



Distributed coloring and the local structure of unit-disk graphs

Louis Esperet, Sébastien Julliot, Arnaud de Mesmay

► To cite this version:

Louis Esperet, Sébastien Julliot, Arnaud de Mesmay. Distributed coloring and the local structure of unit-disk graphs. Theoretical Computer Science, 2023, 944, pp.113674. <10.1016/j.tcs.2022.12.024>. <hal-03269425>

HAL Id: hal-03269425

<https://hal.science/hal-03269425v1>

Submitted on 31 Mar 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC 4.0 - Attribution - Non-commercial use - International License



Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs



Distributed coloring and the local structure of unit-disk graphs ☆,☆☆

Louis Esperet^{a,*}, Sébastien Julliot^a, Arnaud de Mesmay^b

^a Laboratoire G-SCOP (CNRS, Grenoble-INP), Grenoble, France

^b LIGM, CNRS, Univ. Gustave Eiffel, ESIEE Paris, F-77454 Marne-la-Vallée, France

ARTICLE INFO

Article history:

Received 13 December 2021

Received in revised form 13 December 2022

Accepted 16 December 2022

Available online xxxx

Keywords:

Unit-disk graphs

Distributed coloring

Average degree

ABSTRACT

Coloring unit-disk graphs efficiently is an important problem in the global and distributed setting, with applications in radio channel assignment problems when the communication relies on omni-directional antennas of the same power. In this context it is important to bound not only the complexity of the coloring algorithms, but also the number of colors used. In this paper, we consider two natural distributed settings. In the location-aware setting (when nodes know their coordinates in the plane), we give a constant time distributed algorithm coloring any unit-disk graph G with at most $4\omega(G)$ colors, where $\omega(G)$ is the clique number of G . This improves upon a classical 3-approximation algorithm for this problem, for all unit-disk graphs whose chromatic number significantly exceeds their clique number. When nodes do not know their coordinates in the plane, we give a distributed algorithm in the LOCAL model that colors every unit-disk graph G with at most $5.68\omega(G) + 1$ colors in $O(\log^* n)$ rounds. This algorithm is based on a study of the local structure of unit-disk graphs, which is of independent interest. We conjecture that every unit-disk graph G has average degree at most $4\omega(G)$, which would imply the existence of a $O(\log n)$ round algorithm coloring any unit-disk graph G with (approximately) $4\omega(G)$ colors in the LOCAL model. We provide partial results towards this conjecture using Fourier-analytical tools.

© 2022 Elsevier B.V. All rights reserved.

1. Introduction

A *unit-disk graph* is a graph G whose vertex set is a collection of points $V \subseteq \mathbb{R}^2$, and such that two vertices $u, v \in V$ are adjacent in G if and only if $\|u - v\| \leq 1$, where $\|\cdot\|$ denotes the Euclidean norm. Unit-disk graphs are a classical model of wireless communication networks, and are a central object of study in distributed algorithms (see the survey [21] for an extensive bibliography on this topic). A classical way to design distributed communication protocols avoiding interferences is to find a proper coloring of the underlying unit-disk graph: the protocol then lets each vertex of the first color communicate with their neighbors, then each vertex of the second color, etc. Clearly the efficiency of the protocol depends on the number

☆ This article belongs to Section A: Algorithms, automata, complexity and games, Edited by Paul Spirakis.

☆☆ The authors are partially supported by ANR Projects GATO (ANR-16-CE40-0009-01), GrR (ANR-18-CE40-0032), MIN-MAX (ANR-19-CE40-0014) and SoS (ANR-16-CE40-0009-01), and by LabEx PERSYVAL-lab (ANR-11-LABX-0025). A preliminary version of this work appeared in the proceedings of the 17th International Symposium on Algorithms and Experiments for Wireless Sensor Networks (ALGOSENSORS 2021) [6].

* Corresponding author.

E-mail addresses: louis.esperet@grenoble-inp.fr (L. Esperet), arnaud.de-mesmay@univ-eiffel.fr (A. de Mesmay).

of colors used, so it is important to minimize the total number of colors (in addition to optimizing the complexity of the distributed coloring algorithm).

The *chromatic number* of a graph G , denoted by $\chi(G)$, is the smallest number of colors in a proper coloring of G . The *clique number* of G , denoted by $\omega(G)$, is the largest size of a clique (a set of pairwise adjacent vertices) in G . Note that for any graph G we have $\omega(G) \leq \chi(G)$, but the gap between the parameters can be arbitrarily large in general (see [20] for a recent survey on the relation between ω and χ for various graph classes). However, for unit-disk graphs it is known that $\chi(G) \leq 3\omega(G) - 2$ [15] (see also [9] for a different proof), and improving the multiplicative constant 3 is a longstanding open problem. It should be noted that while computing $\chi(G)$ for a unit-disk graph G is NP-hard, computing $\omega(G)$ for a unit-disk graph G can be done in polynomial time [5].

In the LOCAL model, introduced by Linial [13], the graph G that we are trying to color models a communication network: its vertices are processors of infinite computational power and its edges are communication links between (some of) these nodes. The vertices exchange messages with their neighbors in a certain number of synchronous rounds of communication (the round complexity), and then (in the case of graph coloring) each vertex outputs its color in a proper coloring of G . When in addition each vertex knows its coordinates in the plane, we call the model the location-aware LOCAL model (more details about these models will be given in Section 2.1).

In the location-aware LOCAL model, the following is a classical result [11,12,22,23] (see also [21] for a survey on *local algorithms*, which are algorithms that run in a constant number of rounds).

Theorem 1.1 ([11,12,22,23]). *A coloring of any unit-disk graph G with at most $3\chi(G)$ colors can be obtained in a constant number of rounds by a deterministic distributed algorithm in the location-aware LOCAL model.*

We prove the following complementary result, which improves the number of colors as soon as $\chi(G) \geq \frac{4}{3}\omega(G)$. Note that there exists an infinite family of unit disk graphs G for which $\chi(G) \geq \frac{3}{2}\omega(G)$ [14]. Our algorithm is inspired by the proof of [9] showing that unit-disk graphs satisfy $\chi \leq 3\omega$.

Theorem 1.2. *A coloring of any unit-disk graph G with at most $4\omega(G)$ colors can be computed in a constant number of rounds by a deterministic algorithm in the location-aware LOCAL model.*

Given two integers $p \geq q \geq 1$, a $(p:q)$ -coloring of a graph G is an assignment of q -element subsets of $[p]$ to the vertices of G , such that the sets assigned to any two adjacent vertices are disjoint. The *fractional chromatic number* $\chi_f(G)$ is defined as the infimum of $\{\frac{p}{q} \mid G \text{ has a } (p:q)\text{-coloring}\}$ [18] (it can be proved that this infimum is indeed a minimum). Observe that a $(p:1)$ -coloring is a (proper) p -coloring, and that for any graph G , $\omega(G) \leq \chi_f(G) \leq \chi(G)$. The fractional chromatic number is often used in scheduling as an alternative to the chromatic number when resources are fractionable, which is the case for communication protocols. It was proved in [8] that for any unit-disk graph G , $\chi_f(G) \leq 2.155\omega(G)$. Here we give an efficient distributed implementation of this result.

Theorem 1.3. *There is a constant $q \in \mathbb{N}$, such that in any unit-disk graph G , there exists an integer p with $\frac{p}{q} \leq 2.156\omega(G)$ such that a $(p:q)$ -coloring of G can be computed in $O(1)$ rounds by a deterministic distributed algorithm in the location-aware LOCAL model.*

We now turn to the abstract setting, where vertices do not have access to their coordinates in the plane. For a real number $n > 0$, let $\log^* n$ be the number of times we have to iterate the logarithm, starting with n , to reach a value in $(0, 1]$. Since paths are unit-disk graphs and coloring n -vertex paths with a constant number of colors takes $\Omega(\log^* n)$ rounds in the LOCAL model [13], coloring unit-disk graphs of bounded clique number with a bounded number of colors also takes $\Omega(\log^* n)$ rounds in the LOCAL model. Recalling that for any unit-disk graph G , $\omega(G) \leq \chi(G) \leq 3\omega(G)$, a natural question is the following.

Question 1.4. What is the minimum real $c > 0$ such that a coloring of any n -vertex unit-disk graph G with $c \cdot \omega(G)$ colors can be obtained in $O(\log^* n)$ rounds in the LOCAL model?

Using the folklore result that any unit-disk graph G has maximum degree at most $6\omega(G) - 6$ (see [9]), together with the fact that unit-disk graphs of maximum degree Δ can be colored efficiently with $\Delta + 1$ colors in the LOCAL model [19], we deduce that unit-disk graphs G can be colored efficiently with $6\omega(G)$ colors in the LOCAL model. We obtain the following improved version by studying the local structure of unit-disk graphs, using techniques that might be of independent interest.

Theorem 1.5. *Every unit-disk graph G can be colored with at most $5.675\omega(G) + 1$ colors by a deterministic distributed algorithm in the LOCAL model, running in $O(\log^* n)$ rounds.*

In relation to Question 1.4, it is natural to study the power of graph coloring algorithms in unit-disk graphs in a different (less restrictive) range of round complexity.

Question 1.6. What is the minimum real $c > 0$ such that a coloring of any n -vertex unit-disk graph G with $c \cdot \omega(G)$ colors can be obtained in $O(\log n)$ rounds in the LOCAL model?

An interesting property of the $O(\log n)$ range of round complexity (compared to the $O(\log^* n)$ range) is that it allows to solve coloring problems for graphs of bounded average degree (rather than bounded maximum degree). The *average degree* of a graph G is the average of its vertex degrees. The *maximum average degree* of a graph G is the maximum average degree of a subgraph H of G . In [2], Barenboim and Elkin gave, for any $\epsilon > 0$, a deterministic distributed algorithm coloring n -vertex graphs of maximum average degree d with at most $(1 + \epsilon)d + 3$ colors in $O(\frac{d}{\epsilon} \log n)$ rounds (the result was proved in terms of arboricity rather than average degree).

While the chromatic number and degeneracy of unit-disk graphs (as a function of the clique number) are well studied topics, it seems that little is known about the average degree of unit-disk graphs. We conjecture the following:

Conjecture 1.7. *Every unit-disk graph G has average degree at most $4\omega(G)$.*

It can be checked that the constant 4 is best possible by considering uniformly distributed points in the plane. In this case each vertex has degree equal to some density constant $c > 0$ times the area of a disk of radius 1, so the average degree is $c \cdot \pi$. On the other hand any clique is contained in a region of diameter at most 1, and the area of such a region is known to be maximized for a disk of radius $\frac{1}{2}$ [3], i.e., the graph has clique number $c \cdot \pi/4$, giving a ratio of 4 between the average degree and the clique number.

