

The Clique Problem in Ray Intersection Graphs

Sergio Cabello · Jean Cardinal · Stefan Langerman

Received: 20 December 2012 / Revised: 28 June 2013 / Accepted: 15 August 2013 /
Published online: 4 September 2013
© Springer Science+Business Media New York 2013

Abstract Ray intersection graphs are intersection graphs of rays, or halflines, in the plane. We show that any planar graph has an even subdivision whose complement is a ray intersection graph. The construction can be done in polynomial time and implies that finding a maximum clique in a segment intersection graph is NP-hard. This solves a 21-year old open problem posed by Kratochvíl and Nešetřil (Comment Math Univ Carolinae 31(1):85–93, 1990).

Keywords Geometric intersection graphs · Segment intersection graphs · Maximum clique · NP-hardness

1 Introduction

The intersection graph of a collection of sets has one vertex for each set, and an edge between two vertices whenever the corresponding sets intersect. Of particular interest are families of intersection graphs corresponding to geometric sets in the plane. In this contribution, we will focus on *segment intersection graphs*, intersection graphs of line segments in the plane.

S. Cabello
Department of Mathematics, IMFM, FMF, University of Ljubljana, Ljubljana, Slovenia
e-mail: sergio.cabello@fmf.uni-lj.si

J. Cardinal (✉) · S. Langerman
Computer Science Department, Université Libre de Bruxelles (ULB), Brussels, Belgium
e-mail: jcardin@ulb.ac.be

S. Langerman
e-mail: slanger@ulb.ac.be

In a seminal paper, Kratochvíl and Nešetřil [12] proposed studying the complexity of two classical combinatorial optimization problems, the maximum independent set and the maximum clique, in geometric intersection graphs. While those problems are known to be hard to approximate in general graphs (see for instance [5, 9]), their restriction to geometric intersection graphs may be more tractable. They proved that the maximum independent set problem remains NP-hard for segment intersection graphs, even if those segments have only two distinct directions. It was also shown that in that case, the maximum clique problem can be solved in polynomial time. The complexity of the maximum clique problem in general segment intersection graphs was left as an open problem, and remained so until now. In their survey paper “On six problems posed by Jarik Nešetřil” [3], Bang-Jensen *et al.* describe this problem as being “among the most tantalizing unsolved problem in the area”.

Some progress has been made in the meanwhile. In 1992, Middendorf and Pfeiffer [13] showed, with a simple proof, that the maximum clique problem is NP-hard for intersection graphs of 1-intersecting curve segments that are either line segments or curves made of two orthogonal line segments. They also gave a polynomial time dynamic programming algorithm for the special case of line segments with endpoints of the form $(x, 0), (y, i)$, with $i \in \{1, \dots, k\}$ for some fixed k . Another step was made by Ambühl and Wagner [2] in 2005, who showed that the maximum clique problem is NP-hard for intersection graphs of ellipses of fixed, arbitrary, aspect ratio. Unfortunately, this ratio must be bounded, which excludes the case of segments.

Our results. We prove that the maximum clique problem in segment intersection graphs is NP-hard, even if a representation as a set of segments is given as input. In fact, we prove the stronger result that the problem is NP-hard even in *ray intersection graphs*, defined as intersection graphs of rays, or halflines, in the plane. This complexity result is a consequence of the following structural lemma: every planar graph has an even subdivision whose complement is a ray intersection graph. Furthermore, the corresponding set of rays has a natural polynomial size representation. Hence solving the maximum clique problem in this graph allows us to recover the maximum independent set in the original planar graph, a task well known to be NP-hard [8]. The construction is detailed in Sect. 2.

Related work. We prove that the complement of some subdivision of any planar graph can be represented as a segment intersection graph. Whether the complement of every planar graph is a segment intersection graph remains an open question. In 1998, Kratochvíl and Kuběna [10] showed that the complement of any planar graph is the intersection graph of a set of convex polygons. More recently, Francis *et al.* [7] proved that the complement of any partial 2-tree is a segment intersection graph. Partial 2-trees are planar, and in particular every outerplanar graph is a partial 2-tree. The representability of planar graphs by segment intersection graphs, formerly known as Scheinerman’s conjecture, was proved recently by Chalopin and Gonçalves [4].

