Contributions to Discrete Mathematics

Volume 8, Number 2, Pages 60–74
ISSN 1715-0868

COLOURING STABILITY TWO UNIT DISK GRAPHS

HENNING BRUHN

Abstract. We prove that every stability two unit disk graph has chromatic number at most 3/2 times its clique number.

1. Introduction

A unit disk graph (or UDG for short) is defined on a point set in the plane, where two points are considered as adjacent vertices if their distance is at most one. As a very basic model for wireless devices, unit disk graphs have attracted quite a lot of interest and there exists an extensive literature on the subject. One of the earliest applications is due to Hale [9] who considered them in the context of the frequency assignment problem. There, the task is to assign different frequencies to the wireless devices in order for them to communicate with a base station without interference. An edge between two points signifies the devices being close enough to interfere with each other, so that an assigned frequency corresponds to a stable set in the graph. The frequency assignment problem then becomes a graph colouring problem, and naturally, it is desired to assign as few frequencies as possible.

In this article, I will treat the colouring of a unit disk graph from a structural point of view. The chromatic number of unit disk graphs can be bounded above in terms of the clique number. The best known bound is due to Peeters:

Theorem 1.1 (Peeters [15])**.** *A unit disk graph G can be coloured with at most $3\omega(G)-2$ colours.*

How good is this bound? Malesińska, Piskorz and Weißenfels [11] gave the following construction, which has the highest ratio of χ and ω among all known unit disk graphs. Consider the class of graphs C_n^k on vertex set $\{0,\dots,n-1\}$ where i and j are adjacent if $|i-j|\leq k-1$ (where we calculate mod n). Observe that we may realize any C_n^k as a unit disk graph by placing the vertices at equidistant points on a circle of appropriate radius.

Received by the editors November 11, 2011, and in revised form July 24, 2013.
2010 *Mathematics Subject Classification.* 05C15, 05C62.
Key words and phrases. Unit disk graphs, colouring, clique number, stability.

Moreover, we see that C^k_{3k-1} has stability $\alpha = 2$ and clique number $\omega = k$, from which we deduce for the chromatic number that

$$\chi(C^k_{3k-1}) \geq \frac{3k-1}{2} = \frac{3}{2}\omega - \frac{1}{2}.$$

Clearly, between an upper bound of essentially 3ω and the $3\omega/2$ of the Malesińska et al. example there is a sizable gap. In this article, we will close this gap for unit disk graphs of stability at most two:

Theorem 1.2. *A unit disk graph G with $\alpha(G) \leq 2$ can be coloured with at most $3\omega(G)/2$ colours.*

The theorem shows that, at the very least, the example of Malesińska et al. cannot be easily improved. Moreover, I contend that it gives some evidence for believing that the true bound is closer to $3\omega/2$ than to Peeters' bound of $3\omega - 2$.

There is some more evidence for this belief. There are two classes of UDGs, for which it is known that their chromatic number is bounded by $3\omega/2$. The first of these is the class of triangle-free UDGs: A triangle-free UDG is planar, see Breu [2], and thus by Grötzsch' theorem, is 3-colourable [8]. The second class consists of augmentations of induced subgraphs of the triangular lattice in the plane. McDiarmid and Reed [14] show that these can be coloured with at most $(4\omega + 1)/3$ colours. Another piece of evidence is provided by McDiarmid [13], who investigates fairly general models of random unit disk graphs. In that context, it turns out that with high probability, the chromatic number is very close to the clique number. Finally, considering fractional instead of ordinary colourings, Gerke and McDiarmid [6] prove that the fractional chromatic number is bounded by $2.2\omega(G)$ for any unit disk graph G.

Optimisation problems in UDGs, and in particular, colouring UDGs algorithmically have attracted some attention. We just mention Marathe et al [12] who give a 3-approximation colouring algorithm and the result by Clark, Colbourn and Johnson [4] that 3-colourability remains NP-complete for UDGs. We refer to Balasundaram and Butenko [1] for a survey on several optimisation problems in UDGs.

The paper is organised as follows. After briefly stating some of the basic definitions that we are going to use we will proceed with the proof of our main result, Theorem 1.2, in Section 3. The key lemma on which the proof of Theorem 1.2 rests will be deferred to Section 5. In Section 4 we will discuss which geometric insights will be exploited.

2. Definitions

For general graph-theoretic notation and concepts we refer to Diestel [5]. Let G be a graph. A *clique* of G is a subgraph in which any two vertices are adjacent. A *stable set* of G is a subgraph or vertex set so that no two vertices are adjacent. The size of the largest clique is denoted by $\omega(G)$, while the size of the largest stable set is $\alpha(G)$, the *stability of* G. Every *colouring* will be proper, that is, an assignment of colours so that no two adjacent vertices receive the same colour. We denote the chromatic number, the minimum number of colours to colour G, by $\chi(G)$, and define $\overline{\chi}(G)$, the *clique partition number*, to be the chromatic number of the completement of G. A vertex v is *complete to* some vertex set X if every vertex in X is adjacent to v. A vertex set U is *complete to* X is every vertex in U is complete to X.

A *unit disk* is a closed disk of radius 1 in the plane. Unit disk graphs can be represented in two ways: In the intersection model the vertices are unit disks in the plane and two of them are adjacent if and only if the disks intersect; and in the distance model, the vertex set is a point set in the plane, and any two vertices are adjacent if and only if their distance is at most 1. *We will work exclusively with the distance model.*

Moreover, we always see a unit disk graph as a concrete geometric object, that is, the vertex set is indeed a subset of points in $\mathbb{R}^2$. So, every vertex *is* a point in the plane. As a consequence we do not allow two vertices to be represented by the same point. It is not hard to check, however, that this is no restriction for our purposes: Our main theorem remains valid if this requirement is dropped.