Using the result of Barenboim and Elkin [2] mentioned above, Conjecture 1.7 would imply the existence of a deterministic distributed coloring algorithm using $(4 + \epsilon)\omega(G)$ colors in $O(\frac{\omega(G)}{\epsilon} \log n)$ rounds (for fixed $\epsilon > 0$ and sufficiently large $\omega(G) = \Omega(1/\epsilon)$).

Unfortunately we are quite far from proving Conjecture 1.7 at the moment. Our best result so far is the following.

Theorem 1.8. *Every unit-disk graph G has average degree at most $5.68\omega(G)$.*

Our final result shows that Conjecture 1.7, if true, is more subtle than it might seem. In the example above showing the optimality of Conjecture 1.7, the largest cliques are formed by sets of points that are all contained within some disk of radius $\frac{1}{2}$. We may naturally define the *disk clique number* of a unit-disk graph G , denoted by $\omega_D(G)$, to be the largest size of a clique contained within such a disk. Note that $\omega_D(G)$ depends on the embedding of G in the plane, not just of the underlying graph (contrary to the clique number $\omega(G)$). One may wonder whether Conjecture 1.7 holds for $\omega_D(G)$ instead of $\omega(G)$ – by the same argument as for Conjecture 1.7, the constant 4 would be best possible. It is however not the case and we show the following stronger lower bound:

Theorem 1.9. *There exists a unit-disk graph G of average degree at least $4.0905\omega_D(G)$.*

Here, 4.0905 is roughly $4 \cdot (1 + \frac{J_1(2B)}{2B})$, where J_1 is the Bessel function of the first kind and B its first zero. Theorem 1.9 shows that any approach to prove Conjecture 1.7 needs to account for possible shapes of cliques different than the ones contained in disks of radius $\frac{1}{2}$. As the constant suggests, our proof technique relies on Fourier analysis and special functions and we believe that they could be of independent interest in the study of unit-disk graphs, and more generally of intersection graphs of other objects in the plane. In particular they also allow us to prove Conjecture 1.7 when the distribution of points is sufficiently close to the uniform distribution (in the sense that the support of its Fourier transform is contained in a small disk around 0).

Organization of the paper. We start with a presentation of the LOCAL model and some basic results on coloring and unit-disk graphs in Section 2. Section 3 is devoted to proving our main results in the location-aware setting, Theorem 1.2 and Theorem 1.3. In Section 4, we study the local structure of unit-disk graphs and deduce our coloring result in the LOCAL model, Theorem 1.5, together with our upper bound on the average degree of unit-disk graphs, Theorem 1.8 above. In Section 5, we introduce Fourier-analytical tools to approach Conjecture 1.7 and prove Theorem 1.9.

2. Preliminaries

2.1. Distributed models of communication

All our results are proved in the LOCAL model, introduced by Linial [13]. The underlying network is modeled as an n -vertex graph G whose vertices have unbounded computational power, and whose edges are communication links between the corresponding vertices. In the case of deterministic algorithms, each vertex of G starts with an arbitrary unique identifier (an integer between 1 and n^α , for some constant $\alpha \geq 1$, such that all integers assigned to the vertices are distinct). For randomized algorithms, each vertex starts instead with a collection of (private) random bits. The vertices then exchange messages (possibly of unbounded size) with their neighbors in synchronous rounds, and after a fixed number of rounds

(the *round complexity of the algorithm*), each vertex v outputs its local “part” of a global solution to a combinatorial problem in G , for instance its color $c(v)$ in some proper k -coloring c of G .

It turns out that with the assumption that messages have unbounded size, after t rounds we can assume without loss of generality that each vertex v “knows” its neighborhood $B_t(v)$ at distance t (the set of all vertices at distance at most t from v). More specifically v knows the labeled subgraph of G induced by $B_t(v)$ (where the labels are the identifiers of the vertices), and nothing more, and the output of v is based solely on this information (see [13]).

The goal is to minimize the round complexity. Since nodes have infinite computational power, the paragraph above shows that any problem can be solved in a number of rounds equal to the diameter of the graph, which is at most n when G is connected. The goal is to obtain algorithms that are significantly more efficient, i.e., of round complexity $O(\log n)$, or even $O(\log^* n)$.

In this paper, we will also consider the location-aware LOCAL model, which is a variant of the LOCAL model in which the n -vertex graph modeling the communication network is a unit-disk graph embedded in the plane, and every vertex knows its coordinates in the embedding.

2.2. Distributed coloring

Consider a graph G . In the $(\deg + 1)$ -list coloring problem, each vertex v is given a list $L(v)$ of colors such that $|L(v)| \geq d(v) + 1$, where $d(v)$ denotes the degree of v in G , and the goal is to color each vertex v with a color from its list $L(v)$, so that any two adjacent vertices receive different colors.

A graph G has *bounded growth* if there is a function f such that for any integer $r \geq 1$, and any vertex $v \in V(G)$, the maximum number of independent (i.e., pairwise non-adjacent) vertices among the vertices at distance at most r from v is at most $f(r)$. It was proved in [19] that any graph of bounded growth and maximum degree Δ can be colored with $\Delta + 1$ colors in $O(\log^* n)$ rounds in the LOCAL model (where the hidden constant in the $O(\cdot)$ notation only depends on the function f bounding the growth). As noted by an anonymous reviewer, the proof of [19] extends verbatim to the $(\deg + 1)$ -list coloring problem.

Theorem 2.1 ([19]). *There exists a deterministic distributed algorithm in the LOCAL model that solves the $(\deg + 1)$ -list coloring problem in any n -vertex graph of bounded growth in $O(\log^* n)$ rounds.*

It can be checked that unit-disk graphs have bounded growth (the function f in this case is quadratic, see [19]), which yields the following immediate corollary.

Corollary 2.2. *There exists a deterministic distributed algorithm in the LOCAL model that solves the $(\deg + 1)$ -list coloring problem in any n -vertex unit-disk graph in $O(\log^* n)$ rounds.*

2.3. Unit-disk graphs

For $0 \leq r \leq 1$, and a point v let $D_r(v)$ be the disk of radius r centered in v and let $C_r(v)$ the circle of radius r centered in v . Given a unit-disk graph G embedded in the plane, a vertex v , and a real $0 \leq r \leq 1$, we denote by $d_r(v)$ the number of neighbors of v lying in $D_r(v)$, and by $x_r(v)$ the number of neighbors of v lying on $C_r(v)$. Note that $d_1(v)$ is precisely $d(v)$, the degree of v in G , and for every $0 \leq r \leq 1$, $d_r(v) = \sum_{s \in [0, r]} x_s(v)$ (note that since $s \mapsto x_s(v)$ has finite support, this sum is well defined).

It is well known that a disk of radius 1 can be covered by 6 regions of diameter 1 (see [9]), and thus the neighborhood of each vertex of G can be covered by 6 cliques (and thus $d(v) \leq 6\omega(G)$ for each vertex v of G). A *Reuleaux triangle* R is the intersection of 3 disks of radius 1, centered in the three vertices of an equilateral triangle T of side length 1 (see the green region in Fig. 1). The three vertices of T are also called vertices of R . Note that Reuleaux triangles have diameter 1, and 6 Reuleaux triangles are enough to cover a disk of radius 1. Moreover, points close to the center of the disk are covered by more triangles than points on the outer circle. This can be used to prove the following.

Lemma 2.3. *For each vertex v of a unit-disk graph G embedded in the plane, $\sum_{r \in [0, 1]} (2 - r)x_r(v) \leq 6\omega(G)$.*

Proof. Let v be a vertex of G . Note that each neighbor of v lies in $D_1(v)$ (the disk of radius 1 centered in v). Let x be an arbitrary point of $C_1(v)$. For $\theta \in [0, 2\pi)$, let R_θ be the Reuleaux triangle having y, v, z as vertices, in clockwise order (with y and z two vertices of $C_1(v)$ at distance 1), and such that the angle $\angle xvy$ is equal to θ (R_θ is depicted in green in Fig. 1). Consider also a neighbor u of v lying on $C_r(v)$, the circle of radius r centered in v . For θ chosen uniformly at random in the interval $[0, 2\pi)$, the probability that u is covered by R_θ is the length of $C_r(v) \cap R_\theta$ (the bold arc of Fig. 1) divided by $2\pi r$ (the circumference of $C_r(v)$). This probability is thus

$$\frac{(2 \arccos(r/2) - \pi/3)r}{2\pi r} = \frac{1}{\pi} \arccos(r/2) - 1/6.$$




Fig. 1. The intersection (in bold) of a circle C_r of radius r with a Reuleaux triangle R_θ (the green region). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Note that $x \mapsto \arccos x$ is concave in $[0, \frac{1}{2}]$, and thus $r \mapsto f(r) := \frac{1}{\pi} \arccos(r/2) - 1/6$ is concave in $[0, 1]$. As a consequence, for any $r \in [0, 1]$, $f(r) \geq (f(1) - f(0))r + f(0) = (\frac{1}{6} - \frac{1}{3})r + \frac{1}{3} = \frac{1}{3} - \frac{r}{6}$.

It follows that each neighbor of v at distance at most r from v is covered by R_θ with probability at least $\frac{1}{3} - \frac{r}{6}$. Therefore, the expected number of neighbors of v covered by R_θ is at least $\sum_{r \in [0, 1]} (\frac{1}{3} - \frac{r}{6}) x_r(v)$. Since each R_θ has diameter 1, the vertices of R_θ induce a clique in G and thus R_θ contains at most $\omega(G)$ vertices. It follows that $\sum_{r \in [0, 1]} (\frac{1}{3} - \frac{r}{6}) x_r(v) \leq \omega(G)$ and thus $\sum_{r \in [0, 1]} (2 - r) x_r(v) \leq 6\omega(G)$, as desired. $\square$

The following is a direct consequence of Lemma 2.3.

Corollary 2.4. For each vertex v of a unit-disk graph G embedded in the plane and $r \in [0, 1]$, $d(v) + (1 - r)d_r(v) \leq 6\omega(G)$. In particular $d_{1/2}(v) \leq 12\omega(G) - 2d(v)$.

Proof. By Lemma 2.3, $6\omega(G) \geq \sum_{s \in [0, 1]} (2 - s)x_s(v) = \sum_{s \in [0, r]} (2 - s)x_s(v) + \sum_{s \in [r, 1]} (2 - s)x_s(v) \geq (2 - r)d_r(v) + d(v) - d_r(v) = (1 - r)d_r(v) + d(v)$. By taking $r = \frac{1}{2}$, we obtain $d(v) + \frac{1}{2}d_{1/2}(v) \leq 6\omega(G)$, and thus $d_{1/2}(v) \leq 12\omega(G) - 2d(v)$, as desired. $\square$

In the next section will need the following useful observation of [9] (see also [8]).

Lemma 2.5 ([9]). Let G be a unit-disk graph embedded in the plane, such that the y -coordinates of any two vertices of G differ by at most $\sqrt{3}/2$. Then $\chi(G) = \omega(G)$.

We deduce the following easy corollary (which holds in the LOCAL model, so it does not require nodes to know their own coordinates in the plane).