The maximum independent set problem in intersection graphs has been studied by Agarwal and Mustafa [1]. In particular, they proved that it can be approximated within a factor $n^{1/2+o(1)}$ in polynomial time for segment intersection graphs. This has been recently improved by Fox and Pach [6], who described, for any $\varepsilon > 0$, an n^ε -approximation algorithm. In fact, their technique also applies to the maximum

clique problem, and therefore n^ε is the best known approximation factor for this problem too.

In 1994, Kratochvíl and Matoušek [11] proved that the recognition problem for segment intersection graphs is NP-hard, and in PSPACE.

Notation. For any natural number m we use $[m] = \{1, \dots, m\}$. In a graph G , a *rotation system* is a list $\pi = (\pi_v)_{v \in V(G)}$, where each π_v fixes the clockwise order of the edges of $E(G)$ incident to v . When G is an embedded planar graph, the embedding uniquely defines a rotation system, which is often called a combinatorial embedding. For the rest of the paper we use *ray* to refer to an *open ray*, that is, a ray does not contain its origin. Therefore, whenever two rays intersect they do so in the relative interior of both. Since our construction does not use degeneracies, we are not imposing any restriction by considering only open rays. A subdivision of a graph G is said to be *even* if each edge of G is subdivided an even number of times.

2 Construction

Let us start with an overview of the approach. We first construct a set of curves that will form the *reference frame*. This construction is quite generic and depends only on a parameter $k \in \mathbb{N}$. We then show that the complement of any tree has a special type of representation, called *snooker representation*, which is constructed iteratively over the levels of the tree. The number of levels of the tree is closely related to k , the parameter used for the reference frame. We then argue that if G is a planar graph that consists of a tree T and a few, special paths of length two and three, then the complement of G can be represented as an intersection graph of rays by extending a snooker representation of T . Finally, we argue that any planar graph has an even subdivision that can be decomposed into a tree and a set of paths of length two and three with the required properties.

We first describe the construction using real coordinates. The construction does not rely on degeneracies, and thus we can slightly perturb the coordinates used in the description. This perturbation is enough to argue that a representation can be computed in polynomial time. Then, using a relation between the independence number of a graph G and an even subdivision of G , we obtain that computing a maximum clique in a ray intersection graph is NP-hard, even if a representation is given as part of the input.

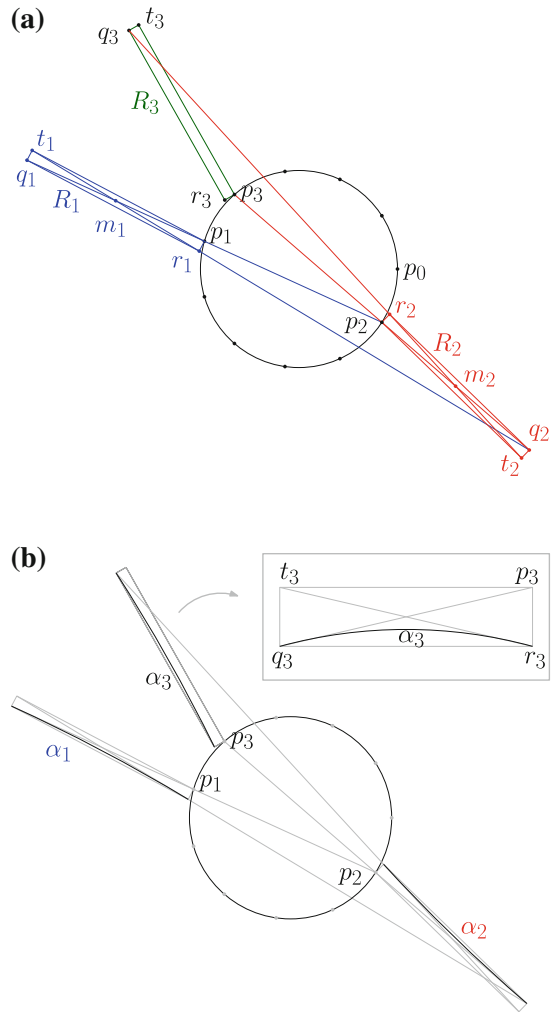
2.1 Reference Frame

Let $k \geq 3$ be an odd number to be chosen later. We set $\theta = \frac{k-1}{k}\pi$, and define for $i = 0, \dots, k-1$ the points

$$p_i = (\cos(i \cdot \theta), \sin(i \cdot \theta)).$$

The points p_i lie on a unit circle centered at the origin; see Fig. 1a. In what follows, all indices are taken modulo k , hence in particular $p_k = p_0$. For each $i \in [k-1]$ we let q_i be the point $p_i + (p_i - p_{i+1})$, symmetric to p_{i+1} with respect to p_i , let m_i be the midpoint between p_i and q_i , and let t_i be, among the two points along the line through q_{i+1} and m_i with the property $|m_i t_i| = |m_i p_i|$, the one that is furthest from