For two points $x, y \in \mathbb{R}^2$, we denote the Euclidean distance in the plane by $\mathrm{dist}(x, y)$. If X is a point set in $\mathbb{R}^2$ then we let $\mathrm{conv}(X)$ be the convex hull of the points in X. As a shorthand we set $\mathrm{conv}(G) := \mathrm{conv}(V(G))$. We say that a line L *separates* a point p from a point set X if p lies in one of the closed half-planes defined by L, while X is contained in the other. For any set X and any x we use $X + x$ to denote $X \cup \{x\}$

3. Proof of main theorem

The proof Theorem 1.2 rests on the Gallai-Edmonds decomposition as well as a key lemma, which will be proved in the course of the following two sections.

Lemma 3.1. *Let G be a unit disk graph with $\alpha(G) \leq 2$. Then $V(G)$ can be partitioned into three cliques, two of which have different cardinalities.*

Let me remark that the lemma is motivated by the structure of the Malesińska et al example. Moreover, as a by-product, we obtain that a stability two unit disk graph has clique partition number $\overline{\chi} \leq 3$.

For the well-known Gallai-Edmonds decomposition, which we state below, we refer to Lovasz and Plummer [10]. We briefly recall the basic notions of matching theory. A *matching* of a graph G is a set of edges so that no two of its edges share an endvertex. The matching is *perfect* if every vertex is incident with a matching edge; and it is *near-perfect* if this is the case for every vertex except one. The graph G is *factor-critical* if $G - v$ has a perfect matching for every vertex v.

Theorem 3.2 (Gallai-Edmonds decomposition)**.** *For any graph G denote by Z the set of those vertices v for which there exists a maximum-size matching missing v. Let $X := (\bigcup_{z\in Z} N(z)) \setminus Z$, and $R := G - (Z \cup X)$. Let M be a maximum-size matching. Then:*

- *Every component of $G[Z]$ is factor-critical and M restricts to a near-perfect matching on every one of them.*
- *M restricted to R is a perfect matching of R.*
- *Every vertex x of X is incident with an edge $e_x \in M$ whose other endvertex lies in a component of $G[Z]$. No two such edges e_x and e_y are incident with the same component.*

The Gallai-Edmonds decomposition combined with Lemma 3.1 allow us to prove the main theorem:

Proof of Theorem 1.2. Let H be the complement of G. Let M be a maximum-size matching of H, let Z, R, X be as in the Gallai-Edmonds decomposition of H, and denote by $\mathcal{O}$ the set of components of Z. Then $M_R := M \cap E(R)$ is a perfect matching, $M_K := M \cap E(K)$ is a near-perfect matching for every $K \in \mathcal{O}$, and and every $x \in X$ is incident with an edge in M whose other endvertex lies in some $K \in \mathcal{O}$. Let us denote the set of matching edges incident with vertices in X by M_X, and let $\mathcal{O}_X$ be the set $K \in \mathcal{O}$ incident with an edge in M_X. Finally, set $\mathcal{O}' := \mathcal{O} \setminus \mathcal{O}_X$. Thus

$$M = M_R \cup M_X \cup \bigcup_{K\in\mathcal{O}} M_K,\ \mathcal{O} = \mathcal{O}_X \cup \mathcal{O}' \text{ and } |M_X| = |\mathcal{O}_X|.$$

By Lemma 3.1, R can be partitioned into three stable sets A_R, B_R, C_R, which we may choose so that $|A_R| \geq \max(|B_R|, |C_R|)$. Moreover, by the same lemma, every $K \in \mathcal{O}$ can be partitioned into three stable sets A_K, B_K, C_K so that the three sets do not have the same size. Choosing them with $|A_K| \geq |B_K| \geq |C_K|$ implies

$$|C_K| + 1 \leq |A_K|. \tag{3.1}$$

Observe that $A := A_R \cup \bigcup_{K\in\mathcal{O}} A_K$ is a stable set of H. Now, we see that

$$2|M_R| = |V(R)| = |A_R| + |B_R| + |C_R| \leq 3|A_R|,$$

while for every $K \in \mathcal{O}$ we obtain with (3.1)

$$2|M_K| = |V(K)| - 1 = |A_K| + |B_K| + |C_K| - 1 \leq 3|A_K| - 2.$$

From this it follows that

$$\begin{aligned}
2|M| &= 2|M_R| + 2|M_X| + 2\sum_{K\in\mathcal{O}} |M_K| \\
&\leq 3|A_R| + 2|M_X| + \sum_{K\in\mathcal{O}} (3|A_K| - 2) \\
&= 3|A| + 2|M_X| - 2|\mathcal{O}| \\
&= 3|A| - 2|\mathcal{O}'|.
\end{aligned}$$

The matching M together with the set of unmatched vertices, one for each $K \in \mathcal{O}'$, yields a clique partition of H of size $|M| + |\mathcal{O}'|$. Hence

$$2\overline{\chi}(H) \leq 2(|M| + |\mathcal{O}'|) \leq 3|A| - 2|\mathcal{O}'| + 2|\mathcal{O}'| \leq 3\alpha(H).$$

We deduce $\chi(G) \leq 3\omega(G)/2$, which finishes the proof. □

4. Basic geometric facts

Before beginning with the proof of the key lemma, let me collect in this section the basic geometric facts that we will need.

The geometry of UDGs is not linear. For example, the subset of points of distance at most 1 to a given set of vertices, which is the intersection of several unit disks, can be very complex indeed. Sometimes this inherent non-linearity can be avoided. That is, instead of exploiting a concrete realisation of the unit disk graph, it is sometimes possible to deduce the desired conclusion only by appealing to abstract properties shared by all UDGs. If this is possible it might even result in cleaner arguments. The MaxClique algorithm by Raghavan and Spinrad[16], for instance, is such an example.

