit disk graph.

1 Introduction

The *frequency assignment problem* arises in *radio* or *mobile telephony* networks when different radio *transmitters* interfere to each other while operating on the same or closely related frequency channels. The most common model for an instance of frequency assignment is the *interference graph*. Each vertex of the interference graph represents a transmitter. If simultaneous broadcasting of two transmitters may cause an interference, then they are connected in the interference graph by an edge. One can associate the coverage area of a transmitter with a disk of a particular diameter, and afterwards, to model the interference graph as a graph whose edges connect transmitters at "close" distance, e.g. when their scopes intersect. Then, the underlying interference graph is a *disk graph*, that

[*] Partially supported by EU ARACNE project HPRN-CT-1999-00112, by GAČR 201/99/0242 and by the Ministry of Education of the Czech Republic as project LN00A056.
[**] Partially supported by the DFG-Graduiertenkolleg "Effiziente Algorithmen und Mehrskalenmethoden".
[***] The work of this author was done while he was a visiting postdoc at DIMATIA-ITI partially supported by GAČR 201/99/0242 and by the Ministry of Education of the Czech Republic as project LN00A056.

is the intersection graph of disks in the plane, or, sometimes, a *unit disk graph*, when all disks are of the same diameter [10].

Regarding the assumption that a pair of "close" transmitters should be assigned different frequencies, the freqency assignment is equivalent to the problem of coloring the interference graph. However, in [10] it was observed that the signal propagation may affect the interference even in distant regions (but with decreasing intensity). Hence, not only "close" transmitters should get different frequencies, but also frequencies used at some distance should be appropriately separated. In this case, the frequency assignment can be modeled as the problem of *distance constrained labeling*, or so called $L_{(p_1,\ldots,p_k)}$*-labeling* of the interference graph [12]:

Definition 1. *Let* $p_1, \ldots, p_k$ *be a sequence of positive integers called* distance constraints. *The* $L_{(p_1,\ldots,p_k)}$*-labeling of a graph* G *is a mapping* c: $V(G) \to \{1, \ldots, \lambda\}$ *such that the following holds:*

$$\forall i : 1 \leq i \leq k, \forall u, v \in V(G) : dist_G(u, v) \leq i \Rightarrow |c(u) - c(v)| \geq p_i.$$

The minimum number for which an $L_{(p_1,\ldots,p_k)}$-labeling of G exists, is denoted by $\chi_{(p_1,\ldots,p_k)}(G)$. Notice that $\chi_{(1)}(G) = \chi(G)$, where $\chi(G)$ is the chromatic number of G. Also for $p_1 = p_2 = \cdots = p_k = 1$, $\chi_{(p_1,\ldots,p_k)}(G) = \chi(G^k)$, where G^k is the k-th power of G, i.e. a graph which arise from G by adding edges connecting vertices at distance at most k.

Related works on graph labeling problems. A special case $k = 2$ and the distance constraints $(p_1, p_2) = (2, 1)$ was intensively studied for different graph classes in [1, 3, 7, 8, 16]. The exact value of $\chi_{(2,1)}$ can be derived for cycles and paths, and there are polynomial algorithms which compute the value $\chi_{(2,1)}$ for trees and co-graphs [3]. The problem of recognizing graphs such that $\chi_{(2,1)} \leq \lambda$ is *NP*-complete for all fixed $\lambda \geq 4$ [6]. For planar graphs, the problem of deciding $\chi_{(2,1)} \leq 9$ was shown to be *NP*-complete in [1].

Related works on disk graphs. The disk graph and unit disk graph recognition problem is *NP*-hard [2, 11]. Hence, algorithms that require the corresponding disk graph representation are substantially weaker than those which work only with graphs.

In [4] there is shown that the 3-coloring of unit disk graph is *NP*-complete (even if the input is a set of unit disks). However, for unit disk graph coloring problem there is a 5-competitive on-line algorithm [13, 14], and in the case of given disks representation, there is a 3-approximation off-line algorithm [2, 14]. Furthermore, there is a 5-approximate off-line algorithm for the disk graph coloring problem, which uses the first-fit technique on disks that are ordered from those with the biggest diameter to the smallest one [13]. From another side, one cannot expect the existence of an on-line coloring algorithm of a constant competitive ratio in the case of general disk graphs, since there is no such one even for planar graphs [9], and every planar graph is also a disk graph.

Our results. For fixed distance constraints $(p_1, \ldots, p_k)$ and fixed diameter ratio σ we present a constant-competitive on-line labeling algorithm which require the

disk representation. This is the first result on-line labeling algorithm known to our knowledge. We prove that when the representation is not given or when the diameter ratio is not bounded then no on-line algorithm with fixed competitive ratio exists.

For our algorithm we present upper and lower bounds on its competitive ratio, and show that in the case of unit disk graphs the algorithm is asymptotically optimal.

Finally, we explore the case of distance constraints $(2,1)$, and present two off-line approximation algorithms. We present robust labeling algorithm, i.e. an algorithm that does not require the disk representation and either outputs a feasible labeling, or answers the input is not a unit disk graph.

The following table summarizes the known and new upper bounds on the competitive or performance ratio on coloring and labeling problems on unit disk graphs (UDG), on disk graphs with diameter ratio bounded by σ (DG_σ) and on general disk graphs (DG).