Fig. 1 The reference frame. **a** The points p_0, \dots, p_{k-1} and the rectangles R_i , for $k = 11$, **b** The circular arcs α_i



q_{i+1} . We define R_i to be the rectangle with vertices p_i, t_i , and q_i . The fourth vertex of R_i is implicit and is denoted by r_i . Any two rectangles R_i and R_j are congruent with respect to a rotation around the origin. We have constructed the rectangles R_i in such way that, for any $i \in [k - 1]$, the line supporting the diagonal $p_i q_i$ of R_i contains p_{i+1} and the line supporting the diagonal $r_i t_i$ of R_i contains q_{i+1} .

For each $i \in [k - 1]$, let α_i be the circular arc tangent to both diagonals of R_i with endpoints q_i and r_i , and contained in R_i (see Fig. 1b). Note that the curves α_i and the rectangles R_i have been chosen so that any line that intersects α_i twice or is tangent to the curve α_i must intersect the curve α_{i+1} . For any $i \in [k - 2]$, let Γ_i be the set of rays that intersect α_i twice or are tangent to α_i and have their origins on α_{i+1} . We also define Γ_0 as the set of rays with origin on α_1 and passing through p_0 . The rays of Γ_i that are tangent to α_i will play a special role. In fact, we will only use rays of Γ_i that are “near-tangent” to α_i .




Fig. 2 Illustration for the proof of Lemma 2.1

Lemma 2.1 When $|j - i| > 1$, any ray from Γ_i intersects any ray from Γ_j .

Proof Because of symmetry, it is enough to argue that all rays in Γ_1 intersect all rays of Γ_j for $j \geq 3$. We provide a synthetic proof of this fact, illustrated on Fig. 2. For this, it is enough to show that the even points p_4, p_6, \dots are below the line through t_1 and q_2 . By construction, this can be reduced to proving that p_4 is below the line through t_1 and q_2 . Let x be the point on the segment $q_1 p_1$ at distance $|p_1 p_4|$ from p_1 . Since $|p_1 p_4| > \frac{1}{2}|p_1 p_2| = |p_1 m_1|$, the point x lies on the segment $q_1 m_1$. Because $\angle(p_2, p_1, p_4) = \frac{\pi}{k}$ we have $\angle(q_1, p_1, p_4) = \pi - \frac{\pi}{k}$. Since the triangle $\Delta(x, p_1, p_4)$ is isosceles because $|xp_1| = |p_1 p_4|$, and $\angle(q_1, p_1, p_4) = \pi - \frac{\pi}{k}$ we have $\angle(p_1, x, p_4) = \frac{\pi}{2k}$. Since $\angle(p_1, p_2, p_3) = \frac{\pi}{k}$ and p_3, p_2, q_2 are collinear, then $\angle(p_1, p_2, q_2) = \pi - \frac{\pi}{k}$ and, using that $\Delta(p_1, p_2, q_2)$ is isosceles, we have $\angle(p_2, p_1, q_2) = \frac{\pi}{2k}$. From $\angle(p_2, p_1, q_2) = \frac{\pi}{2k} = \angle(p_1, x, p_4)$ and the collinearity of x, p_1, p_2 , we conclude that $p_1 q_2$ and $x p_4$ are parallel. Since m_1 lies between x and p_1 , then the line through m_1 and q_2 passes above p_4 . \square

Lemma 2.2 Any ray of Γ_{i+1} tangent to α_{i+1} at the point $x \in \alpha_{i+1}$ intersects any ray from Γ_i , except those having their origin at x .

Proof Any ray of Γ_{i+1} tangent to α_{i+1} intersects the boundary of R_{i+1} in the segments $r_{i+1} p_{i+1}$ and $q_{i+1} t_{i+1}$, while any ray of Γ_i intersects the segment $t_{i+1} p_{i+1}$. \square

Note that the whole construction depends only on the parameter k . We will refer to it as *reference frame*.