Abstract properties of UDGs include the fact that a UDG may not contain any induced $K_{1,6}$ or that

(4.1) *the common neighbourhood of any two non-adjacent vertices induces a co-bipartite graph.*

This is an observation first made by Clark, Colbourn and Johnson [4]. Although (4.1) is a fairly powerful property, even in conjunction with stability $\alpha = 2$, it is not enough to guarantee a chromatic number of $\chi \leq 3\omega/2$. To see this, consider the following graph CS_k, which is a subgraph of a graph appearing in Chudnovsky and Seymour [3]. Let CS_k be defined on four disjoint cliques each of which is comprised of k vertices: $\{a_1, \ldots, a_k\}$, $\{b_1, \ldots, b_k\}$, $\{c_1, \ldots, c_k\}$ and $\{d_1, \ldots, d_k\}$. Additionally, for $i, j = 1, \ldots, k$ with $i \neq j$ we define the following adjacencies: Let a_i be adjacent with b_j and d_j and with c_i; let b_i be adjacent with c_j and a_j and with d_i; let c_i be adjacent with d_j and b_j and with a_i; and let d_i be adjacent with a_j and c_j and with b_i. All other pairs of vertices are non-adjacent.

Clearly, the stability of CS_k is equal to 2, and if $k \geq 3$ then $\omega(\mathrm{CS}_k) = k+1$ and $\chi(\mathrm{CS}_k) = 2k$. It is not entirely obvious but also not overly difficult to check that CS_k satisfies (4.1).

To sum up, directly exploiting the geometry of a UDG might be hard due to the inherent non-linearity, while the other approach of using only abstract properties appears to fail. So, what can be done? We will work with a concrete geometric realisation, that is, the vertices will have concrete positions in the plane, but we will in some sense linearise the adjacencies. To show that two given vertices are adjacent we will never try to calculate their distance directly but rather use the following two principles, Lemma 4.1 and Lemma 4.2, that are of a more combinatorial flavour.

Let us say that for distinct vertices u, v, x, y in a unit disk graph, the two edges uv and xy are *crossing* if $\mathrm{conv}(u, v)$ intersects $\mathrm{conv}(x, y)$.

Lemma 4.1 (Breu [2])**.** *Let u, v, x, y be four distinct vertices in a unit disk graph G. If uv and xy are crossing edges then $G[u, v, x, y]$ contains a triangle.*

Lemma 4.2. *Let a vertex v of a unit disk graph be adjacent to two vertices u and w. Then v is adjacent to x for every vertex $x \in \mathrm{conv}(u, v, w)$.*

Proof. The vertex v is adjacent to every vertex in the unit disk centered at v. This disk clearly contains $\mathrm{conv}(u, v, w)$. □

Having stated that we will work only with these two principles rather than with concrete distances, let me turn around and immediately violate that rule. This becomes necessary as we will need to distinguish two classes of unit disk graphs: Those that have two vertices that are far apart and those that fit into a small disk. This will be done in the next two lemmas—from then on, however, we will adhere to the rule.

Lemma 4.3. *Let G be a unit disk graph. Then either G has two vertices of distance greater than $\sqrt{3}$ or there is a unit disk that contains all of G.*

Proof. We assume that all pairs of vertices of G have distance at most $\sqrt{3}$. It is straightforward to see that there is a point x whose maximal distance to $V(G)$ is minimal. More formally, there is an $x \in \mathbb{R}^2$ so that setting $d := \max\{\mathrm{dist}(x, v) : v \in V(G)\}$ we obtain $\mathrm{dist}(y, v) \geq d$ for all points $y \in \mathbb{R}^2$ and all vertices v.

Let W be the set of vertices v with $\mathrm{dist}(x, v) = d$. We will see that $d \leq 1$, which means that all of G is contained in the unit disk with centre x. For this, we claim that

$$x \in \mathrm{conv}(W). \tag{4.2}$$

Suppose $x \notin \mathrm{conv}(W)$. Then there are two vertices w, w' in W, so that the line L through w and w' does not contain x and separates x from W, and so that $x \notin L$. Choose $\varepsilon > 0$ small enough so that $\mathrm{dist}(v, x) + \varepsilon < d$ for all $v \in V(G) \setminus W$. On the line through x that is orthogonal to L, let x' be the point between x and $\mathrm{conv}(W)$ of distance ε to x. To see that $\max\{\mathrm{dist}(x', v) : v \in V(G)\} < d$ consider any $u \in W$, which then lies on the segment between w and w' on a circle of radius d and centre x. The angle at x' in the triangle with vertices u, x', x is at least $\pi/2$, which means that $\mathrm{dist}(u, x') < \mathrm{dist}(u, x)$. On the other hand, for any $v \in V(G) \setminus W$ we also have $\mathrm{dist}(x', v) < d$ by choice of ε. Therefore, x' contradicts the minimal choice of x.

Now, if x is the convex combination of two vertices w_1 and w_2 in W then $2d = \mathrm{dist}(w_1, w_2) \leq \sqrt{3}$, which implies $d < 1$. If this is not the case, then, by Carathéodory's theorem, x lies in the convex hull of three vertices w_1, w_2, w_3 of W. There are i, j with $1 \leq i < j \leq 3$ so that the angle θ at x in the triangle with vertices x, w_i, w_j is at least $2\pi/3$ but no more than π by (4.2). Then we get

$$\frac{\sqrt{3}}{2} = \sin\left(\frac{\pi}{3}\right) \leq \sin\left(\frac{\alpha}{2}\right) = \frac{\mathrm{dist}(w_i, w_j)}{2d}.$$

Thus, $d \leq 1$ as $\mathrm{dist}(w_i, w_j) \leq \sqrt{3}$. □

The case when a unit disk graph has two vertices u, v of distance at least $\sqrt{3}$ is particularly easy. We will see that in this case we can get the key lemma with only a small effort.