Execution	Off-line		On-line	
Repres.	+	−	+	−
Coloring:				
UDG	3 [14]	3[1]	5 [13, 14]	5[1]
DG_σ	5 (as DG)	5 (as DG)	Theorem 2 [∗]	YES [5]
DG	5 [13]	5[1]	NO [5]	NO [9]
$L_{(2,1)}$-labeling:				
UDG [2]	$12 \to 9$ [∗]	$10.6 \to 10$ [∗]	$25 \to 12.5$ [∗]	—
$L_{(p_1,\ldots p_k)}$-labeling:				
UDG	Th. 2^4, C. 3^3 [∗]	Corollary 4^3 [∗]	Theorem 2 [∗]	NO [∗]
DG_σ	Theorem 2^4 [∗]	?	Theorem 2 [∗]	NO [∗]
DG	?	?	NO [∗]	NO [∗]

Positive results are marked either by "YES" or by the appropriate statement. "NO" means that no algorithm with fixed performance/competitive ratio exists. The sign "?" marks an open problem. The results presented in this paper are highlighted by "[∗]". The list of further explanations follows:

[1] The algorithm working without the disk representation can be derived from those which uses the representaion as shown in Section 4.2.

[2] This rows shows the values of the upper bound in the worth case. Since our results are better for graphs with large cliques, we give also the limit of the upper bound as the clique size grows to infinity.

[3] Only for $k = 2, (p_1, p_2) = (2, 1)$.

[4] Every on-line algorithm can be executed off-line.

Throughout the extended abstract various proofs are omitted or postponed into appendix due to space restriction.

2 Preliminaries

For a set of geometric objects, the corresponding *intersection graph* is the undirected graph whose vertices are objects and an edge connects two vertices if the corresponding objects intersect. Let $\mathcal{E}$ be a 2-dimensional Euclidean plane with the coordinates x, y. Let $\mathcal{D} = \{D_1, ..., D_n\}$ be a set of n disks in $\mathcal{E}$, where each D_i is uniquely determined by its center in (x_i, y_i) and by the diameter $d_i \in \mathbb{R}_+$. The intersection graph $G_\mathcal{D}$ of the set $\mathcal{D}$ is called a *disk graph*, more formally $V(G_\mathcal{D}) = \mathcal{D}, E(G_\mathcal{D}) = \{(D_i, D_j) \in \binom{\mathcal{D}}{2} : D_i \cap D_j \neq \emptyset\}$. When all disks of $\mathcal{D}$ have unit diameter, i.e. $d_i = 1$ for all $D_i \in \mathcal{D}$, then $G_\mathcal{D}$ is called a *unit disk graph*. In both cases, $\mathcal{D}$ is called the *disk representation* of $G_\mathcal{D}$. The value $\sigma(\mathcal{D}) = \frac{\max d_i}{\min d_i}$ is called the *diameter ratio* of $\mathcal{D}$.

We say that an algorithm $\mathcal{A}$ is a *ρ-approximation off-line $L_{(p_1,...,p_k)}$-labeling algorithm* if for a given graph $G_\mathcal{D}$ it runs in polynomial time and outputs an $L_{(p_1,...,p_k)}$-labeling of G such that the maximum label that $\mathcal{A}$ uses is at most $\rho\chi_{(p_1,...,p_k)}(G_\mathcal{D})$. The value ρ is called the *approximation ratio* of $\mathcal{A}$.

We say that an algorithm $\mathcal{B}$ is an *on-line $L_{(p_1,...,p_k)}$-labeling algorithm* if it labels the vertices of a graph $G_\mathcal{D}$ in an externally determined sequence $D_1 \prec ... \prec D_n$. At the time t the algorithm $\mathcal{B}$ has to irrevocably assign a label to D_t, while it can only see the edges connecting vertices $D_1, ..., D_t$. Such an algorithm $\mathcal{B}$ is a *ρ-competitive on-line $L_{(p_1,...,p_k)}$-labeling algorithm* if for every graph $G_\mathcal{D}$ and any ordering $\prec$ on $V(G_\mathcal{D})$ it always outputs an $L_{(p_1,...,p_k)}$-labeling with the maximum label at most $\rho\chi_{(p_1,...,p_k)}(G_\mathcal{D})$. We call the constant ρ the *competitive ratio* of $\mathcal{B}$.

Our on-line algorithm is based on the following partition of the plane $\mathcal{E}$. We call a *hexagonal tiling* on $\mathcal{E}$ the partition of the plane into the set $\mathcal{C}$ of hexagonal cells of diameter one. Each cell $C_{i,j} \in \mathcal{C}$ is characterized by two integer coordinates i, j, and is a simplex delimited by the following lines:

$$\begin{aligned} 2i - j - 1 < &\quad \tfrac{2}{3}\sqrt{3}x &&\leq 2i - j + 1 \\ i + j - 1 < &\quad \tfrac{2}{3}(\sqrt{3}x + 3y) &&\leq i + j + 1 \\ -i + 2j - 1 < &\quad \tfrac{2}{3}(-\sqrt{3}x + 3y) &&\leq i + 2j + 1. \end{aligned}$$

Each point of the plane belongs to exactly one cell $C_{i,j}$. Observe also, that each cell contains exactly two adjacent corners of the bounding hexagon, and the distance between every two points inside the same cell is at most one (see Figure 3 in the Appendix).

We denote the *plane distance* of two points $p, p' \in \mathcal{E}$ by $dist_\mathcal{E}(p, p')$. Similarly the plane distance of cells $C_{i,j}$ and $C_{k,l}$ is defined as

$$dist_\mathcal{E}(C_{i,j}, C_{k,l}) = \inf\{dist_\mathcal{E}(p, p') : p \in C_{i,j}, p' \in C_{k,l}\}.$$

3 On-line distance constrained labeling

In this section we explore the on-line distance constrained labeling problem for disk graphs.