2.2 Complements of Trees

Let T be a tree with a rotation system π_T , let r be a vertex in T and let $rs \in E(T)$ be an arbitrary edge incident to r . The triple (π_T, r, rs) induces a natural linear order $\tau = \tau(\pi_T, r, rs)$ on the vertices of T . This order τ corresponds to the order followed by a breadth-first traversal of T from r with the following additional restrictions:




Fig. 3 In this example, assuming that π is as drawn, the linear order τ is $r, v_1^1, v_1^2, \dots, v_1^5, v_2^1, \dots$

- (i) s is the second vertex;
- (ii) the children of any vertex v are visited according to the clockwise order π_v ;
- (iii) if $v \neq r$ has parent v' , the first child of u of v is such that vu is the successor of vv' in the clockwise order π_v .

We say that vertices v and v' at the same level are *consecutive* when they are consecutive in τ . See Fig. 3. The linear order will be fixed through our discussion, so we will generally drop it from the notation. Henceforth, whenever we talk about a tree T and a linear order τ on $V(T)$, we assume that τ is the natural linear order induced by a triple (π_T, r, rs) . In fact, the triple (π_T, r, rs) is implicit in τ . For any vertex v we use v^+ for its successor and v^- for its predecessor.

A *snooker representation* of the complement of an embedded tree T with linear order τ is a representation of T with rays that satisfies the following properties:

- (a) Each vertex v at level i in T is represented by a ray γ_v from Γ_i . Thus, the origin of γ_v , denoted by a_v , is on α_{i+1} . Note that this implies that k is larger than the depth of T .
- (b) If a vertex u has parent v , then γ_u passes through the origin a_v of γ_v . (Here it is relevant that we consider all rays to be open, as otherwise γ_u and γ_v would intersect.) In particular, all rays corresponding to the children of v pass through the point a_v .
- (c) The origins of rays corresponding to consecutive vertices u and v of level i are consecutive along α_{i+1} . That is, no other ray in the representation has its origin on α_{i+1} and between the origins γ_u and γ_v .

Lemma 2.3 *The complement of any embedded tree with a linear order τ has a snooker representation.*

Proof Consider a reference frame with k larger than the depth of T . The construction we provide is iterative over the levels of T . Note that, since we provide a snooker

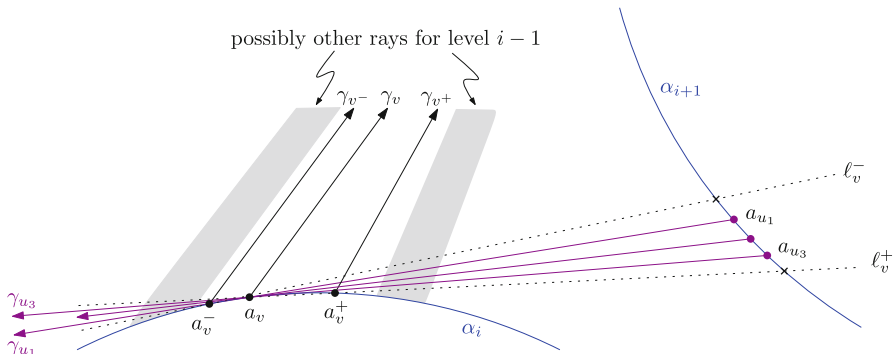


Fig. 4 An example showing the construction of a snooker representation when v has three children and v^+ and v^- are at the same level as v

representation, it is enough to tell for each vertex $v \neq r$ the origin a_v of the ray γ_v . Property (b) of the snooker representation provides another point on the ray γ_v , and thus γ_v is uniquely defined. The ray γ_r for the root r is the ray of Γ_0 with origin a_r at the center of α_1 .