Lemma 4.4. *Let G be a unit disk graph. If u and v are two vertices of distance at least $\sqrt{3}$ then $N(u) \cap N(v)$ is a clique.*

Proof. Consider any two vertices u and v for which $N(u) \cap N(v)$ contains two non-adjacent vertices. The two points of greatest distance in the intersection of the unit disk centered at u and the unit disk centered at v are the two points where the boundaries of the two unit disks meet. Let this distance be s; and observe that $s > 1$ as otherwise $N(u) \cap N(v)$ would be a clique. Now, we obtain

$$\mathrm{dist}(u, v) = \sqrt{4 - s^2} < \sqrt{3}.$$

□

Finally, let us get rid of the special case when the unit disk graph is contained in a line. While this case is easy—the graph then becomes a linear interval graph—it still leads to some unnecessary complications. From an observation of Gräf, Stumpf and Weißenfels [7] it follows that a unit disk graph can always be rescaled so that there is an $\varepsilon > 0$ so that moving any vertex by at most ε from its original position does not change the (abstract) graph. So, by perturbing the graph slightly we may always assume that

(4.3) *no three vertices are collinear.*

5. The key lemma

In this section we will prove a slightly stronger version of the key lemma:

Lemma 5.1. *Let G be a unit disk graph with $\alpha(G) \leq 2$. Then $V(G)$ is the union of three cliques, two of which contain a common vertex.*

Lemma 5.1 implies Lemma 3.1: Let A, B, C be three cliques whose union is $V(G)$, and that are disjoint except for a vertex v that is contained in A and B but not in C. We assume furthermore that $|A| \geq |B|$. Then $A, B \setminus \{v\}, C$ is a clique partition of $V(G)$ in which not all of the cliques have the same size.

Let us start by considering a special case. We say that a unit disk graph G is *hollow* if the interior of $\mathrm{conv}(G)$ does not contain any vertex of G. Thus, assuming the graph has at least three vertices it follows from (4.3) that all vertices of G appear on the boundary of the polygon $\mathrm{conv}(G)$. We fix one of the circular orders, say the clockwise order, in which the vertices appear on the boundary. We will use the usual interval notation for the vertices of a hollow unit disk graph. So, for two distinct vertices u, v we denote by $[u, v]$ the set of all vertices that appear in clockwise order on the boundary starting with u up to v. If $u = v$, we set $[u, v] = \{u\}$. We further define $(u, v] := [u, v] \setminus \{u\}$, $[u, v) := [u, v] \setminus \{v\}$ and $(u, v) := [u, v] \setminus \{u, v\}$. We say that u and v are *consecutive* if $u \neq v$ and either $(u, v) = \emptyset$ or $(v, u) = \emptyset$.

Hollow unit disk graphs have an advantage over general unit disk graphs. To decide whether two edges uv and xy cross reduces to determining the order of the endvertices on the boundary: The edges cross if and only if the endvertices are interleaved, that is, if and only if both (u, v) and (v, u) meet $\{x, y\}$.

Lemma 5.2. *Let G be a hollow unit disk graph. Assume G to have three distinct non-edges $\{x_1, y_1\}$, $\{x_2, y_2\}$ and $\{x_3, y_3\}$ so that $x_1, y_1, x_2, y_2, x_3, y_3$ appear in this order on the boundary of* $\operatorname{conv}(G)$. *Then $\alpha(G) \geq 3$.*

Proof. If x_1 and x_2 fail to be adjacent, set $s = x_2$. If x_1 and x_2 are adjacent then $y_1y_2 \notin E(G)$; otherwise x_1x_2 and y_1y_2 would be crossing but G does not contain any triangle on these four vertices, which is impossible by Lemma 4.1. Setting $s = y_1$, we obtain in both cases that $s \in (x_1, y_2)$ and that x_1s and sy_2 are non-edges. By symmetry, we find a $t \in (s, y_3)$ so that t is not a neighbour of s nor of y_3. We consider the four vertices x_1, s, t, y_3. Unless $\alpha(G) \geq 3$, we have that x_1 is adjacent to t and y_3 adjacent to s. Then x_1t and y_3s are two crossing edges whose endvertices do not induce any triangle, which contradicts Lemma 4.1. $\square$

We now prove the key lemma for hollow unit disk graphs.

Lemma 5.3. *Let G be a hollow unit disk graph with $\alpha(G) \leq 2$. Then for every vertex v, there are vertices v^+ and v^- so that $[v^-, v]$, $[v, v^+]$ and $V(G) \setminus [v^-, v^+]$ are cliques.*

Proof. First, we may clearly exclude the case when G is complete. Let y^+ be the last vertex in clockwise direction from v so that $[v, y^+)$ forms a clique, that is, we choose y^+ so that $[v, y^+)$ is a clique and maximal among all such cliques. By choice of y^+, there exists a $x^+ \in [v, y^+)$ that is non-adjacent to y^+. Similarly, we denote by y^- the last vertex in counterclockwise direction so that $(y^-, v]$ is a clique, and we let x^- be a non-neighbour of y^- in $(y^-, v]$. Define v^+ to be the clockwise predecessor of y^+, that is, we choose v^+ to be the vertex for which $[v, v^+] = [v, y^+)$.