We start with the following observation. Let $(p_1, ..., p_k)$ be a sequence of distance constraints ($k \geq 2$ and $p_2 \geq 1$) and let $\mathcal{B}$ be an arbitrary on-line $L_{(p_1,...,p_k)}$-labeling algorithm for disk graphs. Let $\mathcal{D}$ be a set of mutually disjoint disks, i.e. $G_{\mathcal{D}}$ has no edges.

Consider the two following cases. First assume that the disk representation is not a part of the input. Secondly, $\mathcal{B}$ the disk representation is given as a part of the input, but the diameter ratio $\sigma(\mathcal{D})$ is not known during the execution of $\mathcal{B}$.

In both cases, the algorithm $\mathcal{B}$ labels all vertices of $G_{\mathcal{D}}$ by distinct labels. Otherwise any new vertex may create a path of length two between an arbitrary pair of already labeled vertices (in the second case by extending $\mathcal{D}$ by a disk of large diameter). The maximal label used by $\mathcal{B}$ on $G_{\mathcal{D}}$ is at least $|\mathcal{D}|$, although $\chi_{(p_1,...,p_k)}(G_{\mathcal{D}}) = 1$.

Due to the above arguments, we consider the case when the disk representation $\mathcal{D}$ of $G_{\mathcal{D}}$ is a part of the input and when the corresponding diameter ratio $\sigma(\mathcal{D})$ is known, i.e. it is bounded by a constant $\sigma \geq 1$.

First we introduce a special circular labeling of cells and show how our on-line labeling algorithm uses such labeling. Later we derive the upper bound on its competitive ratio and show that the algorithm is asymptotically optimal for the class of unit disk graphs with at least one edge. Finally, we illustrate the algorithm performance for distance constraints $(2, 1)$.

3.1 Circular labeling

Let $\mathcal{D}$ be a set of disks with the diameter ratio $\sigma(\mathcal{D})$ and $G_{\mathcal{D}}$ be the corresponding disk graph. We suppose that the coordinates of plane $\mathcal{E}$ are scaled such that the smallest disk has the unit diameter. Moreover, through this section we assume that the diameter ratio σ and distance constraints $(p_1, ..., p_k)$ are fixed, i.e. for every selection of these parameters we design a specific algorithm.

Let $\mathcal{C}$ be the set of hexagonal cells. We say that a mapping $\varphi : \mathcal{C} \to \{1, 2, ..., l\}$ is a *circular l-labeling of $\mathcal{C}$* (with respect to $(p_1, ..., p_k)$ and $\sigma \geq 1$) if the following condition

$$dist_{\mathcal{E}}(C, C') \leq i \cdot \sigma \Rightarrow \min\{|\varphi(C) - \varphi(C')|, l - |\varphi(C) - \varphi(C')|\} \geq p_i,$$

is satisfied for all cells $C, C' \in \mathcal{C}$ and all $i \in \{1, ..., k\}$.

An example of a circular 25-labeling (with respect to $(p_1, p_2) = (2, 1), \sigma = \frac{\sqrt{7}}{2}$) is depicted in Figure 1. In fact, the cells with equal labels are at plane distance at least $2\sqrt{3}$ and those pairs labeled by consecutive labels have plane distance at least $\frac{\sqrt{7}}{2} \doteq 1.32$.

For general case, the existence of such labelings is guaranteed by the following theorem.

Theorem 1. *For every k-tuple of distance constraints $(p_1, ..., p_k)$, every $\sigma \geq 1$ and*

$$l^* = 1 + 6\Big(2p_1 - 1 + \sum_{i=2}^{\lfloor \frac{4k\sigma+4}{3} \rfloor} i \cdot (2p_{\lceil \frac{3i-4}{4\sigma} \rceil} - 1)\Big),$$

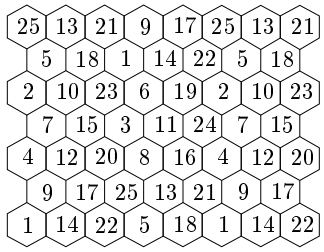

Fig. 1. An example of a 25-circular labeling ($(p_1, p_2) = (2, 1), \sigma = \frac{\sqrt{7}}{2}$)

there exists a circular l^*-labeling of $\mathcal{C}$ and can be found in $O(l^*k^4)$ time.

3.2 The algorithm

Assume now that the parameters of algorithm, i.e. distance constraints $(p_1, ..., p_k)$ and diameter ratio $\sigma \geq 1$ are fixed.

Algorithm On-line disks labeling (ODL)
Input: Set of disks $\mathcal{D}$, in an arbitrarily order $D_1 \prec ... \prec D_n$.
Output: An $L_{(p_1,...,p_k)}$-labeling $c : \mathcal{D} \to \mathbb{Z}^+$ of $G_{\mathcal{D}}$.

1. Find a circular l-labeling $\varphi : \mathcal{C} \to \{1, ..., l\}$
(with respect to $(p_1, ..., p_k)$ and σ).
2. For each k from 1 to n:
   2a. Decide into which cell $C_{i,j}$ the disk D_k belongs.
   2b. Define the label $c(D_k) = \varphi(C_{i,j}) + l \cdot |\mathcal{D}_{i,j}|$.
   2c. Set $\mathcal{D}_{i,j} = \mathcal{D}_{i,j} \cup \{D_k\}$.