Consider any level $i > 1$ and assume that we have a representation of the vertices at level $i - 1$. Consider a vertex v at level $i - 1$ and let u_1, \dots, u_d denote its d children. If the successor v^+ of v is also at level $i - 1$, we take $a_v^+ = a_{v^+}$, and else we take a_v^+ to be an endpoint of α_i such that no other origin is between the endpoint and a_v (see Fig. 4). Similarly, if the predecessor v^- of v is at level $i - 1$, we take $a_v^- = a_{v^-}$, and else we take a_v^- to be an endpoint of α_i such that no other origin is between the endpoint and a_v . (If v is the only one vertex at level i , we also make sure that $a_v^- \neq a_v^+$.) Let ℓ_v^+ be the line through a_v and a_v^+ . Similarly, let ℓ_v^- be the line through a_v and a_v^- . We then choose the points a_{u_1}, \dots, a_{u_d} on the portion of α_{i+1} contained between ℓ_v^+ and ℓ_v^- such that the $d + 2$ points $\ell_v^- \cap \alpha_{i+1}, a_{u_1}, \dots, a_{u_d}, \ell_v^+ \cap \alpha_{i+1}$ are regularly spaced. Since the ray γ_{u_j} has origin a_{u_j} and passes through a_v , this finishes the description of the procedure. Because a_{u_j} lies between ℓ_v^+ and ℓ_v^- , the ray γ_{u_j} either intersects α_i twice or is tangent to α_i , and thus $\gamma_{u_j} \in \Gamma_i$.

Recall that from Lemma 2.1, any ray from Γ_i intersects any ray from Γ_j when $|j - i| > 1$. Therefore, vertices from levels i and j , where $|i - j| > 1$, intersect. For vertices u and v at levels $i - 1$ and i , respectively, Lemma 2.2 and the choices for a_v imply that γ_v intersects γ_u if and only if u is not the parent of v in T . For vertices u and v at the same level i , the rays γ_u and γ_v intersect: if they have the same parent w , then they intersect on a_w , and if they have different parents, the order of their origins a_u and a_v on α_{i+1} and the order of their intersections with α_i are reversed. \square

2.3 A Tree with a Few Short Paths

Let T be an embedded tree with a linear order τ . An *admissible extension* of T is a graph P with the following properties:

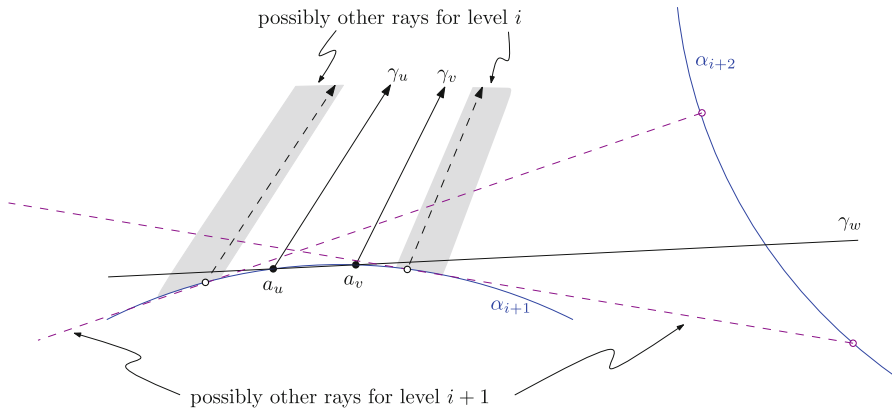


Fig. 5 Case 1 in the proof of Lemma 2.4: adding a path with one internal vertex

- P is a union of vertex-disjoint paths;
- each maximal path in P has 3 or 4 vertices;
- the endpoints of each maximal path in P are leaves of T that are consecutive and at the same level;
- the internal vertices of any path in P are not vertices of $V(T)$.

Note that $T + P$ is a planar graph because we only add paths between consecutive leaves.

Lemma 2.4 *Let T be an embedded tree and let P be an admissible extension of T . The complement of $T + P$ is a ray intersection graph.*

Proof We construct a snooker representation of T using Lemma 2.3 where k , the number of levels, is the depth of T plus 2 or 3, whichever is odd. We will use a local argument to represent each maximal path of P , and each maximal path of P is treated independently. It will be obvious from the construction that rays corresponding to vertices in different paths intersect. We distinguish the case where the maximal path has one internal vertex or two.

Consider first the case of a maximal path in P with one internal vertex. Thus, the path is uwv where u and v are consecutive leaves in T and $w \notin V(T)$ is not yet represented by a ray. The origins a_u and a_v of the rays γ_u and γ_v , respectively, are distinct and consecutive along α_{i+1} because we have a snooker representation. We thus have the situation depicted in Fig. 5. We can then just take the ray γ_w to be the line through a_u and a_v . (This line can also be a ray with an origin sufficiently far away.) This line intersects the ray of any other vertex different from γ_u and γ_v .