If $y^- \in [v, v^+]$ then set $v^- = y^+$. This ensures that $[v^-, v]$ is a subset of the clique $(y^-, v]$. As $V(G) \setminus [v^-, v^+]$ is empty in this case, we are done. So, we will assume that $y^- \notin [v, v^+] = [v, y^+)$. Let v^- be the vertex for which $[v^-, v] = (y^-, v]$. Suppose that $V(G) \setminus [v^-, v^+] = [y^+, y^-]$ is not a clique. Thus, there exists non-adjacent $r, s \in [y^+, y^-]$, where $s \notin [y^+, r]$. Then x^+, y^+, r, s, y^-, x^- yield three pairs of non-adjacent vertices that are as in Lemma 5.2, which is impossible. $\square$

Next, we will see how the general case can be deduced from the hollow case. For this, we first note two simple consequences of Lemma 4.2:

Lemma 5.4. *Let K be a clique in a unit disk graph. Then*
(i) *$K+u$ is a clique for every vertex $u \in \mathrm{conv}(K)$;*
(ii) *every vertex v that is complete to K is adjacent to every vertex in $\mathrm{conv}(K)$.*

Proof. (i) The assertion is trivially true if K comprises at most two vertices, so we assume it to have more. Consider any vertex $k \in K$. By (4.3), there is a triangle krs with $r, s \in K$ that contains u in its convex hull. Now the assertion follows from Lemma 4.2.

(ii) This follows directly from (i) with K replaced by $K' = K + u$. $\square$

We quickly exclude one more easy case in the proof of Lemma 5.1.

Lemma 5.5. *Let G be a unit disk graph with $\alpha(G) \leq 2$, and let u and v be two non-adjacent vertices. Assume that all vertices of G lie on one side of the line through u and v, that is, the line through u and v does not separate any two points of $\mathrm{conv}(G)$. Then $V(G)$ is the union of three cliques, two of which contain a common vertex.*

Proof. As $\alpha(G) \leq 2$, every vertex of G except for u, v is a neighbour of u or a neighbour of v. Moreover, both $N(u) \setminus N(v)$ and $N(v) \setminus N(u)$ are cliques. So, if $N(u) \cap N(v)$ is a clique as well then $V(G)$ is the union of the three cliques $(N(u) \setminus N(v)) + u$, $(N(u) \cap N(v)) + u$ and $N(v) \cap N(u)$.

In order to show that $N(u) \cap N(v)$ is a clique, consider two common neighbours x, y of u and v. If $x \in \mathrm{conv}(u, y, v)$ or if $y \in \mathrm{conv}(u, x, v)$ then x and y are adjacent by Lemma 4.2. So, suppose that neither is the case. In a similar way, it follows from $uv \notin E(G)$ that neither u nor v can be contained in the interior of the convex hull of the other three. Thus, all four vertices lie on the boundary of $\mathrm{conv}(u, v, x, y)$. Because $\mathrm{conv}(G)$ lies on one side of the line through u and v, we deduce that one of the two pairs of edges, uy, vx and ux, vy cross. That x and y are adjacent now follows from Lemma 4.1. $\square$

In our proof of the key lemma there is only one obstacle left, which we will overcome with the help of the next lemma. The lemma is based on the insight that, provided the graph is contained in a unit disk, the vertices on the boundary of $\mathrm{conv}(G)$ largely determine the behaviour of the interior vertices. Slightly more precisely, we know from Lemma 5.3 that the outer vertices can be partitioned into three cliques, and we will see that each interior vertex can easily be assigned to one of these cliques—with the exception of two small zones of vertices. Handling these two zones will be the main difficulty.

Lemma 5.6. *Let G be a unit disk graph with $\alpha(G) \leq 2$. If G is contained in a unit disk then $V(G)$ is the union of three cliques, two of which have a common vertex.*

Proof. We will show that there is a vertex b on the boundary of $\mathrm{conv}(G)$, and three cliques B^+, B^-, R that cover $V(G)$ so that $b \in B^+ \cap B^-$.

By assumption, all of G is contained in a unit disk, which means that there is a point $p \in \mathbb{R}^2$ of distance at most 1 to every vertex in G. Now, adding p as a vertex to G is entirely harmless: $G + p$ still has stability at most two and, assuming we find cliques as desired in $G + p$ with $b \neq p$, we obtain such cliques for G by simply deleting p from the cliques. Thus, we may assume that

(5.1) *G has a vertex p that is adjacent to every other vertex.*

As the lemma trivially holds if G has at most two vertices, we assume from now on that there are at least three vertices in G. Denote by F the set of vertices of G on the boundary of $\mathrm{conv}(G)$. Then the graph induced by F is a hollow unit disk graph, and we will continue to use the interval notation for this induced subgraph of G, and only for this graph. This means that a set $[u, v]$ is always understood with respect to the hollow unit disk graph on F, and that in particular, $[u, v] \subseteq F$.

In light of Lemma 5.5 we may assume that

(5.2) *two consecutive vertices on the boundary of* $\mathrm{conv}(G)$ *are adjacent.*

Pick any vertex $b \in F$ other than p. By Lemma 5.3 we may choose $b^+, b^- \in F$ so that $[b, b^+]$, $[b^-, b]$ and $F \setminus [b^-, b^+]$ are cliques that meet at most in b, and such that $F \setminus [b^-, b^+]$ is minimal subject to this condition. Observe that $b^+ = b^-$ is impossible: This would entail $b = b^+ = b^-$ but both of $[b, b^+]$ and $[b^-, b]$ contain two vertices by (5.2). Thus, $F \setminus [b^-, b^+] = (b^+, b^-)$. If $(b^+, b^-) \neq \emptyset$ let r^+, r^- be so that $[r^+, r^-] = (b^+, b^-)$. If, on the other hand, $(b^+, b^-) = \emptyset$ then we put $r^+ = b^-$ and $r^- = b^+$.

We now have that

(5.3) *every pair of consecutive vertices in F lies in one of the following cliques;* $[b, b^+]$, $[b^+, r^+]$, (b^+, b^-), $[r^-, b^-]$ *and* $[b^-, b]$.