The next lemma follows directly from the properties of circular l-labeling.

Lemma 1. *Suppose that at the first step the algorithm ODL finds a circular l-labeling (with respect to $(p_1, ..., p_k)$ and σ). Then for any set of disks $\mathcal{D}$ of diameter ratio $\sigma(\mathcal{D}) \leq \sigma$, the algorithm ODL produces a feasible $L_{(p_1,...,p_k)}$-labeling of $G_{\mathcal{D}}$ and the maximal label used by this algorithm is at most $l\omega(G_{\mathcal{D}})$.*

The analysis of our algorithms is based on the following fact: If a set of vertices X induces a complete subgraph of a graph G, then the maximal label used for vertices of X in any labeling is at least $p_1(|X| - 1) + 1$.

Lemma 2. *For any k-tuple $(p_1, ..., p_k)$ of distance constraints and any set of disks $\mathcal{D}$,*

$$\chi_{(p_1,...,p_k)}(G_{\mathcal{D}}) \geq p_1(\omega(G_{\mathcal{D}}) - 1) + 1 \geq p_1(\max_{i,j}\{|V(G_{\mathcal{D}_{i,j}})|\} - 1) + 1,$$

where $\omega(G_{\mathcal{D}})$ is the size of the maximum clique in $G_{\mathcal{D}}$.

Proof. Let $\mathcal{D}$ be a set of disks and $C_{i,j}$ a cell of $\mathcal{C}$. Let $\mathcal{D}_{i,j}$ be the set of disks of $\mathcal{D}$ with centers in $C_{i,j}$. Then, the corresponding subgraph $G_{\mathcal{D}_{i,j}}$ of $G_{\mathcal{D}}$ is isomorphic to a complete graph and the number of disks in $|\mathcal{D}_{i,j}|$ is at most $\omega(G_{\mathcal{D}})$. □

Combining Theorem 1 with Lemmas 1 and 2, we get the main result of this section.

Theorem 2. *For every $(p_1, ..., p_k)$ and every $\sigma \geq 1$ the competitive ratio of algorithm ODL is bounded by*

$$\rho \leq \min_{\mathcal{D}} \frac{\omega(G_{\mathcal{D}}) \cdot l^*}{(\omega(G_{\mathcal{D}}) - 1) \cdot p_1 + 1},$$

where the minimum is taken over all sets of disks of diameter ratio at most σ.

Therefore, ODL is an l^*-competitive on-line $L_{(p_1,...,p_k)}$-labeling algorithm for the class of disks graphs $G_{\mathcal{D}}$ of the diameter ratio $\sigma(\mathcal{D}) \leq \sigma$. Moreover, if disk graphs have at least one edge then the competitive ratio $\rho \leq \frac{2l^*}{p_1}$, and if $\omega(G_{\mathcal{D}}) \to \infty$ then $\rho \to \frac{l^*}{p_1}$.

If we apply the above theorem on the labeling depicted in Fig. 1, we get:

Corollary 1. *For the distance constraints $(2, 1)$, the algorithm ODL is $\frac{50}{3} \doteq 16.67$-competitive on disk graphs with at least one edge and $\sigma \leq 1.32$. Moreover, the competitive ratio tends to 12.5 when $\omega(G_{\mathcal{D}}) \to \infty$.*

Now we present a lower bound on the competitive ratio of any on-line algorithm solving the distance constrained labeling problem for disk graphs, which have a representation of bounded diameter ratio.

Theorem 3. *For any k-tuple $(p_1, ..., p_k)$ of distance constraints $(k \geq 2)$, any $\sigma \geq 1$ and any $\varepsilon > 0$, there is no $(\bar{\rho} - \varepsilon)$-competitive on-line $L_{(p_1,...,p_k)}$-labeling algorithm for the class of disk graphs with representation of diameter ratio at most σ, where*

$$\bar{\rho} = 1 + \frac{\sigma^2}{9} \max_{i=2,...,k} \{i^2 p_i\}.$$

Theorems 2 and 3 imply that:

Corollary 2. *For any $(p_1, ..., p_k)$, $k \geq 2$ and any $\sigma \geq 1$, the competitive ratio of algorithm ODL on disk graph (with given representation, diameter ratio at most σ and with at least one edge) is at most $O(\log k)$ times larger than the competitive ratio of* any *on-line $L_{(p_1,...,p_k)}$-labeling algorithm. Therefore, under these restricted conditions, the algorithm ODL is asymptotically optimal.*

As the last remark we propose a slightly different strategy with a smaller channel separation used for disks in the same location, though it does not lead to a better competitive ratio of the final algorithm. The idea is as follows. First scale the cells of $\mathcal{C}$ such that every cell has diameter σ. Then, find a feasible

circular l-labeling $\varphi : \bar{\mathcal{C}} \to \{1, ..., l\}$. After that, run the algorithm ODL on the set $\mathcal{D}$ of non-unit disks using the scaled cells. As the result, we find a feasible $L_{(p_1,...,p_k)}$-labeling of the disk graph $G_{\mathcal{D}}$, but the disks belonging to the same cell do not induce a complete graph. On the other hand, the area of any cell of diameter σ can be covered by at most $\sigma^2 + t\sigma$ unit cells, where t is a suitable constant that does not depend on σ. Then the number of vertices in any $G_{\mathcal{D}_{i,j}}$ is at most $\sigma^2 + t\sigma$ times more than the maximum clique size $\omega(G_{\mathcal{D}})$. Thus, it follows that

$$\chi_{(p_1,...,p_k)}(G_{\mathcal{D}_{i,j}}) \geq p_1(\omega(G_{\mathcal{D}_{i,j}}) - 1) \geq p_1\left(\left\lceil \frac{V(G_{\mathcal{D}_{i,j}})}{\sigma^2 + t\sigma} \right\rceil - 1\right)$$

and we use the same arguments as above.