Corollary 2.6. Let G be a unit-disk graph embedded in the plane, such that the y -coordinates of any two vertices of G differ by at most $\sqrt{3}/2$, and the x -coordinate differ by at most ℓ , for some real number ℓ . Then G can be colored with $\chi(G) = \omega(G)$ colors by a deterministic distributed algorithm running in $O(\ell)$ rounds in the LOCAL model.

Proof. Take any shortest path $P = v_1, v_2, \dots, v_k$ in G . Then for any $1 \leq i \leq k - 2$, $\|v_i - v_{i+2}\| > 1$, since otherwise P would not be a shortest path in G . Since y -coordinates differ by at most $\sqrt{3}/2$, it follows that the x -coordinates of v_i and v_{i+2} differ by at least $\frac{1}{2}$, and the x -coordinates of the vertices v_i with i odd are monotone (say increasing with loss of generality). Hence, the x -coordinates of v_1 and v_k differ by at least $k/4$, and by definition, $k \leq 4\ell$. As a consequence, any connected component of G has diameter at most 4ℓ , and thus any connected component can be colored optimally by a deterministic distributed algorithm running in $O(\ell)$ rounds in the LOCAL model. $\square$

3. Location-aware coloring

In the proceedings version of this paper, we claimed that for any $0 < \epsilon \leq 1$, a coloring of any unit-disk graph G with at most $(3 + \epsilon)\omega(G) + 6$ colors can be computed in $O(1/\epsilon)$ rounds by a deterministic distributed algorithm in the location-aware LOCAL model [6, Theorem 8]. Unfortunately, we recently discovered an error in our proof of [6, Lemma 3], on which [6, Theorem 8] was based, and at the moment we are not able to fix the proof. Here we prove the following weaker version instead.

Theorem 1.2. A coloring of any unit-disk graph G with at most $4\omega(G)$ colors can be computed in a constant number of rounds by a deterministic algorithm in the location-aware LOCAL model.




Fig. 2. Covering $V(G)$ with horizontal stripes in the proof of Theorem 1.2.




Fig. 3. In grey, the region $\mathcal{R}(x, y)$ defined in the proof of Theorem 1.3.

Proof. In the remainder, by the *distance between two regions* R_1 and R_2 of the plane, we mean the minimum distance between a point of R_1 and a point of R_2 . We start by covering $V(G)$ with consecutive horizontal stripes $\mathcal{S}_1, \mathcal{S}_2, \dots$, each of height $\sqrt{3}/2$ (see Fig. 2). Note that any two stripes $\mathcal{S}_j$ and $\mathcal{S}_{j+3}$ are at distance at least $2 \cdot \sqrt{3}/2 = \sqrt{3} > 1$ apart, so there are no edges connecting a vertex lying in $\mathcal{S}_j$ and a vertex lying in $\mathcal{S}_{j+3}$.

Let $\mathcal{R}$ be a union of rectangles of height $\sqrt{3}/2$ and length 1 (depicted in white in Fig. 2), with the following properties:

- (1) each rectangle of $\mathcal{R}$ is included in some horizontal stripe $\mathcal{S}_i$, $i \geq 1$,
- (2) any two consecutive rectangles of $\mathcal{R}$ on a stripe $\mathcal{S}_i$ lie at distance 5 apart,
- (3) the distance between the ends of a stripe and the closest rectangle of $\mathcal{R}$ on that stripe is at most 5,
- (4) any two rectangles of $\mathcal{R}$ lie at distance more than 1 apart.

Note that (4) can be obtained by simply shifting the rectangles of $\mathcal{R}$ in $\mathcal{S}_i$ to the right to obtain the rectangles of $\mathcal{R}$ in $\mathcal{S}_{i+1}$ (and possibly adding a new rectangle of $\mathcal{R}$ at the beginning of $\mathcal{S}_{i+1}$).

For any $i \geq 1$, let $\mathcal{S}'_i = \mathcal{S}_i \setminus \mathcal{R}$. Note that each $\mathcal{S}'_i$ consists of a sequence of rectangles of height $\sqrt{3}/2$ and length at most 5, all at distance more than 1 apart. It follows from this observation and Corollary 2.6 that for each $i \in \{0, 1, 2\}$, the subgraph G_i of G induced by the vertices lying in the union of all $\mathcal{S}'_j$ with $j \equiv i \pmod{3}$ has chromatic number at most $\omega := \omega(G)$, and can be colored with ω colors in $O(1)$ rounds in the location-aware LOCAL model. We will color each of these 3 graphs G_i , $i \in \{0, 1, 2\}$, with a disjoint set of at most ω colors. It remains to color the vertices of G lying in $\mathcal{R}$. Using Corollary 2.6, the set of vertices lying in each rectangle of $\mathcal{R}$ can be colored with ω colors in $O(1)$ rounds, and we can use the same set of ω colors for each rectangle of $\mathcal{R}$, as all these rectangles are at distance more than 1 apart. In total we have used at most $3 \cdot \omega + \omega = 4\omega$ colors, as desired. $\square$

It was proved in [8] that for any unit-disk graph G , $\chi_f(G) \leq 2.155 \omega(G)$. We now give an efficient distributed version of this result.

To give an intuition of the proof, note that the fractional chromatic number of G can be equivalently defined as the infimum $x > 0$ such that there is a probability distribution over the independent sets of G , such that with respect to this distribution, each vertex has probability at least $1/x$ to be in a random independent set. A way to define such a good distribution in a unit-disk graph G is to choose a random point (x, y) , and then consider the grey rectangles depicted in Fig. 3. The union of all vertices lying in grey rectangles induces a union of perfect graphs of bounded diameter and can thus be colored in a constant number of rounds (in the location-aware LOCAL model) with $\omega(G)$ colors. We now take a random color class among the $\omega(G)$ color classes, and it can be easily checked that each vertex has probability at least $\frac{\sqrt{3}/2}{1+\sqrt{3}/2} \cdot \frac{1/\epsilon}{1+1/\epsilon} \cdot \frac{1}{\omega(G)} \geq \frac{1}{(1+\epsilon)2.155\omega(G)}$ to lie in this random independent set (see [8] for a version of this argument with horizontal stripes instead of rectangles).

We now give an efficient distributed version of this argument that does not use randomization, and in particular allows to output (p, q) -colorings with bounded p and q , which is a desirable property in the context of local algorithms [1,4].

Theorem 1.3. *There is a constant $q \in \mathbb{N}$, such that in any unit-disk graph G , there exists an integer p with $\frac{p}{q} \leq 2.156 \omega(G)$ such that a $(p:q)$ -coloring of G can be computed in $O(1)$ rounds by a deterministic distributed algorithm in the location-aware LOCAL model.*

Proof. Fix any real $\epsilon > 0$. Let G be a unit-disk graph embedded in the plane, and let $\omega = \omega(G)$. Given $(x, y) \in \mathbb{R}^2$, let $R(x, y)$ be the axis-parallel rectangle of height $\sqrt{3}/2$, length $1/\epsilon$, and bottom-left corner (x, y) . Let $\mathcal{R}(x, y)$ be the union of all rectangles $R(x', y')$, for $x' \in x + (1 + 1/\epsilon) \cdot \mathbb{Z}$ and $y' \in y + (1 + \sqrt{3}/2) \cdot \mathbb{Z}$ (see the grey area in Fig. 3). It will be convenient to assume that each rectangle contains its top and right boundaries and excludes its bottom and left boundaries. In particular, any two points from distinct rectangles of $\mathcal{R}(x, y)$ are at distance more than 1 apart. Thus it follows from Corollary 2.6 that for any $(x, y) \in \mathbb{R}^2$, the subgraph of G induced by the vertices lying in $\mathcal{R}(x, y)$ can be colored with at most ω colors in $O(1/\epsilon)$ rounds in the location-aware LOCAL model.

Fix some integer $r \geq 1$. For any $0 \leq i \leq r - 1$, we consider the region $\mathcal{R}_i := \mathcal{R}(\frac{i}{r}(1 + 1/\epsilon), \frac{i}{r}(1 + \sqrt{3}/2))$. Note that for each point (x', y') of the plane, there are at most $\lceil r \cdot \frac{1}{1+1/\epsilon} \rceil$ integers $0 \leq i \leq r - 1$ such that x' lies outside of the projection of $\mathcal{R}_i$ on the horizontal axis, and there are at most $\lceil r \cdot \frac{1}{1+\sqrt{3}/2} \rceil$ integers $0 \leq i \leq r - 1$ such that y' lies outside of the projection of $\mathcal{R}_i$ on the vertical axis. As a consequence, each point of the plane lies in at least

$$r - \lceil r \cdot \frac{1}{1+\sqrt{3}/2} \rceil - \lceil r \cdot \frac{1}{1+1/\epsilon} \rceil \geq r - (r \cdot \frac{1}{1+\sqrt{3}/2} + 1) - (r \cdot \frac{1}{1+1/\epsilon} + 1) \quad (1)$$

$$= r - r \cdot \frac{1}{1+\sqrt{3}/2} - r \cdot \frac{1}{1+1/\epsilon} - 2 \quad (2)$$

regions $\mathcal{R}_i$. Note that for any $\epsilon > 0$, $\frac{1}{1+1/\epsilon} = \frac{\epsilon}{\epsilon+1} \leq \epsilon$, and thus (2) is at least $(0.464101 - \epsilon)r - 2$. It follows that for sufficiently large (but constant) r and sufficiently small (but constant) $\epsilon > 0$, the quantity (2) is at least $0.4641r$. Now, for any $0 \leq i \leq r - 1$, we consider a set S_i of ω colors (such that all sets S_i are pairwise disjoint) and color the subgraph of G induced by the vertices lying in $\mathcal{R}_i$ with colors from S_i . By the paragraph above, this can be done in $O(1/\epsilon) = O(1)$ rounds in the location-aware LOCAL model. In total, at most $p := r\omega$ colors are used, and since each point of the plane is covered by at least $0.4641r$ regions $\mathcal{R}_i$, each vertex receives at least $q := \lceil 0.4641r \rceil$ different colors. We obtain a $(p:q)$ -coloring of G with $p/q \leq \frac{\omega}{0.4641} \leq 2.156\omega$, as desired. $\square$

4. Coloring without coordinates

We start by giving a bound on the average of the degrees of two adjacent vertices. Before proceeding with the proof of our result, we give a high-level intuition of our approach. A unit-disk graph can be viewed as multiset of points in the plane (several points might be located at the same coordinates), or equivalently as a set of points carrying some integral positive weights, or equivalently again, as a function $\mu : \mathbb{R}^2 \rightarrow \mathbb{N}^{\geq 0}$, where $\mu((x, y)) = i \geq 0$ if and only if the unit-disk graph has i vertices located at (x, y) . Let us call $\mu((x, y))$ the *weight* of the point (x, y) . The assumption that the underlying unit-disk graph has clique number at most ω can be described by an (infinite) set of *linear* inequalities of the following form: for any region R of the plane with diameter 1, the sum of the weights of the points of R is at most ω . The degree of a vertex v can also be described as a linear combination of the weights: it is the sum of the weights of the points in $D_1(v)$. Therefore, the average of the degrees of two vertices u and v can also be described by a linear combination of the weights. We want to prove an upper bound on this quantity for all unit-disk graphs of maximum clique size ω , so what we can do is set our weights as variables, add the (linear) constraints implying that all cliques have size at most ω , and maximize the (linear) objective function defined as the average of the degrees of two adjacent vertices. The optimum will be the desired upper bound on the average of the degrees of two adjacent vertices in a unit-disk graph of maximum clique size ω . However, we quickly run into a number of complications.