Consider now the case of a maximal path in P with two internal vertices. Thus, the path is $uwv'v$ where u and v are consecutive leaves in T and $w, w' \notin V(T)$. In this case, the construction depends on the relative position of the origins a_u and a_v , and we distinguish two scenarios: (i) shifting the origin a_u of ray γ_u towards a_v while maintaining the slope introduces an intersection between γ_u and the ray for the parent of u or (ii) shifting the origin a_v of ray γ_v towards a_u while maintaining the slope

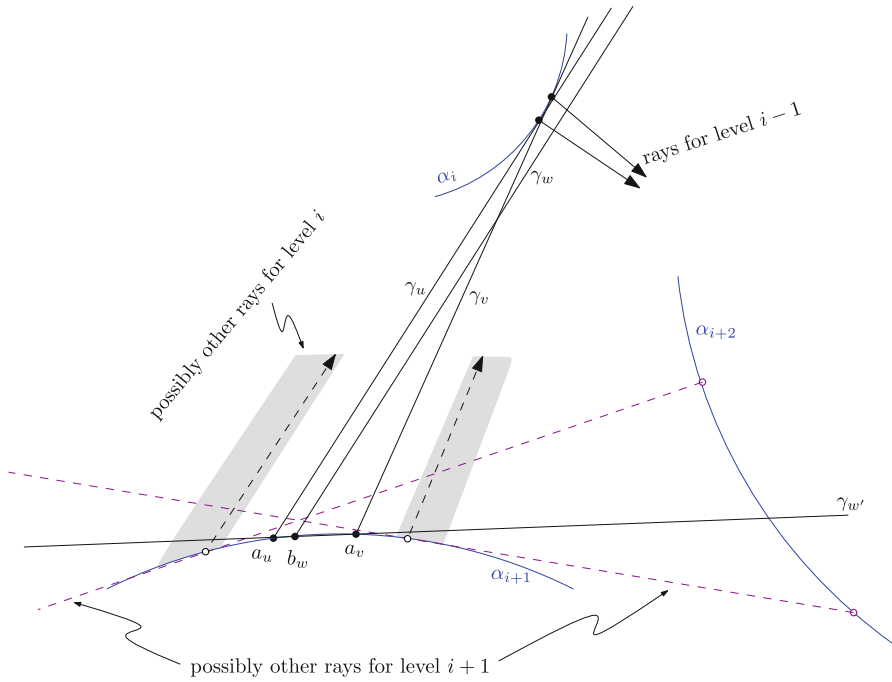


Fig. 6 Case 2(i) in the proof of Lemma 2.4: adding a path with two internal vertices

introduces an intersection between γ_v and the ray for the parent of v . Note that exactly one of the two options must occur.

Let us consider only scenario (i), since the other one is symmetric; see Fig. 6. We choose a point b_w on α_{i+1} between a_u and a_v very near a_u and represent w with a ray γ_w parallel to γ_u with origin b_w . Thus γ_w does not intersect γ_u but intersects any other ray because we are in scenario (i). Finally, we represent w' with the line $\gamma_{w'}$ through points b_w and a_v . With this, $\gamma_{w'}$ intersects the interior of any other ray but γ_w and γ_v , as desired. \square

Lemma 2.5 Any embedded planar graph G has an even subdivision $T + P$, where T is an embedded tree and P is an admissible extension of T . Furthermore, such T and P can be computed in polynomial time.

Proof Let r be an arbitrary vertex in the outer face f of G . Let B be the set of edges in a BFS tree of G from r . With a slight abuse of notation, we also use B for the BFS tree itself. Let $C = E(G) \setminus B$. In the graph G^* dual to G , the edges $C^* = \{c^* \mid c \in C\}$ are a spanning tree of G^* , which with a slight abuse of notation we also denote by C^* . We root C^* at the node f^* , corresponding to the outer face. This is illustrated in Fig. 7.