4 Off-line labeling with constraints (2,1)

In this section we explore the off-line distance labeling problem only for one particular selection of distance constrains $(p_1, p_2) = (2, 1)$. First we discuss the case when the graph and its disc representation are known. Then we present a robust algorithm which do not require the disc representation.

4.1 The approximation algorithm

We first discuss the case that the entire graph $G_{\mathcal{D}}$ is given together with the representation $\mathcal{D}$ as the part of the input. The main idea of the off-line approximation algorithm is rather simple: We "cut" the plane into strips of small width, label unit disk graphs induced by strips and combine the labelings of strip graphs.

An unit disk graph $G_{\mathcal{D}}$ is a $\frac{1}{\sqrt{2}}$-*strip graph* if there is a disk representation $\mathcal{D}$ of $G_{\mathcal{D}}$ such that the centers of disks in $\mathcal{D}$ are in a strip of width $\frac{1}{\sqrt{2}}$. In other words, there is a mapping $f\colon V(G_{\mathcal{D}}) \to \mathbb{R} \times [0, \frac{1}{\sqrt{2}}]$ such that $(u, v) \in E(G_{\mathcal{D}})$ if and only if $dist_{\mathcal{E}}(f(u), f(v)) \leq 1$.

Lemma 3. *Let G be a $\frac{1}{\sqrt{2}}$-strip graph and let v be a vertex such that the unit disk corresponding to v (in some representation $\mathcal{D}$) has the smallest x-coordinate. Then the cardinality of the vertex set*

$$N_{G^2}(v) = \{u \in V(G) - \{v\}\colon dist_G(u, v) \leq 2\}$$

is at most $3\omega(G) - 1$.

A vertex ordering $v_1 \prec ... \prec v_n$ of a $\frac{1}{\sqrt{2}}$-strip graph G is *increasing* if there is a disk representation $\mathcal{D}$ such that $i < j$ if and only if x-coordinate of a disk in $\mathcal{D}$ corresponding to vertex v_i is at most x-coordinate of a disk corresponding to v_j.

By Lemma 3, the first-fit coloring algorithm that process vertices of the second power G^2 of a $\frac{1}{\sqrt{2}}$-strip graph G in an increasing order, uses at most distinct $3\omega(G)$ colors. This coloring of G^2 is equivalent to an $L_{(1,1)}$-labeling of G. By multiplying all labels of $L_{(1,1)}$-labeling by 2, we obtain an $L_{(2,2)}$-labeling of G. Finally,

$$\chi_{(2,1)}(G) \leq \chi_{(2,2)}(G) \leq 2\chi_{(1,1)}(G) = 2\chi(G^2) \leq 6\omega(G).$$

Using the first-fit approach, the $L_{(2,1)}$-labeling of G can be obtained in $O(nm)$ time ($n = |V(G)|, m = |E(G)|$). Notice, that all labels used for this labeling are *even*.

Now we are ready to describe the approximate labeling algorithm for unit disk graphs and constraints $(2,1)$. Without loss of generality we may assume that G is connected and has at most two vertices (i.e. $\omega(G) \geq 2$).

Algorithm Distance Labeling (DL)
Input: Unit disk representation of a unit disk graph G.
Output: Labeling of c of $V(G)$.

1. Partition a plane into $k = O(n)$ strips $S_1, ..., S_k$ of width $\frac{1}{\sqrt{2}}$, numbered from top to bottom, such that S_1 contains a disk with maximal y-coordinate and S_k the one with minimal. This partition induces partition of G into $\frac{1}{\sqrt{2}}$-strip graphs $G_1, ..., G_k$. (In a case of disks with centers in two strips ties are broken arbitrarily.)
2. For each $i \in \{1, ..., k\}$ find an $L_{(2,1)}$-labeling of G_i using even labels and with the maximum label bounded by $6\omega(G_i) \leq 6\omega(G)$.
3. Change the labels of graph G_i by increasing them by the number $\sharp_{(i \bmod 6)}$, where

$$(\sharp_0, ..., \sharp_6) = (0,\ 6\omega(G),\ 12\omega(G),\ -1,\ 6\omega(G) - 1,\ 12\omega(G) - 1)$$

Theorem 4. *For any unit disk graph G, algorithm DL produces an $L_{(2,1)}$-labeling and the maximal label used by this algorithm is at most $18\omega(G)$.*

Corollary 3. *The approximation ratio of the algorithm DL is bounded by 12 and tends to 9 as $\omega(G)$ grows to infinity.*

Observe that $\frac{1}{\sqrt{2}}$-strips were used in the description of the algorithm to simplify the explanation. To avoid irrationality, $\frac{1}{\sqrt{2}}$-strips in the algorithm can be replaced by c-strips, where c is any rational number between $\frac{2}{3}$ and $\frac{1}{\sqrt{2}}$. Also the algorithm can be generalized easily to an algorithm producing an $L_{(p,1)}$-labeling ($p \geq 1$) with the maximal label used at most $9p\omega(G)$.