- In order to be solved efficiently, a linear program needs to allow for rational solutions. But this is fine, we can simply relax the condition that the weights are integral by allowing rational weights, and compute the optimum. It will still give an upper bound on the integral version of the linear program.
- There is an infinite number of constraints, as there is an infinite number of regions of diameter 1 in the plane. So here too, we relax the problem and only consider a well chosen (finite) set of regions of diameter 1 and set as (linear) constraints that the sum of the weights in each of these regions is at most ω . Again, the optimum with respect to this smaller set of inequalities will be at least the optimum for the original problem, and thus still be an upper bound on the average of the degrees of two adjacent vertices.
- There is an infinite number of variables (all points in $D_1(u) \cup D_2(v)$). Here, we argue that because of our choice of finitely many regions for the constraints, there are optimal solutions whose support is discrete (located at the intersection of the boundaries of these regions). Moreover, using symmetry arguments, we show that it suffices to consider a very small number of points, and eventually our proof boils down to solving two linear programs with a handful of variables and inequalities.

We now proceed with the formal proof of our result.




Fig. 4. Two adjacent vertices u and v in the first part of the proof of Lemma 4.1.

Lemma 4.1. *For every two vertices u and v in a unit-disk graph G embedded in the plane such that $\frac{1}{2} \leq \|u - v\| \leq 1$, we have $\frac{1}{2}(d(u) + d(v)) \leq 5.675 \omega(G)$.*

Proof. Let $\delta = \|u - v\| \in [\frac{1}{2}, 1]$ and $\omega = \omega(G)$. Given a bounded subset R of the plane, we denote by $\mu(R)$ the number of vertices of G lying in R . Note that for any vertex x of G , $d(x) = \mu(D_1(x))$, and thus it is enough to prove that $\frac{1}{2}(\mu(D_1(u)) + \mu(D_1(v))) \leq 5.675 \omega(G)$. We will give two different upper bounds on $\frac{1}{2}(\mu(D_1(u)) + \mu(D_1(v)))$, and the minimum of the two will be at most 5.675ω for every $\delta \in [\frac{1}{2}, 1]$. The second upper bound will be stronger than the first, except when δ , the distance between u and v , is close to 1 (see Fig. 5 for a comparison of the two upper bounds as a function of $\delta \in [\frac{1}{2}, 1]$).

For the first bound, consider the region $S_u \subset D_1(u)$ (bounded by the fat red dashed curve) and the region $S_v \subset D_1(v)$ (bounded by the fat blue curve) in Fig. 4. S_u is defined as follows: take 6 points $z_1, \dots, z_6$ appearing in clockwise order on $C_1(u)$, such that any two consecutive points z_i, z_{i+1} (with indices modulo 6) form an equilateral triangle with u , and such that z_1 and z_2 are symmetric with respect to the line (uv) . Now define S_u as the union of the five Reuleaux triangles with vertices $uz_2z_3, uz_3z_4, uz_4z_5, uz_5z_6$, and uz_6z_1 . The region S_v is defined as the symmetric of S_u with respect to the perpendicular bisector of the segment $[u, v]$. Note that by definition, each of S_u and S_v can be covered by 5 regions of diameter 1, and thus

$$\mu(S_u) \leq 5\omega \quad \text{and} \quad \mu(S_v) \leq 5\omega. \quad (3)$$

By Lemma 2.3, we also have

$$\sum_{r \in [0,1]} (2-r) \mu(C_r(u)) \leq 6\omega \quad \text{and} \quad \sum_{r \in [0,1]} (2-r) \mu(C_r(v)) \leq 6\omega \quad (4)$$

(recall that $C_r(u)$ is the circle of radius r centered in u). Note that $d(u) = \mu(D_1(u)) = \sum_{r \in [0,1]} \mu(C_r(u))$ and $d(v) = \mu(D_1(v)) = \sum_{r \in [0,1]} \mu(C_r(v))$.

Subject to these inequalities (and with ω fixed), our goal is to maximize $\frac{1}{2}(\mu(D_1(u)) + \mu(D_1(v)))$ over all point sets in the plane. As explained in the introduction of this section, it is convenient to relax the problem and optimize over points sets carrying some nonnegative (non necessarily integral) weights. In this context, $\mu(R)$ is simply defined as the sum of the weights of the points lying in R (the objective function $\frac{1}{2}(\mu(D_1(u)) + \mu(D_1(v)))$ and the linear constraints (3) and (4), which were defined using μ , are modified accordingly). Thus we have a linear program with objective function $\frac{1}{2}(\mu(D_1(u)) + \mu(D_1(v)))$ (which we seek to maximize), whose variables are the weights of the points in the plane, with linear constraints given by (3) and (4). The solution of this linear program will give us an upper bound for our original (integral) optimization problem.

We now argue that some optimal solution of the linear program defined above has finite support. We define the following points (see Fig. 4):

- w_1 is a point in $(D_1(u) \cap D_1(v)) \setminus (S_u \cup S_v)$, such that $\|u - w_1\| + \|v - w_1\|$ (the sum of the distances of w_1 to u and v) is maximized,
- w_2 is one of the two points in the intersection of $C_1(u)$ and $C_1(v)$,
- w_3 is a point of $(C_1(v) \cap S_u) \setminus S_v$, which is arbitrarily close to S_v .
- w'_3 is the symmetric of w_3 with respect to the perpendicular bisector of the segment $[u, v]$.

Observe that the linear inequalities (3) and (4) and the objective function $\frac{1}{2}(\mu(D_1(u)) + \mu(D_1(v)))$ are symmetric with respect to the line (uv) and to the perpendicular bisector of the segment $[uv]$. This allows us to restrict ourselves to optimal solutions that have the following additional properties:




Fig. 5. Two upper bounds on $\frac{1}{2\omega}(d(u) + d(v))$ when $\|u - v\| = \delta$. The function $\frac{f_5(\delta)}{\omega}$ is plotted with a red dashed curve, and the function $\frac{f_4(\delta)}{\omega}$ is plotted in blue.

- they are symmetric with respect to the perpendicular bisector of the segment $[uv]$ (i.e., we can interchange u and v without affecting the solution), and
- their support lies in one of the two half-planes defined by (uv) , say the upper half-plane.

For the first property, it suffices to take the average of some optimal solution and the symmetric image of this solution with respect to the perpendicular bisector of the segment $[uv]$, and for the second property, it suffices to take the union of the solution in the upper half-plane and the symmetric of the solution in the lower half-plane with respect to (uv) .

As $D_1(u) \setminus D_1(v)$ and $D_1(v) \setminus D_1(u)$ are symmetric with respect to the perpendicular bisector of the segment $[uv]$, it follows that in some optimal solution as above we have $\mu(D_1(u) \setminus D_1(v)) = \mu(D_1(v) \setminus D_1(u))$. If this quantity is non-zero, we can modify this solution by deleting all the points in the support of the solution that are in the symmetric difference of $D_1(u)$ and $D_1(v)$ and adding weight $\mu(D_1(u) \setminus D_1(v)) = \mu(D_1(v) \setminus D_1(u))$ to the point w_2 . The objective function remains unchanged and (3) and (4) are still satisfied, so we obtain a (symmetric) optimal solution whose support lies in the intersection of $D_1(u)$ and $D_1(v)$. Now, if a symmetric weighted point set satisfies this property as well as (3) and (4), and contains a point z (and its symmetric z') distinct from w_1, w_2, w_3 (and their symmetric images), then z and z' can be moved locally so that

- the boundaries of S_u and S_v are not crossed in the motion, and
- $\|u - z\| + \|v - z\|$ increases (and by symmetry, $\|u - z'\| + \|v - z'\|$ increases).

The first property implies that the inequalities of (3) remain valid after the motion. The second property implies that the left-hand sides of the two inequalities of (4) are non-increasing: the contribution of z and z' to $\sum_{r \in [0,1]} (2-r)\mu(C_r(u))$ is

$$(2 - \|u - z\|)\mu(z) + (2 - \|u - z'\|)\mu(z') = (4 - (\|u - z\| + \|v - z\|))\mu(z),$$

where we have used that $\mu(z) = \mu(z')$ and $\|u - z'\| = \|v - z\|$ by symmetry. The same holds for the second inequality by symmetry. It follows that the two inequalities of (4) are still valid after the motion, while the objective function remains unchanged. This shows that we can assume without loss of generality that in some optimal solution, the support of the weighted point set is included in $\{w_1, w_2, w_3, w'_3\}$, and the weight of w_3 is equal to the weight of w'_3 .

Once this has been observed, it remains to solve the following finite linear program (here $x_1 = \mu(\{w_1\})$, $x_2 = \mu(\{w_2\})$, and $x_3 = \mu(\{w_3\}) = \mu(\{w'_3\})$).

$$\begin{aligned} \max_{x_1, x_2, x_3} \quad & x_1 + x_2 + 2x_3 \\ \text{s.t.} \quad & x_2 + x_3 \leq 5\omega \\ & \left(2 - \sqrt{\frac{\delta^2}{4} + \left(1 - \sqrt{1 - \frac{\delta^2}{4}}\right)^2}\right) \cdot x_1 + x_2 + \left(3 - \sqrt{1 + \delta^2 - \delta\sqrt{3}}\right) \cdot x_3 \leq 6\omega \\ & x_1, x_2, x_3 \geq 0 \end{aligned}$$

where we have used $d(u, w_1) = d(v, w_1) = \sqrt{\delta^2/4 + (1 - \sqrt{1 - \delta^2/4})^2}$ and $d(u, w_3) = d(v, w'_3) = \sqrt{1 + \delta^2 - \delta\sqrt{3}}$ (these expressions are obtained from repeated applications of Pythagoras' theorem starting from $d(u, v) = \delta$ and $d(u, w_3) = d(u, w'_3) = d(u, w_3) = d(u, w'_3) = 1$, and using that w_1 is at distance 1 from two specific points of $C_1(u)$ and $C_1(v)$ forming a rectangle with u and v). Let us denote by $f_5(\delta)$ the optimum of this linear program. We compute $f_5(\delta)$ for $\delta \in [\frac{1}{2}, 1]$ numerically using SageMath (see [7] for the source code); the function $\frac{f_5(\delta)}{\omega}$ is plotted in Fig. 5 (red dashed curve).