We define for each edge $e \in C$ the number k_e of subdivisions it will undertake using a bottom-up approach. Any edge $e \in C$ that is incident to a leaf of C^* gets assigned $k_e = 4$. For any other edge $e \in C$, we define k_e as 2 plus the maximum $k_{e'}$ over all

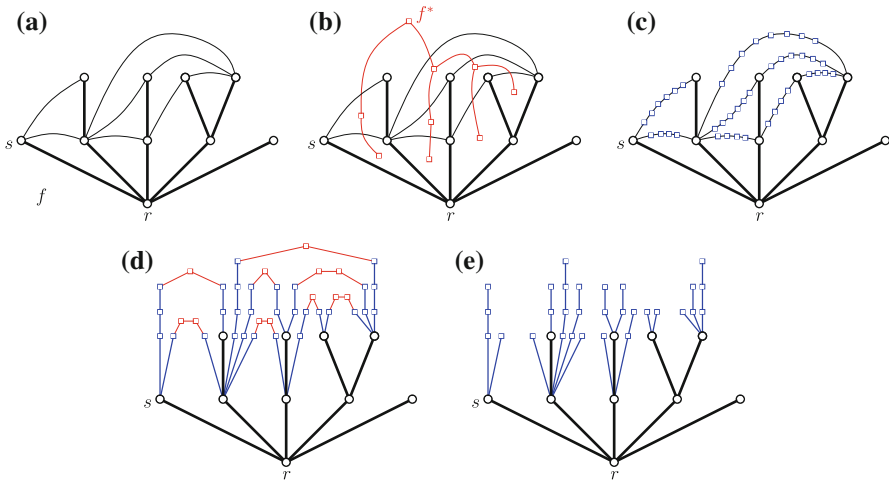


Fig. 7 Figure for the proof of Lemma 2.5. **a** A planar graph with a BFS from r in bold. **b** The spanning tree C^* of the dual graph. **c** The subdivision of the edges from C . **d** The resulting subdivided graph with P marked in red and thinner lines. **e** The resulting graph T after the removal of P , drawn so that the heights of the vertices correspond to their levels (Color figure online)

descendants $(e')^* \in C^*$ of e^* in C^* . Let H be the resulting subdivision. For any edge $e \in C$, let Q_e be the path in H that corresponds to the subdivision of e . We use in H the combinatorial embedding induced by G .

We can now compute the tree T and the paths P . Since B is a BFS tree, every edge $e \in C$ connects two vertices that are either at the same level in B , or at two successive levels. For every edge $e \in C$ that connects two vertices at the same level in B , let P_e be the length-three subpath in the middle of Q_e . For every edge $uv \in C$ that connects vertex u at level i to vertex v at level $i + 1$ in B , let P_e be the length-two subpath in the middle of $Q_e - u$. We then take $P = \bigcup_{e \in C} P_e$ and take T to be the graph $H - P$, after removing isolated vertices. In T we use the rotation system inherited from H and use the edge rs to define the linear order, where rs is an edge in the outer face f .

It is clear from the construction that $T + P = H$ is an even subdivision of G . We have to check that P is indeed an admissible extension of T . The maximal paths of P are vertex disjoint and connect leaves of P because the paths $Q_e, e \in C$, are edge-disjoint and each P_e is strictly contained in the interior of Q_e . Since in $H - P$ we removed isolated vertices, it is clear that no internal vertex of a path of P is in T . The graph T is indeed a tree because, for every edge $e \in C$, we have removed some edges from its subdivision Q_e . Since B is a BFS tree and P_{uv} is centered within Q_{uv} , when u and v are at the same level, or within $Q_{uv} - u$, when u is one level below v , the maximal paths of P connect vertices that are equidistant from r in H .

It remains to show that the endpoints of any maximal path in P are consecutive vertices in T . This is so because of the inductive definition of k_e . The base case is when $e \in C$ is incident to a leaf of C^* , and is subdivided $k_e = 4$ times. The edge e either connects two vertices at the same level in B , or at two successive levels. In both cases it can be checked that P_e connects two consecutive vertices in T . The inductive case is as follows. We let k_e be equal to 2 plus the maximum $k_{e'}$ over all descendants $(e')^* \in C^*$

of e^* in C^* . By induction, all the corresponding $P_{e'}$ connect two consecutive vertices of T , say at level i in T . By definition, P_e will connect two consecutive vertices of T at level $i + 1$.

The construction of T and P only involves computing the BFS tree B , the spanning tree C^* , and the values k_e , which can clearly be done in polynomial time. \square

Combining Lemmas 2.4 and 2.5 directly yields the following.