4.2 Robust approximation algorithm

The main purpose of this section is to present the approximation labeling algorithm which doesn't need a geometric representation of a unit disk graph as part of input. (Let us remind that it is NP-hard to recognize unit disk graphs.)

In [15] the following notion of *robust algorithms* is discussed: an algorithm which solves an optimization problem on class is called robust if it satisfies the following conditions.

1. Whenever the input is in class C, the algorithm finds the correct solution.
2. If the input is not in the class C, then the algorithm either finds the correct solution, or answers that the input is not in the class.

Based on the ideas of [4], a robust algorithm computing the maximal clique of a unit disk graph is given in [15].

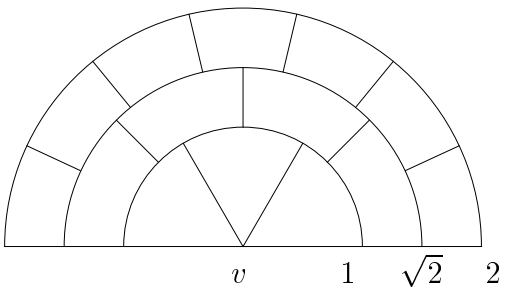

Fig. 2. The plane partition around v

Lemma 4. *Every unit disk graph G has a vertex v such that the set*

$$N_G(v) = \{u \neq v \colon \{u, v\} \in E(G)\}$$

has cardinality at most $3\omega(G) - 3$ and the set

$$N_{G^2}(v) - N_G(v)$$

contains at most $11\omega(G)$ vertices.

Proof. The proof of the lemma is based on the plane partition around vertex v shown in Fig. 2. □

We say that a vertex ordering $v_1 \prec ... \prec v_n$ of G is *good* if for every $2 \leq i \leq n$

(i) $|N_G(v_i) \cap \{v_1, ..., v_{i-1}\}| \leq 3\omega(G) - 3$;
(ii) $|(N_{G^2}(v_i) - N_G(v_i)) \cap \{v_1, ..., v_{i-1}\}| \leq 11\omega(G)$.

Notice, that by Lemma 4 every unit disk graph has a good vertex ordering. Also for a graph G with n vertices one can in $O(n^3)$ time either find a good vertex ordering, or conclude that there is no good ordering for G.

Now we are ready to describe algorithm RDL which can be regarded as robust distance labeling approximation algorithm. It doesn't use the geometric

representation of a graph G and either concludes that G is not unit disk graph, either (we prove it below) it finds an $L_{(2,1)}$-labeling of G with the maximum label at most $\leq 20\omega(G) - 8$.

Algorithm Robust Distance Labeling (RDL)
Input: Graph G in adjacency list.
Output: Either $L_{(2,1)}$-labeling c of $V(G)$, either conclusion that G is not unit disk graph.

1. Run the robust algorithm to compute $\omega(G)$. This algorithm either computes $\omega(G)$, either concludes that G is not unit disk graph.
2. Find a good vertex ordering $v_1 \prec ... \prec v_n$. If there is no such an ordering, then conclude that G is not unit disk graph.
3. Label vertices sequentially in order $\prec$ as follows:
   3a. Assume that vertices $v_1, ..., v_{i-1}$ are already labeled.
   3b. Let $\lambda \geq 1$ be the smallest integer which is not used as a label of vertices in $N_{G^2}(v_i) \cap \{v_1, ..., v_{i-1}\}$ nor is a member of the set $\bigcup_{j \in \{1,...,i-1\}: v_j \in N_G(v_i)} \{c(v_j) - 1, c(v_j), c(v_j) + 1\}$.
   3c. Label v_i by $c(v_i) = \lambda$.

Theorem 5. *For any graph G, Algorithm RDL either produces an $L_{2,1}$-labeling with maximum label $\leq 20\omega(G) - 8$, either concludes that G is not unit disk graph.*

Proof. Suppose that algorithm output that G is not unit disk graph. If it was done after the first step, then G has no edge ordering $\prec_e$ and therefore is not unit disk graph. If the algorithm halts at the second step, then its conclusion is verified by Lemma 4.

Suppose that RDL outputs a labeling. Let us first show that the maximum label used by the algorithm is $\leq 20\omega(G) - 8$. We proceed by induction. The vertex v_1 is labeled by 1, since both sets declared in 3b are empty. Suppose that we have labeled vertices $v_1, ..., v_{i-1}$. When we assign a label to v_i, then by (i) it has at most $3\omega(G) - 3$ labeled neighbors and at most $9\omega(G) - 9$ labels are unavailable because of these neighbors. (If one neighbor of v_i has label x then labels $x - 1, x$ and $x + 1$ are forbidden for v_i.) By (ii), at most $11\omega(G)$ labeled vertices are at distance two from v_i. Therefore, the number of unavailable labels is at most $20\omega(G) - 9$. Since we have $20\omega(G) - 8$ labels, $c(v_i) \leq 20\omega(G) - 8$. □

Again, this robust algorithm can be generalized for $L_{(p,1)}$-labeling ($p \geq 1$).