We now turn to our second upper bound on $\frac{1}{2}(\mu(D_1(u)) + \mu(D_1(v)))$. Consider the region $R_u \subset D_1(u)$ (bounded by the fat red dashed curve) and the region $R_v \subset D_1(v)$ (bounded by the fat blue curve) in Fig. 6. The region R_u is defined as




Fig. 6. Two adjacent vertices u and v in the second part of the proof of Lemma 4.1.

follows: take 6 points $z_1, \dots, z_6$ appearing in clockwise order on $C_1(u)$, such that any two consecutive points z_i, z_{i+1} (with indices modulo 6) form an equilateral triangle with u , and such that z_1 lies on the line (uv) , and the angle $\angle z_1 u v$ is equal to 0. Now define R_u as the union of the four Reuleaux triangles with vertices uz_2z_3 , uz_3z_4 , uz_4z_5 , and uz_5z_6 . The region R_v is defined as the symmetric of R_u with respect to the perpendicular bisector of the segment $[u, v]$. By definition, each of the regions R_u and R_v can be covered by four regions of diameter 1, and thus

$$\mu(R_u) \leq 4\omega \quad \text{and} \quad \mu(R_v) \leq 4\omega. \quad (5)$$

As before, by Lemma 2.3, we also have

$$\sum_{r \in [0,1]} (2-r) \mu(C_r(u)) \leq 6\omega \quad \text{and} \quad \sum_{r \in [0,1]} (2-r) \mu(C_r(v)) \leq 6\omega. \quad (6)$$

Note that $d(u) = \mu(D_1(u)) = \sum_{r \in [0,1]} \mu(C_r(u))$ and $d(v) = \mu(D_1(v)) = \sum_{r \in [0,1]} \mu(C_r(v))$.

We now define a number of points and a new region (see Fig. 6). Points w_1, w_2, w_3, w'_3 are defined similarly as in the first part of the proof, using R_u and R_v instead of S_u and S_v .

- w_1 is a point in $(D_1(u) \cap D_1(v)) \setminus (R_u \cup R_v)$, such that $\|u - w_1\| + \|v - w_1\|$ (the sum of the distances of w_1 to u and v) is maximized,
- w_2 is one of the two points in the intersection of $C_1(u)$ and $C_1(v)$,
- w_3 is a point of $(C_1(v) \cap R_u) \setminus R_v$, which is arbitrarily close to R_v ,
- w'_3 is the symmetric of w_3 with respect to the perpendicular bisector of the segment $[u, v]$,
- $w_4 \in C_1(v)$ and $w'_4 \in C_1(u)$ are symmetric with respect to the perpendicular bisector of the segment $[u, v]$, and the distance between w_4 and w'_4 is equal to 1,
- the region R_{uv} (depicted in green in Fig. 6) is the union of triangles uvw_4 and uvw'_4 , circular sectors $w_4vw_2 \subset D_1(v)$ and $w_2uw'_4 \subset D_1(u)$, and the symmetric region with respect to (uv) . For convenience we assume that R_{uv} does not contain w_2 .

By definition, $d(u, w'_4) = d(u, w_2) = d(w_4, w'_4) = d(v, w_4) = 1$, and thus each of the top and bottom halves of the region R_{uv} has diameter 1. Hence,

$$\mu(R_{uv}) \leq 2\omega. \quad (7)$$

The objective is again to maximize $\frac{1}{2}(\mu(D_1(u)) + \mu(D_1(v)))$ over all weighted point sets satisfying (5), (6), and (7). As before, we can assume that some optimal solution is symmetric with respect to the perpendicular bisector of the segment $[uv]$, that its support lies in the upper half-plane defined by (uv) and is included in the intersection of $D_1(u)$ and $D_1(v)$. If a symmetric weighted point set satisfies this property in addition to (5), (6), and (7), and contains two symmetric points z and z' distinct from w_1, w_2, w_3 , and w_4 (and their symmetric images), then z and z' can be moved locally so that $\|u - z\| + \|v - z\|$ and $\|u - z'\| + \|v - z'\|$ increase and thus no constraint is violated. Together with the symmetry of the constraints and objective function, this shows that we can assume without loss of generality that in some optimal solution, the support of the weighted point set is included in $\{w_1, w_2, w_3, w'_3, w_4, w'_4\}$ and moreover $\mu(\{w_3\}) = \mu(\{w'_3\})$ and $\mu(\{w_4\}) = \mu(\{w'_4\})$. The optimization problem can thus again be formulated as a finite linear program, as follows (here $x_1 = \mu(\{w_1\})$, $x_2 = \mu(\{w_2\})$, $x_3 = \mu(\{w_3\}) = \mu(\{w'_3\})$, and $x_4 = \mu(\{w_4\}) = \mu(\{w'_4\})$).

$$\begin{aligned}
& \max_{x_1, x_2, x_3, x_4} && x_1 + x_2 + 2x_3 + 2x_4 \\
& \text{s.t.} && x_2 + x_3 + x_4 \leq 4\omega \\
& && x_1 + 2x_3 \leq 2\omega \\
& && \left(2 - \sqrt{\frac{\delta^2}{4} + \left(\frac{\sqrt{3}}{2} - \sqrt{1 - \frac{(1+\delta)^2}{4}}\right)^2}\right) \cdot x_1 + x_2 + \left(3 - \sqrt{1 + \delta^2 - \delta}\right) \cdot x_3 \\
& && \quad + \left(3 - \sqrt{1 - \delta}\right) \cdot x_4 \leq 6\omega \\
& && x_1, x_2, x_3, x_4 \geq 0,
\end{aligned}$$

where we have used $d(u, w_1) = d(v, w_1) = \sqrt{\delta^2/4 + \left(\frac{\sqrt{3}}{2} - \sqrt{1 - (1+\delta)^2/4}\right)^2}$, $d(u, w_3) = d(v, w'_3) = \sqrt{1 + \delta^2 - \delta}$, and $d(u, w_4) = d(v, w'_4) = \sqrt{1 - \delta}$ (these expressions are again obtained from repeated applications of Pythagoras' theorem). Let us denote by $f_4(\delta)$ the optimum of this linear program. We compute $f_4(\delta)$ for $\delta \in [\frac{1}{2}, 1]$ numerically using SageMath (see [7] for the code); the function $\frac{f_4(\delta)}{\omega}$ is plotted in Fig. 5 (blue curve).

We can now check that when $\delta \in [\frac{1}{2}, 1]$, $\min\{f_4(\delta), f_5(\delta)\}$ is maximized when $\delta = 1$, with $f_5(1) = 5.6746\omega$ and $f_4(1) = 6\omega$ (note that $f_4(1/2) = 5.6698\omega < f_5(1)$). It follows that for any u, v with $\|u - v\| \in [\frac{1}{2}, 1]$, $\frac{1}{2}(d(u) + d(v)) \leq 5.6746\omega$, as desired. $\square$

This bound easily implies that there is an efficient distributed algorithm coloring G with at most $5.675\omega(G) + 1$ colors.

Theorem 1.5. *Every unit-disk graph G can be colored with at most $5.675\omega(G) + 1$ colors by a deterministic distributed algorithm in the LOCAL model, running in $O(\log^* n)$ rounds.*

Proof. Let $\omega = \omega(G)$, let A be the set of vertices of degree more than 5.675ω , and let B be the remaining vertices. We claim that any connected component of $G[A]$, the subgraph of G induced by A , is a clique. Indeed, any two adjacent vertices in $G[A]$ are at distance at most $\frac{1}{2}$ by Lemma 4.1 and thus if $G[A]$ contains a path uvw , then u and w are at distance $\frac{1}{2} + \frac{1}{2} \leq 1$, and so u and w are adjacent. Since $G[A]$ is a union of cliques, it can be colored with at most ω colors in $O(1)$ rounds (each connected component has diameter at most 2, and contains at most ω vertices). For each vertex $v \in B$, let $L(v)$ be the set of colors from $1, 2, \dots, 5.675\omega + 1$ that do not appear among the colored neighbors of v . Let us denote by $d_A(v)$ and $d_B(v)$ the number of neighbors of v in A and B , respectively. Note that for each $v \in B$, $|L(v)| \geq 5.675\omega + 1 - d_A(v) \geq d_B(v) + 1$, since $d_A(v) + d_B(v) \leq 5.675\omega$. Coloring the vertices of B from their lists $L(v)$ is thus an instance of the $(\deg + 1)$ -list coloring problem, which can be solved by a deterministic algorithm running in $O(\log^* n)$ rounds, by Corollary 2.2. The resulting coloring is a coloring of G with at most $5.675\omega + 1$ colors, as desired. $\square$

As a direct application of Lemma 4.1, we now obtain an improved upper bound on the average degree of any unit-disk graph.

Theorem 1.8. *Every unit-disk graph G has average degree at most $5.68\omega(G)$.*

Proof. Set $\epsilon = 6 - 5.675 = 0.325$ (all the computations in the proof are with respect to this specific choice of ϵ) and $\omega = \omega(G)$, and consider a fixed embedding of G in the plane. By Lemma 4.1, the average of the degrees of any two vertices u, v with $\|u - v\| \in [\frac{1}{2}, 1]$ is most $(6 - \epsilon)\omega$.

Each vertex v of G starts with a charge $w(v) = d(v)/\omega$, so that the average charge is precisely the average degree of G divided by ω . The charge is then moved according to the following rule: for any $\delta \geq 0$, each vertex with degree $(6 - \epsilon - \delta)\omega$ takes $\frac{\delta}{(6 - \epsilon - \delta)\omega}$ from the charge of each neighbor. For each $v \in V(G)$, let $w'(v)$ be the resulting charge of v . Note that the total charge has not changed and thus the average of $w'(v)$ over $v \in V$ is still the average degree of G divided by ω . We now prove that $w'(v) \leq 6 - \epsilon + 0.005 = 5.68$ for each vertex $v \in V$, which directly implies that G has average degree at most 5.68ω .

Consider first a vertex v of degree at most $(6 - \epsilon)\omega$. Then $d(v) = (6 - \epsilon - \delta)\omega$ for some $\delta \geq 0$. By the discharging rule, v takes $\frac{\delta}{(6 - \epsilon - \delta)\omega}$ from the charge of each of its $(6 - \epsilon - \delta)\omega$ neighbors, and might also give some of its charge to its neighbors. Thus $w'(v) \leq 6 - \epsilon - \delta + (6 - \epsilon - \delta)\omega \frac{\delta}{(6 - \epsilon - \delta)\omega} = 6 - \epsilon$.

Consider now a vertex v of degree more than $(6 - \epsilon)\omega$. Then $d(v) = (6 - \epsilon + \delta)\omega$ for some $0 < \delta \leq \epsilon$. By Corollary 2.4, $d_{1/2}(v) \leq 2(\epsilon - \delta)\omega$, and thus $D = D_1(v) \setminus D_{1/2}(v)$ contains at least $(6 - \epsilon + \delta)\omega - 2(\epsilon - \delta)\omega \geq (6 - 3\epsilon + 3\delta)\omega$ vertices.