If $(b^+, b^-) \neq \emptyset$ then trivially every pair of consecutive vertices lies in one of the five sets. In the case of $(b^+, b^-) = \emptyset$ we have $[r^-, b^-] = [b^+, b^-]$.

It remains only to verify that $[b^+, r^+]$ and $[r^-, b^-]$ are cliques as claimed. Indeed, b^+ and r^+, as well as b^- and r^-, are consecutive vertices on the boundary. Thus, by (5.2), $[b^+, r^+]$ is simply the edge b^+r^+, and $[r^-, b^-]$ coincides with the edge b^-r^-. This proves (5.3).

We define

$$B^+ = \mathrm{conv}([b, b^+] + p) \cap V(G), \qquad B^- = \mathrm{conv}([b^-, b] + p) \cap V(G),$$
$$T^+ = \mathrm{conv}(b^+, r^+, p) \cap V(G), \qquad T^- = \mathrm{conv}(r^-, b^-, p) \cap V(G),$$
$$\text{and } R = \mathrm{conv}([r^+, r^-] + p) \cap V(G).$$

See Figure 1 for an illustration. Observe that every vertex is contained in one of the five sets. Moreover, as p is adjacent to every other vertex, we deduce from (5.3) and Lemma 5.4 that B^+, B^-, T^+, T^- and R are cliques.

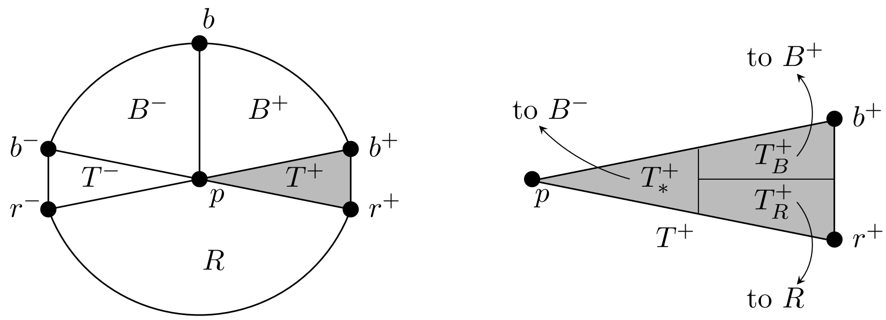

FIGURE 1. The five cliques (left); how to divide up T^+ (right)

Assume for the moment that $(b^+, b^-) = \emptyset$. Then $[b^+, r^+] = [b^+, b^-] = [r^-, b^-]$, which means that $T^+ = T^-$. Thus, we see that $V(G)$ is the union of the three cliques B^+, B^- and T^+, two of which contain b. As we are done in this case, we will assume from now on that

(5.4) $(b^+, b^-) = [r^+, r^-] \neq \emptyset$ *and neither of* $[b, r^+]$ *and* $[r^-, b]$ *is a clique.*

The second part of the assertion follows from the first: If, for instance, $[b, r^+]$ was a clique then $F \setminus [b^-, r^+]$ would be strictly smaller than $F \setminus [b^-, b^+]$, and $[b, r^+]$ and $[b^-, b]$ would still be cliques meeting only in b; this however contradicts the choice of b^+ and b^-.

The rest of the proof will be spent on dividing up $T^+ \cup T^-$ among the other cliques, so that we obtain three cliques that cover all of G. More precisely, we will partition T^+ and T^- into sets T^+_B, T^+_R, T^+_* and T^-_B, T^-_R, T^-_*, respectively, so that $B^+ \cup T^+_B \cup T^-_*$, $B^- \cup T^-_B \cup T^+_*$ and $R \cup T^+_R \cup T^-_R$ are cliques; see Figure 1. As b is contained in the first two of these three, this will complete the proof of the lemma. In order to do so, we define

- $T^+_B := \{t \in T^+ : t$ *is complete to* $B^+\}$;
- $T^+_R := \{t \in T^+ \setminus T^+_B : t$ *is complete to* $R\}$;
- $T^+_* := T^+ \setminus (T^+_B \cup T^+_R)$.

The sets T^-_B, T^-_R and T^-_* are defined symmetrically.

We claim that

(5.5) $R \cup T^+_R \cup T^-_R$ *is a clique.*

Suppose that is not the case. We note that R, T^+ and T^- are cliques and that T^+_R and T^-_R are defined in such a way that each is complete to R. Thus, $R \cup T^+_R \cup T^-_R$ may only fail to be a clique if there exists a pair t^+, t^- of non-adjacent vertices with $t^+ \in T^+_R$ and $t^- \in T^-_R$. By definition of T^+_R, the vertex t^+ is not complete to B^+; otherwise t^+ would be in T^+_B. Since $B^+ \subseteq \mathrm{conv}([b, b^+] + p)$, Lemma 5.4 implies that t^+ has a non-neighbour s^+ in $[b, b^+]$. Symmetrically, we find an $s^- \in [b^-, b]$ that is non-adjacent to t^-.

We will focus on t^+, t^- and the following vertices that appear in this clockwise order on the boundary of $\mathrm{conv}(G)$: b^-, s^-, s^+, b^+. Observe that t^+ is adjacent to s^- and t^- is adjacent to s^+; otherwise we would obtain a contradiction to $\alpha(G) \leq 2$. Moreover, t^+ and b^+ are adjacent since both are elements of the clique T^+. In the same way, we have $t^-b^- \in E(G)$.