Corollary 4. *The approximation ratio of the algorithm RDL is bounded by $\frac{32}{3} \doteq 10,67$ and tends to 10 as $\omega(G)$ grows to infinity.*

Finally, we observe that for the construction of on-line algorithms with constant competitive ratio on unit disk graphs the knowledge of geometrical representation is crucial. For the coloring problem similar robust algorithm on unit disk graphs can be turned into on-line coloring algorithm (with worse competitive ratio). However, this is not the case for the labeling problem. The main

reason why RDL cannot be turned into "first-fit" algorithm is that at the moment when we have to select a suitable label for vertex v_i we need information about all vertices from the set $\{v_1, ..., v_{i-1}\}$ that are at distance two from v_i in G. Unfortunately this information cannot be fully derived from G restricted onto $\{v_1, ..., v_i\}$, moreover any robust on-line algorithm should use at least $|V(G)|-1$ labels, even if G is an independent set: Every pair of vertices, except maybe the last one, should get different labels because there is always a possibility that the last vertex will be adjacent to both of them and they will be at distance two.

References

1. H. L. Bodlaender, T. Kloks, R. B. Tan, and J. van Leeuwen, *λ-coloring of graphs*, in Proceedings of the 17th Annual Symp. on Theoretical Aspects of Computer Science (STACS 2000), H. Reichel and S. Tison, eds., Springer Verlag, Lecture Notes in Computer Science, vol. 1770, 2000, pp. 395–406.
2. H. Breu and D. G. Kirkpatrick, *Unit disc graph recognition is NP-hard*, Computational Geometry: Theory and Applications, 9 (1998), pp. 3–24.
3. G. J. Chang and D. Kuo, *The $L(2,1)$-labeling problem on graphs*, SIAM Journal of Discrete Mathematics, 9 (1996), pp. 309–316.
4. B. N. Clark, C. J. Colbourn, and D. S. Johnson, *Unit disk graphs.*, Discrete Math., 86 (1990), pp. 165–177.
5. personal communication to Thomas Erlebach.
6. J. Fiala, J. Kratochvíl, and T. Kloks, *Fixed-parameter tractability of λ-colorings*, in Graph-Theoretical Concepts in Computer Science, 25th WG '99, Ascona, no. 1665 in Lecture Notes in Computer Science, Springer Verlag, 1999, pp. 350–363.
7. J. P. Georges and D. W. Mauro, *On the size of graphs labeled with a condition at distance two*, Journal of Graph Theory, 22 (1996), pp. 47–57.
8. J. R. Griggs and R. K. Yeh, *Labelling graphs with a condition at distance 2*, SIAM Journal of Discrete Mathematics, 5 (1992), pp. 586–595.
9. A. Gyárfás and J. Lehel, *On-line and first fit colourings of graphs*, Jornal of Graph Theory, 12 (1988), pp. 217–227.
10. W. K. Hale, *Frequency assignment: Theory and applications*, Proc. of the IEEE, 68 (1980), pp. 1497–1514.
11. P. Hliněný and J. Kratochvíl, *Representing graphs by disks and balls.* to appear in Discrete Math.
12. R. A. Leese, *Radio spectrum: a raw material for the telecommunications industry.* 10th Conference of the European Consortium for Mathematics in Industry, Goteborg, 1998.
13. E. Malesińska, *Graph theoretical models for frequency assignment problems*, PhD thesis, Technical University of Berlin, 1997.
14. R. Peeters, *On coloring j-unit sphere graphs*, tech. rep., Dept. of Economics, Tilburg University, 1991.
15. V. Raghavan and J. Spinrad, *Robust algorithms for restricted domains* manuscript submitted to Special Issue of Journal of Algorithms.
16. J. van den Heuvel, R. A. Leese, and M. A. Shepherd, *Graph labeling and radio channel assignment*, Journal of Graph Theory, 29 (1998), pp. 263–283.

Appendix

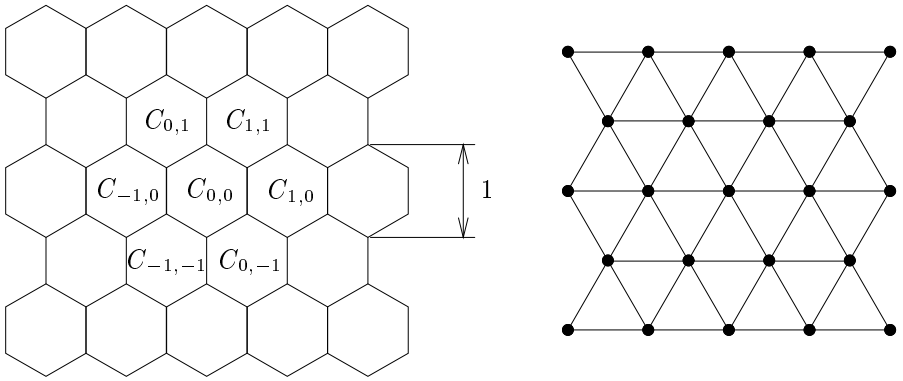

Fig. 3. The hexagonal tiling and the corresponding mesh M

Proof of Theorem 1:

Proof. For the purposes of the proof we will represent cells as vertices of an infinite triangular mesh M shown in Figure 3, where edges connect adjacent cells. For any pair of cells C and C', $dist_M(C, C')$ is the *mesh distance* measured as the number of edges of the shortest path connecting the vertices corresponding to C and C' in M.

Let $a = \lceil \frac{2k\sigma}{\sqrt{3}} \rceil$ and observe, that the cells $C_{i,j}$, $C_{i+a,j}$, $C_{i,j+a}$ and $C_{i+a,j+a}$ have pairwise plane distance greater than $k\sigma$ and we label these cells by the same number. By this method, we have to define a suitable labeling of a^2 cells $C_{i,j}$ with coordinates $i, j \in \{1, .., a\}$.

We order these a^2 cells arbitrarily and label them by first-fit method, with respect to labels of adjoin cells shifted by a.