By Lemma 4.1, each vertex of D has degree at most $(6 - \epsilon - \delta)\omega$. Observe that the function $x \mapsto \frac{x}{6 - \epsilon - x}$ is increasing for our choice of ϵ and for $x \in [0, \epsilon]$, thus each vertex of D takes at least $\frac{\delta}{(6 - \epsilon - \delta)\omega}$ from the charge of v (while v does not receive any charge from its neighbors). It follows that

$$w'(v) \leq 6 - \epsilon + \delta - (6 - 3\epsilon + 3\delta)\omega \frac{\delta}{(6 - \epsilon - \delta)\omega} = 6 - \epsilon + \frac{2\epsilon\delta - 4\delta^2}{6 - \epsilon - \delta} \leq 5.68,$$

for any $0 \leq \delta \leq \epsilon$ (for our choice of ϵ). $\square$

In the next section, we propose an approach to improve upon Theorem 1.8 (and get closer to Conjecture 1.7) using Fourier analysis.

5. Fourier analysis and the average degree of unit-disk graphs

Recall from the introduction that the disk clique number of a unit disk graph is the largest size of a clique contained within a disk of radius $1/2$ (note that this depends on the embedding of the graph in the plane). In this section, we will leverage Fourier-analytic techniques in order to investigate the ratio between the average degree of unit-disk graphs and their disk clique number. Since those tools are not standard in the graph theory literature, we first motivate their introduction informally.

A unit disk graph G is at its core simply a (multi)set of points in the plane and thus can be represented by a (discontinuous) function $f : \mathbb{R}^2 \rightarrow \mathbb{N}$. In this language, the degree of a vertex v can be readily computed as $d(v) := \int_{D_1(v)} f(u) du$, where the integral denotes the counting measure, i.e., a sum; and thus the sum of the degrees will be $\int_{\mathbb{R}^2} f(x) d(x) dx$, yielding average degree

$$\frac{\int_{\mathbb{R}^2} f(x) d(x) dx}{\int_{\mathbb{R}^2} f(x) dx} = \frac{\int_{\mathbb{R}^2} f(x) \int_{D_1(v)} f(u) dudx}{\int_{\mathbb{R}^2} f(x) dx}.$$

On the other hand, the disk clique number can also be formulated nicely, as the number of vertices in a disk of radius $1/2$ centered at x is simply $\int_{u \in D_{1/2}(x)} f(u) du$, and thus the disk clique number is $\max_{x \in \mathbb{R}^2} \int_{D_{1/2}(x)} f(u) du$.

Both the expressions of average degree and disk clique number can be expressed in a more compact way using the language of *convolutions*. The convolution product $f * g$ is defined as $(f * g)(x) = \int_{y \in \mathbb{R}^2} f(y) g(x - y)$. If we denote by $\chi_r : \mathbb{R}^2 \rightarrow \{0, 1\}$ the indicator function of the disk of radius $r > 0$ centered in $\mathbf{0} = (0, 0)$, we then have for the average degree the compact formulation $\int_{\mathbb{R}^2} f \cdot (f * \chi_1) / \|f\|_1$, while the disk clique number is $\|f * \chi_{1/2}\|_\infty$, where we have used the functional analysis norms $\|f\|_1 = \int |f|$ and $\|f\|_\infty = \max f$. Controlling the ratio between these two quantities is the aim of Question 5.1 below.

This functional formulation suggests a natural attempt at maximizing the ratio between the disk clique number and the average degree: we could try to look for a nonnegative function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^{\geq 0}$ (which we would then discretize) of fixed average $\|f\|_1$ and such that the unit disk average $f * \chi_1$ is constant, while f and hopefully $f * \chi_{1/2}$ are not (and thus the maximum value of the latter would be greater than the average value of f). In the one-dimensional case, any nonconstant 1-periodic function would do, but how to reason about this in two dimensions?

The analogy with periodic function in the one-dimensional case suggests that Fourier analysis might be helpful. Furthermore, the language of convolution products gives us a second hint that these quantities would be simpler to investigate in the Fourier domain: indeed, a key property of Fourier transforms is that they turn convolution products into usual products: $\widehat{f * g} = \widehat{f} \cdot \widehat{g}$. The Fourier transforms of the functions χ_r involve special functions called Bessel functions of the first kind, denoted by J_1 . In line with the one-dimensional case, where being 1-periodic amounts to having Fourier coefficients which are nonzero only at integer values, the analogue of 1-periodic functions that we are looking for will turn out to be functions whose Fourier coefficients are only nonzero when J_1 is zero. Discretizing such a function will bring us back to the realm of unit disk graphs, yielding examples with interesting properties.

In the remainder of this section, we formalize this idea properly, leading to a proof of Theorem 1.9. In order to be rigorous, some parts require analytical tools, for example the framework of *tempered distributions*, for which we refer to standard analysis textbooks, e.g., Rudin [17]. But the graph-theoretically minded reader can safely disregard these analytical issues and read through the text using the analogies that we have just described, thinking of a function as merely a (continuous version) of a unit disk graph, and of an integral as a sum.

5.1. The question

A function that will be crucial in the remainder is $\chi_r : \mathbb{R}^2 \rightarrow \{0, 1\}$, the indicator function of the disk of radius $r > 0$ centered in $\mathbf{0} = (0, 0)$.

$$\chi_r(x) = \begin{cases} 1 & \text{if } \|x\| \leq r, \\ 0 & \text{otherwise.} \end{cases}$$

We will be mostly interested in $\chi_{1/2}$ and χ_1 .

Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}^{\geq 0}$ be a Lebesgue-integrable function. As $f(x) = |f(x)|$ for any $x \in \mathbb{R}^2$, we have

$$\int_{\mathbb{R}^2} f(x) dx = \int_{\mathbb{R}^2} |f(x)| dx = \|f\|_1.$$

We will be interested in integrating f on disks of radius $r > 0$ centered in points x . As explained above, this can be described easily with a convolution product.

$$\int_{y \in D_r(x)} f(y) dy = \int_{\|x-y\| \leq r} f(y) dy = \int_{\mathbb{R}^2} f(y) \chi_r(y-x) dy = (f * \chi_r)(x),$$

where $D_r(x)$ denotes the disk of radius r centered in x , and $f * \chi_r$ denotes the convolution product of f and χ_r .

Question 5.1. What is the minimum constant $c > 0$ such that for any Lebesgue-integrable function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^{\geq 0}$,

$$\mu_f := \int_{\mathbb{R}^2} f(x) \cdot (f * \chi_1)(x) dx \leq c \cdot \|f\|_1 \cdot \|f * \chi_{1/2}\|_\infty?$$

Note that $D_1(x)$ can be covered by a constant number of disks of radius $\frac{1}{2}$, so there is a constant $c > 0$ such that $\|f * \chi_1\|_\infty \leq c \cdot \|f * \chi_{1/2}\|_\infty$, and thus the question above is well defined.

Take $R = [0, N]^2$ and let $\mathbf{1}_R$ denote the indicator function of R . Note that $\|\mathbf{1}_R\|_1 = N^2$ and $\|\mathbf{1}_R * \chi_{1/2}\|_\infty = \pi/4$ for $N \geq 1$. On the other hand, for all $x \in [1, N-1]^2$, $(\mathbf{1}_R * \chi_1)(x) = \pi$ and thus $\mu_f = (\pi - o(1))N^2 = (4 - o(1))\|\mathbf{1}_R\|_1 \cdot \|\mathbf{1}_R * \chi_{1/2}\|_\infty$, as $N \rightarrow \infty$. This shows that $c \geq 4$ in Question 5.1. We will see in Section 5.4 a finer example that shows that $c \geq 4.0905$ in Question 5.1.

5.2. Application to unit-disk graphs

Let us explain the connection to our original problem of bounding the average degree of unit-disk graphs as a function of their clique number.

Take a unit-disk graph G embedded in the plane, and consider the associated (finite) point multiset $P = V(G)$ in $\mathbb{R}^2$. By translation, we can assume that $P \subseteq [0, N]^2$, for some real $N > 0$. Fix some integer $M \gg N$, set $\delta := N/M$; and divide $[0, N]^2$ into M^2 squares R_{ij} (for $i, j \in \{1, \dots, M\}$) of size $\delta \times \delta$ (and area δ^2). Define the following function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^{\geq 0}$.

$$f(x) = \begin{cases} 0 & \text{if } x \notin [0, N]^2, \\ \frac{1}{\delta^2} |P \cap R_{ij}| & \text{if } x \in R_{ij}. \end{cases}$$

In words, $f(x)$ is the number of points of P in the square containing x , scaled by a factor δ^2 .

Note that for any $i, j \in \{1, \dots, M\}$,

$$\int_{R_{ij}} f(x) dx = \delta^2 \cdot \frac{1}{\delta^2} |P \cap R_{ij}| = |P \cap R_{ij}|.$$

It follows that $\|f\|_1 = |P| = |V(G)|$. More generally for any fixed disk $D \subseteq \mathbb{R}^2$ of positive radius, $\int_D f(x) dx \rightarrow |P \cap D|$ as $\delta \rightarrow 0$ (or equivalently, as $M \rightarrow \infty$).

As any disk of radius $\frac{1}{2}$ has diameter at most 1, the points of P lying in such a disk form a clique in G , and thus P intersects every such clique in at most $\omega(G)$ points. It follows that $\|f * \chi_{1/2}\|_\infty \leq (1 + o(1))\omega(G)$ (where the $o(1)$ term is with respect to $\delta \rightarrow 0$, or equivalently, as $M \rightarrow \infty$).

For sufficiently small $\delta > 0$, for any point $p \in P$ and $\delta \times \delta$ square R containing p , $\int_R f(x)(f * \chi_1)(x) dx$ counts the number of points of P (possibly with repetitions) coinciding with p , times the number of points in the disk of radius 1 centered in p (which is the same as the number of neighbors of p in G). It follows that as $\delta \rightarrow 0$, $\mu_f \rightarrow \sum_{v \in V(G)} d(v)$, the sum of the degrees of G . So for any constant $c > 0$ answering Question 5.1, we have

$$\sum_{v \in V(G)} d(v) \leq c \cdot |V(G)| \cdot (1 + o(1)) \cdot \omega(G),$$

and thus the average degree of G is at most $(c + o(1)) \cdot \omega(G)$.

What about the reverse direction? Given some function f showing that the answer to Question 5.1 is at least c_0 , for some $c_0 > 0$, we can translate this into a discrete distribution of points in the plane as follows. We consider a sufficiently large grid of dimension $N \times N$, and sufficiently small step $\delta > 0$, and for some sufficiently large real $k > 0$ we place $\lceil kf(x) \rceil$ points of P at each point x of the grid. If f behaves well (for instance if f is continuous and periodic), the ratio between the average degree and the disk clique number (recall that this is the maximum number of points in a disk of radius $\frac{1}{2}$) of the associated unit-disk graph G will be close to c_0 . However it might be possible that maximum cliques in the unit-disk graph G do not come from disks of radius $\frac{1}{2}$, but other shapes of diameter 1 (see Section 5.4). In this case the average degree of G is not necessarily close to $c_0 \cdot \omega(G)$.