Suppose that $t^+b^- \in E(G)$. Then t^+ is adjacent to p, to b^- and to $r^- \in R$ (the last adjacency is because of the definition of T^+_R that ensures that t^+ is adjacent to every vertex in R). With Lemma 5.4 we conclude now that t^+ is complete to T^-, which is impossible as $t^- \in T^-$; thus, t^+ and b^- are non-neighbours, similarly, t^- and b^+ are non-neighbours. Let us sum up the non-adjacencies:

(5.6)
None of the pairs t^-t^+, t^+s^+, t^-s^-, t^+b^- and t^-b^+ is an edge of G.

See also Figure 2.

Next, let us note that s^-, s^+, b^+, b^- are pairwise distinct. Indeed, the fact that t^+ is adjacent to b^+ and s^- but not to s^+ nor to b^- implies that $s^+ \neq s^-$, $s^+ \neq b^+$, $s^- \neq b^-$ and $b^+ \neq b^-$. All other identities are excluded by the fact that s^-, s^+, b^+, b^- appear in this order on F.

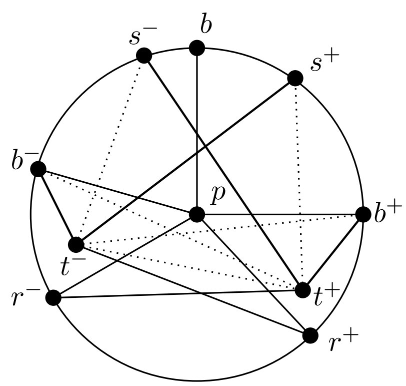

FIGURE 2. The crossing paths of (5.5); non-adjacencies in broken lines. (Not all adjacencies shown.)

To conclude, we find two disjoint paths $s^+t^-b^-$ and $s^-t^+b^+$ with endvertices in F that are interleaved: s^+, b^+, b^-, s^- appear in this clockwise order. Thus an edge of the first path needs to cross an edge of the second path. The possible pairs are s^-t^+ and s^+t^-; b^+t^+ and b^-t^-; s^+t^- and b^+t^+; s^-t^+ and b^-t^-. Suppose that s^-t^+ and s^+t^- cross. Then, by Lemma (4.1), three of the four vertices s^+, s^-, t^+, t^- form a triangle, which, however, is impossible by (5.6). In a similar way, we see that none of the other pairs of edges may cross, and we have therefore found the final contradiction that proves (5.5).

We show next that

(5.7) *T^+_* is complete to $B^- \cup T^-$, and T^-_* is complete to $B^+ \cup T^+$.*

By symmetry it suffices to show that any $t \in T_*^+$ is complete to $B^- \cup T^-$. We first observe that

(5.8) *there are distinct non-neighbours x, y of t so that $[x, y]$ is a clique and t is complete to (y, x).*

To see (5.8), denote by X the set of non-neighbours of t in F. Note that t has by definition of T_*^+ a non-neighbour in $B^+ \subseteq \text{conv}([b, b^+] + p)$ and another one in $R \subseteq \text{conv}([r^+, r^-] + p)$. Lemma 5.4 implies that t has therefore a non-neighbour in $[b, b^+]$ and in $[r^+, r^-]$, and consequently that $|X| \geq 2$. Since $\alpha(G) \leq 2$, the set X has to be a clique. We deduce from Lemma 5.4 that $t \notin \text{conv}(X)$. Thus, there is a line, disjoint from $X + t$, that separates t from X. Let H_t be the corresponding open half-plane containing t and let H_X be the other, which then contains all of X. Pick a point q in $H_t \cap F$, and following the boundary of $\text{conv}(G)$ in clockwise direction let x be the first vertex in X and let y be the last. Then clearly $X \subseteq [x, y]$. Moreover, $t \notin \text{conv}([x, y])$ as $[x, y] \subseteq H_X$.

Now consider the unit disk graph on $[x, y] + t$, which is hollow: No $z \in [x, y] \subseteq F$ lies in the interior of $\text{conv}(G)$, so in particular, no such z lies in the interior of $\text{conv}([x, y])$. Moreover, $t \notin \text{conv}([x, y])$ by choice of x and y. Since y is non-adjacent to t, which in turn is non-adjacent to x, we may apply Lemma 5.2 to the hollow unit disk graph on $[x, y] + t$ in order to deduce that $[x, y]$ is a clique. This finishes (5.8).

We distinguish four cases.

Case 1: $x \in [r^+, b]$ and $y \in [r^+, b)$.

Since $(b, b^+]$ is disjoint from $\{x, y\}$, we have that $(b, b^+] \subseteq (x, y)$ or $(b, b^+] \subseteq (y, x)$. Then $[b, b^+] \subseteq [x, y)$ or $[b, b^+] \subseteq (y, x)$ as $y \neq b$. By definition of T_*^+, t has a non-neighbour in $B^+ \subseteq \text{conv}([b, b^+] + p)$, and thus also in $[b, b^+]$ by Lemma 5.4. As t is complete to (y, x) by (5.8), it follows that $[b, b^+] \subseteq [x, y)$. By (5.8), $[x, y]$ is a clique. Its subset $[b, r^+]$ is therefore a clique as well, which is impossible by (5.4).

Case 2: $\{x, y\} \subseteq [b^-, b^+]$.

The vertex t has by definition of T_*^+ a non-neighbour in $R \subseteq \text{conv}([r^+, r^-] + p)$, which by way of Lemma 5.4 implies that t has a non-neighbour in $[r^+, r^-]$ as well. On the other hand, t is complete to (y, x) by (5.8), which excludes that $[r^+, r^-] \subseteq (y, x)$. Now, it follows from $\{x, y\} \subseteq [b^-, b^+]$ that $[r^+, r^-] \subseteq (x, y)$. Thus, $[b^+, b^-] \subseteq [x, y]$ is a clique by (5.8), which implies with Lemma 5.4 that $T^+ \cup R \cup T^- \subseteq \text{conv}([b^+, b^-] + p)$ is a clique. This, however, is impossible as $t \in T_*^+$ is supposed to have a non-neighbour in R.