Consider that we have to select a label for a cell C. For any $i \geq 1$, there exist exactly $6i$ cells that are at mesh distance i from C, and their *plane distance* from C is at least $\frac{3i-4}{4}$ (it is nonnegative number for $i \geq 2$.) The maximum i we consider is $\lfloor \frac{4k\sigma+4}{3} \rfloor$, because for higher values of i the plane distance become greater than $k\sigma$ and even if some of these cells are already labeled, this cannot influent the first-fit assignment of the label for C.

If any cell at mesh distance i from C is already labeled, then by the definition of the circular l-labeling, at most $2p_{\lceil \frac{3i-4}{4\sigma} \rceil} - 1$ numbers are unavailable for C.

In the worth case at most

$$6\Big(2p_1 - 1 + \sum_{i=2}^{\lfloor \frac{4k\sigma+4}{3} \rfloor} i \cdot (2p_{\lceil \frac{3i-4}{4\sigma} \rceil} - 1)\Big) = l^* - 1$$

numbers are forbidden as a label for C. Hence, among l^* distinct labels there exists at least one feasible candidate. □

Proof of Theorem 3:

Proof. Take any $t \in (1, \sqrt{2})$ and define an auxiliary sequence $a_i = \lfloor \frac{(i-1)\sigma+1}{t\sqrt{2}} + 1 \rfloor$ for $i = 2, ..., k$.

We define a set $\mathcal{D} = \{D_{1,1}, D_{1,2}, ..., D_{a_k,a_k}\}$ of a_k^2 unit disks, where disk $D_{j,l}$ has center (jt, lt) and j, l are integers from the set $\{1, 2, ..., a_k\}$. The graph $G_{\mathcal{D}}$ consist of a_k^2 independent vertices.

Now consider disks D and D' with coordinates j, l and j', l' and take the maximum i such that $|j - j'| \leq a_i$ and $|l - l'| \leq a_i$. It is easy to see that D and D' are at plane distance at most $(i-1)\sigma$, and therefore there exists a set $\overline{\mathcal{D}(D, D')}$ of $(i-1)$ disks of diameter σ such that in $G_{\mathcal{D} \cup \overline{\mathcal{D}(D,D)}}$ the vertices $D_{j,l}$ and $D_{j',l'}$ are at distance i.

Let $\overline{\mathcal{D}}$ be the union of $\overline{\mathcal{D}(D, D')}$ over all $D, D' \in \mathcal{D}$. We have showed that for every $i : 2 \leq i \leq k$ the a_i^2 vertices from $\{D_{j,l} : 1 \leq j, l \leq a_i\}$ are in $G_{\mathcal{D} \cup \overline{\mathcal{D}}}$ mutually at distance at most i Hence any on-line distance labeling algorithm $\mathcal{B}$ needs at least a_i^2 labels that pairwise differ by at least p_i to allow the extension to the labeling from $\mathcal{D}$ to $\mathcal{D} \cup \overline{\mathcal{D}}$

Therefore on the set $\mathcal{D}$ the algorithm $\mathcal{B}$ must use label of size at least

$$1 + \max_{i=2,...,k} \Big\{ \Big(\Big\lfloor \frac{(i-1)\sigma + 1}{t\sqrt{2}} + 1 \Big\rfloor^2 - 1 \Big) p_i \Big\},$$

After setting $t = \frac{3}{2\sqrt{2}}$ and suitable estimating, we get that $\mathcal{B}$ needs a label of size at least $1 + \frac{\sigma^2}{9} \max_{i=2,...,k}\{i^2 p_i\}$, although $\chi_{(p_1,...,p_k)}(G_{\mathcal{D}}) = 1$ □

Proof of Corollary 2:

Proof. First, observe that if the set $\mathcal{D}$ from the proof of Theorem 3 is extended by a pair of intersecting disks (but not intersecting the others), then the same argument applied to this new graph $G_{\mathcal{D}}$ gives that the lower bound on the performance ratio of any on-line labeling algorithm for such disk graphs is at least

$$\frac{1 + \frac{\sigma^2}{9} \max_{i=2,...,k}\{i^2 p_i\}}{1 + p_1} \geq c \frac{\sigma^2 \max_{i=2,...,k}\{i^2 p_i\}}{1 + p_1},$$

and the upper bound given by Theorem 2 is

$$\frac{2 + 12\Big(2p_1 - 1 + \sum_{i=2}^{\lfloor \frac{4k\sigma+4}{3} \rfloor} i \cdot (2p_{\lceil \frac{3i-4}{4\sigma} \rceil} - 1)\Big)}{1 + p_1} \leq c' \frac{\sigma^2 \sum_{i=2}^{k} i p_i}{1 + p_1}.$$

for suitable constants c, c' that do not depend on σ nor on $(p_1, ..., p_k)$.

(One term σ appears due to the change in number of elements from $\lfloor \frac{4k\sigma+4}{3} \rfloor$ to k, the other follows from the same change and i inside the sum.)

Let j be the index for which the term $i^2 p_i$ attains its maximum, i.e. for all $i \geq 2 : p_i \leq \frac{j^2}{i^2} p_j$.

Then

$$\sum_{i=2}^{k} i p_i \leq \sum_{i=2}^{k} \frac{j^2}{i} p_j = j^2 p_j \sum_{i=2}^{k} \frac{1}{i} \leq \max_{i=2,\ldots,k} \{i^2 p_i\} \cdot O(\log k)$$

and the statement of Corollary 2 is proved. □