5.3. Tools

Given a Lebesgue-integrable function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, the Fourier transform of f is given by

$$\hat{f}(\xi) = \int_{\mathbb{R}^2} f(x) e^{-2\pi i x \cdot \xi} dx,$$

where $x \cdot \xi$ denotes the dot product of x and ξ . Note that for any function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^{\geq 0}$,

$$\hat{f}(\mathbf{0}) = \int_{\mathbb{R}^2} f(x) dx = \int_{\mathbb{R}^2} |f(x)| dx = \|f\|_1.$$

When $\hat{f}$ is also integrable, we have that the reverse equality

$$f(x) = \int_{\mathbb{R}^2} \hat{f}(\xi) e^{2\pi i x \cdot \xi} d\xi$$

holds almost everywhere (and it holds everywhere if f is continuous).

It is well known (see for example [16, Example 17.3.1]) that for any $\xi \in \mathbb{R}^2$,

$$\hat{\chi}_1(\xi) = \frac{J_1(2\pi \|\xi\|)}{\|\xi\|},$$

where J_1 denotes the Bessel function of the first kind. Note that $J_1(x) \sim x/2$ when $x \rightarrow 0$, so $\hat{\chi}_1(\mathbf{0}) = \pi$ is well defined by continuity. Observe that for $r > 0$ $\chi_r(x) = \chi_1(x/r)$ for every $x \in \mathbb{R}^2$, and thus by the time-scaling property of Fourier transforms, we have

$$\hat{\chi}_r(\xi) = r^2 \hat{\chi}_1(r\xi) = \frac{r J_1(2\pi r \|\xi\|)}{\|\xi\|}$$

for any $\xi \in \mathbb{R}^2$ (as above, $\hat{\chi}_r(\mathbf{0}) = \pi r^2$ is well defined by continuity).

An important property of the convolution product is that it behaves well under Fourier transforms. For any $\xi \in \mathbb{R}^2$ and $r > 0$,

$$(\widehat{f * \chi_r})(\xi) = \hat{f}(\xi) \hat{\chi}_r(\xi) = \hat{f}(\xi) \frac{J_1(2\pi r \|\xi\|)}{\|\xi\|}.$$

So we have

$$(f * \chi_{1/2})(x) = \int_{\mathbb{R}^2} \hat{f}(\xi) \frac{J_1(\pi \|\xi\|)}{\|\xi\|} e^{2\pi i x \cdot \xi} d\xi \leq \|f * \chi_{1/2}\|_\infty,$$

for any $x \in \mathbb{R}^2$.

If we denote by $\mathbf{1}$ the constant function with $\mathbf{1}(x) = 1$ for any $x \in \mathbb{R}^2$, then $\hat{\mathbf{1}}(\xi) = \delta(\xi)$, where δ denotes the Dirac delta function.¹ In this case the equality above tells us that for any $x \in \mathbb{R}^2$,

$$(\mathbf{1} * \chi_r)(x) = \int_{\mathbb{R}^2} \delta(\xi) \frac{r J_1(2\pi r \|\xi\|)}{\|\xi\|} e^{2\pi i x \cdot \xi} d\xi = \hat{\chi}_r(\mathbf{0}) = \pi r^2,$$

which is just a complicated way to say that the area of a disk of radius r is πr^2 .

For f and g square-integrable, the Parseval formula (see, e.g., [17, Section 7.9]) stipulates that

$$\int_{\mathbb{R}^2} \overline{\hat{f}(x)} \hat{g}(x) dx = \int_{\mathbb{R}^2} \overline{\hat{f}(\xi)} \hat{g}(\xi) d\xi.$$

Assuming that f is square-integrable, then so is $f * \chi_r$ (this follows from the Minkowski integral inequality [10, Inequality 202]), and thus we have

$$\int_{\mathbb{R}^2} \overline{\hat{f}(x)} \cdot (f * \chi_r)(x) dx = \int_{\mathbb{R}^2} \overline{\hat{f}(\xi)} \cdot (\widehat{f * \chi_r})(\xi) d\xi$$

¹ Dirac delta functions can be formally defined using distributions or generalized functions, see for example Rudin [17, Chapter 6].

for any $r > 0$. If moreover, f is a real function, it follows that

$$\mu_f = \int_{\mathbb{R}^2} f(x) \cdot (f * \chi_1)(x) dx = \int_{\mathbb{R}^2} \overline{\hat{f}(\xi)} \cdot \hat{f}(\xi) \frac{J_1(2\pi \|\xi\|)}{\|\xi\|} d\xi = \int_{\mathbb{R}^2} \|\hat{f}(\xi)\|^2 \cdot \frac{J_1(2\pi \|\xi\|)}{\|\xi\|} d\xi$$

On the other hand, we have

$$\int_{\mathbb{R}^2} \|\hat{f}(\xi)\|^2 \cdot \frac{J_1(\pi \|\xi\|)}{2\|\xi\|} d\xi = \int_{\mathbb{R}^2} f(x) \cdot (f * \chi_{1/2})(x) dx \leq \|f * \chi_{1/2}\|_\infty \int_{\mathbb{R}^2} f(x) dx \leq \|f * \chi_{1/2}\|_\infty \cdot \|f\|_1$$

So if we had $J_1(2\xi) \leq \frac{c}{2} \cdot J_1(\xi)$ for any $\xi \geq 0$, and for some (hopefully small) constant $c > 0$, this would directly imply $\mu_f \leq c \|f\|_1 \cdot \|f * \chi_{1/2}\|_\infty$. However, the former does not hold, as the two functions oscillate independently.

An important observation here (in connection with the next section), is that for any $z \in [0, 1.18]$, $J_1(2\pi z) \leq 2 \cdot J_1(\pi z)$, with equality only if $z = 0$. This shows that in the disk $D_{1.18}(\mathbf{0})$ of radius 1.18 centered in $\mathbf{0}$,

$$\int_{D_{1.18}(\mathbf{0})} \|\hat{f}(\xi)\|^2 \cdot \frac{J_1(2\pi \|\xi\|)}{\|\xi\|} d\xi \leq 4 \cdot \int_{D_{1.18}(\mathbf{0})} \|\hat{f}(\xi)\|^2 \cdot \frac{J_1(\pi \|\xi\|)}{2\|\xi\|} d\xi,$$

so if the support of $\hat{f}$ lies in $D_{1.18}(\mathbf{0})$, we have $\mu_f \leq 4 \cdot \|f * \chi_{1/2}\|_\infty \cdot \|f\|_1$, with equality if and only if the support of $\hat{f}$ is $\{\mathbf{0}\}$, which is equivalent to say that f is a constant function. So if we want to find examples where $\mu_f > 4 \cdot \|f * \chi_{1/2}\|_\infty \cdot \|f\|_1$, we need to make sure that the support of $\hat{f}$ intersects the complement of $D_{1.18}(\mathbf{0})$.

5.4. A finer example

Consider the function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^{\geq 0}$ defined by $f((x, y)) := 1 + \sin(2Bx)$, where B is the first positive zero of the Bessel function J_1 , that is B is the smallest positive real such that $J_1(B) = 0$. It is known that $B \approx 3.8317$. We will use Fourier analysis on f , which is definitely not Lebesgue-integrable. Nevertheless, this can be justified using the framework of *tempered distributions*, to which f belongs since it is bounded. For the sake of readability, we do not enter these technical details and refer the reader to standard textbooks on distributions, e.g., Rudin [17, Chapters 6 and 7].

Take some real number $N = k \cdot \frac{\pi}{B}$, for some integer $k \geq 2$ (note that $\frac{\pi}{B}$ is the period of f), and let $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined as $g(x) = f(x) \cdot \mathbf{1}_{[0, N]^2}(x)$. Note that $\|g\|_1 = N^2$ (since N is a multiple of the period of g), and $\|g * \chi_{1/2}\|_\infty = \|f * \chi_{1/2}\|_\infty$ (as g is periodic and we have chosen N large enough). As we have seen above, the value of $\|f * \chi_{1/2}\|_\infty$ can be computed using

$$(f * \chi_{1/2})(x) = \int_{\mathbb{R}^2} \hat{f}(\xi) \frac{J_1(\pi \|\xi\|)}{2\|\xi\|} e^{2\pi i x \cdot \xi} d\xi.$$

In order to use this equality, we observe that for any $\xi \in \mathbb{R}^2$,

$$\hat{f}(\xi) = \delta(\xi) + \frac{1}{2i}(\delta(\xi - \frac{2B}{2\pi}) - \delta(\xi + \frac{2B}{2\pi})),$$

where δ denotes the Dirac delta function. This shows that $\hat{f}(\xi) = 0$ unless $\xi = \mathbf{0}$ or $\xi = \pm \frac{B}{\pi}$ (we observe, in connection with the final comment of the previous section, that $\frac{B}{\pi} \approx 1.22 > 1.18$). If $\xi = \pm \frac{B}{\pi}$, then $J_1(\pi \|\xi\|) = 0$, by definition of B . It follows that the only non-zero term in the integral is for $\xi = \mathbf{0}$, where $\frac{J_1(\pi \|\xi\|)}{2\|\xi\|} e^{2\pi i x \cdot \xi} = \frac{\pi}{4}$. By the definition of the Dirac delta function, the integral evaluates as

$$(f * \chi_{1/2})(x) = \int_{\mathbb{R}^2} \hat{f}(\xi) \frac{J_1(\pi \|\xi\|)}{2\|\xi\|} e^{2\pi i x \cdot \xi} d\xi = \frac{\pi}{4}.$$

This shows that $\|g * \chi_{1/2}\|_\infty = \|f * \chi_{1/2}\|_\infty = \frac{\pi}{4}$.