Case 3: $x \in [r^+, r^-]$ and $y \in [b, b^+]$.

In this case, we find that $[b^-, b] \subseteq (x, y]$, and thus that $[r^-, b] \subseteq [x, y]$ is a clique, which contradicts (5.4).

Case 4: $x \in (b, b^+]$ and $y \in [r^+, r^-]$.

As $[b^-, b] \subseteq (y, x)$, it follows from (5.8) that t is complete to $[b^-, b]$ and thus to $B^- \subseteq \text{conv}([b^-, b] + p)$ by Lemma 5.4. Next, let us show that t is complete to T^- as well. If $r^- = y$ then all of $[r^+, r^-]$ is contained in the clique $[x, y]$ as well as b^+. From Lemma 5.4 we deduce that $R \cup T^+ \subseteq \mathrm{conv}([b^+, r^-] + p)$ is a clique, which is impossible as $t \in T^+_* \subseteq T^+$ cannot, by definition, be complete to R. Thus, $r^- \in (y, x)$, which means that t is adjacent to r^-. As t is also adjacent to $b^- \in B^-$ it follows from Lemma 5.4 that t is complete to $T^- \subseteq \mathrm{conv}(b^-, r^-, p)$, as desired.

The Cases 1–4 cover all possible values for x and y. Indeed, assume first that neither x nor y is equal to b. Then any pair of x, y not treated in Case 1 either lies completely in $(b, b^+]$, or one of x, y lies in $(b, b^+]$ and the other in $[r^+, b)$. The former configuration is covered by Case 2, which also takes care of a part of the latter configuration. It remains to consider the case when one of x, y lies in $(b, b^+]$ and the other in $[r^+, r^-]$. This is dealt with in Cases 3 and 4. So, assume now that $x = b$. Then $y \in [r^+, b)$ falls under Case 1, and $y \in (b, r^+) = (b, b^+]$ under Case 2. Finally, if $y = b$ then $x \in [b^-, b^+]$ is covered by Case 2, while $x \in [r^+, r^-]$ is covered by Case 3.

We therefore have proved (5.7). Now the definitions of T^+_B, T^+_R, T^+_* and T^-_B, T^-_R, T^-_* together with (5.5) and (5.7) imply directly that $B^+ \cup T^+_B \cup T^-_*$, $B^- \cup T^-_B \cup T^+_*$ and $R \cup T^+_R \cup T^-_R$ are cliques. □

We can finally prove our key lemma:

Proof of Lemma 5.1. We need to find three cliques whose union is $V(G)$ so that two of them share a vertex. Assume first that there are two vertices u, v of distance $\geq \sqrt{3}$. Because of $\alpha(G) \leq 2$, we see that $(N(u) \setminus N(v)) + u$ as well as $N(v) \setminus N(u)$ induce cliques. These two together with $(N(u) \cap N(v)) + u$, which is a clique by Lemma 4.4, form three cliques as stated. If all pairs of vertices have distance at most $\sqrt{3}$ then the assertion follows directly from Lemmas 4.3 and 5.6. □

Acknowledgment

I am very grateful for inspiring discussions with Ross Kang, as well as with Naho Fujimoto.

References

1. B. Balasundaram and S. Butenko, *Optimization problems in unit-disk graphs*, Encyclopedia of Optimization (C. Floudas and M. Pardalos, eds.), Springer, 2008, pp. 2832–2844.
2. H. Breu, *Algorithmic aspects of constrained unit disk graphs*, Ph.D. thesis, The University of British Colombia, 1996.
3. M. Chudnovsky and P. Seymour, *The structure of claw-free graphs*, Surveys in combinatorics 2005, vol. 327, London Math Soc Lecture Note, 2005, pp. 153–171.
4. B.N. Clark, C.J. Colbourn, and D.S. Johnson, *Unit disk graphs*, Disc. Math. **86** (1990), 165–177.
5. R. Diestel, *Graph theory* (4th edition), Springer-Verlag, 2010.
6. S. Gerke and C. McDiarmid, *Graph imperfection*, J. Combin. Theory (Series B) **83** (2001), 58–78.

7. A. Gräf, M. Stumpf, and G. Weißenfels, *On coloring unit disk graphs*, Algorithmica **20** (1998), 277–293.
8. H. Grötzsch, *Zur Theorie der diskreten Gebilde. VII. Ein Dreifarbensatz für dreikreisfreie Netze auf der Kugel*, Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg. Math.-Nat. Reihe **8** (1959), 109–120.
9. W.K. Hale, *Frequency assignment: Theory and applications*, Proc. IEEE **68** (1980), 1497–1514.
10. L. Lovász and M.D. Plummer, *Matching theory*, Akadémiai Kiadó - North Holland, 1986.
11. E. Malesińska, S. Piskorz, and G. Weißenfels, *On the chromatic number of disk graphs*, Networks **32** (1998), 13–22.
12. M.V. Marathe, H. Breu, H.B. Hunt, S.S. Ravi, and D.J. Rosenkrantz, *Simple heuristics for unit disk graphs*, Networks **25** (1995), 59–68.
13. C. McDiarmid, *Random channel assignment in the plane*, Random Structures & Algorithms **22** (2003), 187–212.
14. C. McDiarmid and B. Reed, *Channel assignment and weighted coloring*, Networks **36** (2000), 114–117.
15. R. Peeters, *On coloring j-unit spheres*, Tech. Report FEW 512, Tilburg University, 1991.
16. V. Raghavan and J. Spinrad, *Robust algorithms for restricted domains*, J. Algorithms **48** (2003), 160–172.

Universität Ulm, Germany
E-mail address: `henning.bruhn@uni-ulm.de`