It remains to evaluate

$$\mu_g = \int_{\mathbb{R}^2} g(x) \cdot (g * \chi_1)(x) dx = \int_{[0, N]^2} f(x) \cdot (g * \chi_1)(x) dx.$$

We can compute μ_g for N a multiple of π/B by using the Parseval formula as before, this time to f and $g * \chi_1$. Note that f is not square-integrable in this case, but it is a tempered distribution and g has compact support and is thus a test function, so we can still use Parseval formula [17, Chapter 7] in this case (we omit the definitions of tempered distributions and test functions here, the reader is referred to [17] for more details). We first compute that $\hat{g}(\mathbf{0}) = N^2 = \|g\|_1$, $\hat{g}(\frac{2B}{2\pi}) = \frac{N^2}{2i}$ and $\hat{g}(-\frac{2B}{2\pi}) = \frac{-N^2}{2i}$. Now:

$$\begin{aligned}
\mu_g &= \int_{[0,N]^2} f(x) \cdot (g * \chi_1)(x) dx \\
&= \int_{\mathbb{R}^2} f(x) \cdot (g * \chi_1)(x) dx - \int_{[-1,N+1]^2 \setminus [0,N]^2} f(x) \cdot (g * \chi_1)(x) dx \\
&= \int_{\mathbb{R}^2} f(x) \cdot (g * \chi_1)(x) dx - O(N) \\
&= \int_{\mathbb{R}^2} \overline{\hat{f}}(\xi) \cdot \hat{g}(\xi) \cdot \frac{J_1(2\pi \|\xi\|)}{\|\xi\|} d\xi - O(N) \\
&= \int_{\mathbb{R}^2} \left(\delta(\xi) + \frac{1}{2i} \left(\delta(\xi + \frac{2B}{2\pi}) - \delta(\xi - \frac{2B}{2\pi}) \right) \right) \cdot \hat{g}(\xi) \cdot \frac{J_1(2\pi \|\xi\|)}{\|\xi\|} d\xi - O(N) \\
&= \int_{\mathbb{R}^2} \|g\|_1 \left(\lim_{\xi \rightarrow 0} \left(\frac{J_1(2\pi \|\xi\|)}{\|\xi\|} \right) + \frac{1}{4} \cdot \frac{\pi J_1(2B)}{B} + \frac{1}{4} \cdot \frac{\pi J_1(-2B)}{B} \right) - O(N) \\
&= \pi \|g\|_1 \left(1 + \frac{J_1(2B)}{2B} \right) - O(N),
\end{aligned}$$

where we used that $f(x) \cdot (g * \chi_1)(x)$ is zero outside of $[-1, N+1]^2$ and is always at most 2, that $\lim_{\xi \rightarrow 0} \frac{J_1(2\pi \|\xi\|)}{\|\xi\|} = \pi$ and that the J_1 function is even. Now, $\frac{J_1(2B)}{2B} \approx 0.04527$ and $\|g\|_1 = N^2$, and thus we obtain that for N large enough,

$$\mu_g \approx 4.0905 \cdot \|g\|_1 \cdot \|g * \chi_{1/2}\|_\infty.$$

Therefore we have $c \geq 4.0905$ in Question 5.1. Recall that the disk clique number of a unit-disk graph G , denoted by $\omega_D(G)$, is the largest size of a set of points contained within a disk of radius $\frac{1}{2}$ in G . Discretizing the function g introduced above as described in Section 5.2 yields the following theorem.

Theorem 1.9. *There exists a unit-disk graph G of average degree at least $4.0905 \omega_D(G)$.*

Proof. We consider a 2-dimensional grid of dimension $N \times N$ and step $\delta > 0$, and we place $\lceil kg(x) \rceil$ points at each point x of the grid for the function $g(x) := (1 + \sin(2Bx)) \cdot \mathbf{1}_{[0,N]^2}(x)$. This set of points defines a unit-disk graph G . As explained in Section 5.2, for N and k large enough and δ small enough, the ratio between the average degree of G and its disk clique number $\omega_D(G)$ will converge to $\frac{\|g \cdot (g * \chi_1)\|_1}{\|g\|_1 \|g * \chi_{1/2}\|_\infty}$. By the calculations above, this ratio is larger than 4.0905 for sufficiently large N , thus establishing Theorem 1.9. $\square$

Note however that (as alluded to in Section 5.2), the unit-disk graph G obtained from this construction does not contradict Conjecture 1.7. Indeed, numerical computations² suggest that some smooth versions of Reuleaux triangles are slightly denser than disks of radius $\frac{1}{2}$ with respect to f (and g), and in fact these computations seem to indicate that the ratio between the average degree and the clique number in the corresponding unit-disk graph G is close to 3.93 (see Fig. 7 for a picture of a maximum clique with respect to this distribution, with density significantly larger than $\pi/4$).

6. Conclusion

Given a sequence of pairs $\mathcal{P} = (p_i, r_i)_{1 \leq i \leq n}$, where each p_i is a point in the plane and each r_i is a positive real, the *disk graph* on $\mathcal{P}$ is the graph with vertex set $\{p_1, p_2, \dots, p_n\}$, in which two vertices p_i and p_j are adjacent if and only if $\|p_i - p_j\| \leq r_i + r_j$. Disk graphs model wireless communication networks using omni-directional antennas of possibly different powers (where the power of the i -th antenna is proportional to the real number r_i), and disk graphs where all the reals r_i are equal are precisely unit-disk graphs. A natural question is whether results similar to the results we obtain here can be proved for disk graphs. A major difference is that disk graphs do not have their maximum degree bounded by a function of their clique number ω (the class of disk graphs contains the class of all trees for instance). However disk graphs have average degree bounded by a constant times ω , and this can be used to obtain $O(\log n)$ round algorithms coloring these graphs with few colors [2]. In the same spirit as Question 1.6 we can ask the following.

² As the code is not easily readable, due to several (probably necessary) optimization tricks, we chose not to make it publicly available, however we are happy to send it to any interested reader upon request. The computation boils down to solving a max-flow instance in a fairly dense graph on $\sim 10^4$ vertices.




Fig. 7. The points of a maximum clique with respect to $f : (x, y) \mapsto 1 + \sin(2Bx)$. The function f is plotted in red, but not on the same scale as the points of the clique (in blue).

Question 6.1. What is the minimum real $c > 0$ such that a coloring of any n -vertex disk graph G with $c \cdot \omega(G)$ colors can be obtained in $O(\log n)$ rounds in the LOCAL model?

As a final comment, we recall that in the proceedings version of the paper [6] we claimed that for any $\epsilon > 0$, any unit-disk graph G could be colored with $(3 + \epsilon)\omega(G)$ colors in $O(1/\epsilon)$ rounds in the location-aware LOCAL model, but that we later found an error in our original argument, and could only replace $(3 + \epsilon)\omega(G)$ by $4\omega(G)$ (in Theorem 1.2). A problem left open by this work is to prove (or disprove) that $(3 + \epsilon)\omega(G)$ colors are sufficient. A related (purely existential) question is the following. A *flat annulus*, or *cylinder*, of height h and circumference ℓ is the metric space obtained from a Euclidean rectangle of height h and length ℓ by identifying its left and right boundaries. While a rectangle is the Cartesian product of two segments, a flat annulus can be thought of as the product of a cycle with a segment.

Question 6.2. Let $\epsilon > 0$. Let G be a unit-disk graph G whose vertices are embedded in a flat annulus of height $\sqrt{3}/2$ and sufficiently large circumference (as a function of ϵ). Is it true that if $\omega(G)$ is sufficiently large, $\chi(G) \leq (1 + \epsilon)\omega(G)$?

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The first part of the code is available in the arXiv version of the paper, and the second part of the code is available upon request.

Acknowledgement

The authors would like to thank Wouter Cames van Batenburg and François Pirot for the interesting discussions. The authors would also like to express their gratitude to the reviewers of the conference and journal versions of the paper for their helpful comments and suggestions, and to Mohsen Ghaffari for his kind explanations on the status of the $(\deg + \epsilon\Delta)$ -list coloring and $(\deg + 1)$ -list coloring problems.

References

- [1] A. Balliu, F. Kuhn, D. Olivetti, Improved distributed fractional coloring algorithms, in: Proc. of the 25th International Conference on Principles of Distributed Systems (OPODIS 2021), in: LIPIcs, vol. 217, 2021, 18.
- [2] L. Barenboim, M. Elkin, Sublogarithmic distributed MIS algorithm for sparse graphs using Nash-Williams decomposition, Distrib. Comput. 22 (5–6) (2010) 363–379.

- [3] L. Bieberbach, Über eine Extremaleigenschaft des Kreises, Jahresber. Dtsch. Math.-Ver. 24 (1915) 247–250.
- [4] N. Bousquet, L. Esperet, F. Pirot, Distributed algorithms for fractional coloring, in: Proc. of the 28th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2021), 2021.
- [5] B.N. Clark, C.J. Colbourn, D.S. Johnson, Unit disk graphs, Discrete Math. 86 (1–3) (1990) 165–177.
- [6] L. Esperet, S. Julliot, A. de Mesmay, Distributed coloring and the local structure of unit-disk graphs, in: Algorithms for Sensor Systems (ALGOSENSORS), in: Lecture Notes in Computer Science, vol. 12961, 2021.
- [7] L. Esperet, S. Julliot, A. de Mesmay, Distributed coloring and the local structure of unit-disk graphs, arXiv:2106.12322, 2021.
- [8] S. Gerke, C.J.H. McDiarmid, Graph imperfection, J. Comb. Theory, Ser. B 83 (1) (2001) 58–78.
- [9] A. Gräf, M. Stumpf, G. Weißenfels, On coloring unit disk graphs, Algorithmica 20 (1998) 277–293.
- [10] G.H. Hardy, J. Littlewood, G. Pólya, Inequalities, Cambridge University Press, 1952.
- [11] M. Hassinen, J. Kaasinen, E. Kranakis, V. Polishchuk, J. Suomela, A. Wiese, Analysing local algorithms in location-aware quasi-unit-disk graphs, Discrete Appl. Math. 159 (15) (2011) 1566–1580.
- [12] F. Kuhn, The Price of Locality: Exploring the Complexity of Distributed Coordination Primitives, PhD thesis, ETH, Zurich, 2005.
- [13] N. Linial, Locality in distributed graph algorithms, SIAM J. Comput. 21 (1992) 193–201.
- [14] E. Malesinška, S. Piskorz, G. Weißenfels, On the chromatic number of disk graphs, Networks 32 (1) (1998) 13–22.
- [15] R. Peeters, On Coloring j -Unit Sphere Graphs, FEW, vol. 512, Department of Economics, Tilburg University, 1991.
- [16] A.D. Poularikas, The Handbook of Formulas and Tables for Signal Processing, CRC Press LLC, Boca Raton, Florida, 1999.
- [17] W. Rudin, Functional Analysis, McGraw-Hill Book Company, 1973.
- [18] E.R. Scheinerman, D.H. Ullman, Fractional Graph Theory: A Rational Approach to the Theory of Graphs, Dover Publications, 2013.
- [19] J. Schneider, R. Wattenhofer, An optimal maximal independent set algorithm for bounded-independence graphs, Distrib. Comput. 22 (2010) 349–361.
- [20] A. Scott, P. Seymour, A survey of χ -boundedness, J. Graph Theory 95 (2020) 473–504.
- [21] J. Suomela, Survey of local algorithms, ACM Comput. Surv. 45 (2) (2013) 24.
- [22] A. Wiese, Local approximation algorithms in unit disk graphs, Master's thesis, Technische Universität Berlin, 2007.
- [23] A. Wiese, E. Kranakis, Local construction and coloring of spanners of location aware unit disk graphs, Discrete Math. Algorithms Appl. 1 (4) (2009) 555–588.