Theorem 2.6 *Any planar graph has an even subdivision whose complement is a ray intersection graph. Furthermore, this subdivision can be computed in polynomial time.*

2.4 Polynomial-Time Construction

So far, we used real coordinates in our construction. We next argue how to have a construction using only integers with a polynomial number of bits.

Lemma 2.7 *In the construction of Lemma 2.4, if $n = |V(T + P)|$,*

- *any two origins of rays are at distance at least $n^{-O(n)}$ and at most $O(n)$;*
- *the distance between any ray and any other origin is at least $n^{-O(n)}$.*

Proof Recall that the circle containing points p_0, p_1, \dots has radius 1. The rectangles R_i are all congruent and have two diagonals of length $|p_0 p_1| = \Theta(1)$. The small side of rectangle R_i has size $\Theta(1/n)$ and both diagonals of R_i form an angle of $\Theta(1/n)$. It follows that the center of the circle supporting α_i has coordinates $\Theta(n)$. For points from different curves α_i there is at least a separation of $\Theta(1/n)$.

We first bound the distance between the origins for the rays representing vertices of T . Let us refer by *features* on the curve α_i the origins of rays lying on α_i and the extremes of α_i . Let δ_i be the minimum separation between any two features on the curve α_i . In the curve α_1 there are three features: the two extremes of α_1 and the origin a_r of γ_r , which is in the middle of α_1 . Since α_1 have length $\Omega(1)$, it follows that $\delta_1 = \Omega(1)$.

We will bound the ratio δ_{i+1}/δ_i for $i \geq 1$. Consider the construction of Lemma 2.3 to determine the features on α_{i+1} . By induction, any two consecutive features along α_i are separated at least by δ_i . Since α_i is supported by a circle of radius $\Theta(n)$, the lines ℓ_v^+ and ℓ_v^- form an angle of at least $\Omega(\delta_i/n)$. This implies that the points $\ell_v^- \cap \alpha_{i+1}, a_{u_1}, \dots, a_{u_d}, \ell_v^+ \cap \alpha_{i+1}$ have a separation of $\Omega(\delta_i/(nd))$. It follows that $\delta_{i+1} = \Omega(\delta_i/(nd)) = \Omega(\delta_i/n^2)$, and thus $\delta_{i+1}/\delta_i = \Omega(1/n^2)$. Since T has depth at most n , all features are at distance at least $n^{-O(n)}$.

We can now argue that the origins of the rays used in the construction of Lemma 2.4 also have a separation of $n^{-O(n)}$. In Case 1, we just add a line through previous features. In Case 2, it is enough to place b_w at distance $|a_u a_v|/n$ from a_u , and thus the features keep a separation of at least $n^{-O(n)}$.

For the proof of the second item, consider a ray on level i and the point a_v at which it hits α_i . Now consider another origin b on α_i . The angle x made by the line $a_v b$ and the ray is bounded from below (like the angle between ℓ_v^- and ℓ_v^+ in Fig. 4) by n^{-cn} for some constant c , and so is the distance d between a_v and b . Now the distance between

b and the ray is $d \sin x \geq d2x/\pi \geq n^{-2cn}(2/\pi) = n^{-O(n)}$, as claimed, where the first inequality comes from the relation $\sin x \geq 2x/\pi$ for $x \in [0, \pi/2]$. \square

We can now give the following algorithmic version of Lemma 2.4.

Lemma 2.8 *Let T be an embedded tree and let P be an admissible extension of T . We can find in polynomial time a family of rays described with integer coordinates whose intersection graph is isomorphic to the complement of $T + P$.*

Proof Recall that, after constructing the reference frame, each point in the construction is created either as (i) an intersection between a line through two existing points and a curve α_i , or (ii) by equally spacing $O(n)$ points between two existing points on a curve α_i . Consider that each point is moved by a distance $< \varepsilon$ after it is constructed. This causes any point further constructed from those to move as well. In case (ii), the new points would be moved by no more than ε as well. In case (i) however, the error could be amplified. Let a and b be two previously constructed points, and suppose they are moved to a' and b' , within a radius of ε . By the previous lemma, the distance between a and b is at least $n^{-O(n)}$ and so the angle between the line ab and $a'b'$ is at most $\varepsilon n^{O(n)}$. Because the radius of the supporting circle of α_i is $\Theta(n)$, the distance between $ab \cap \alpha_i$ and $a'b' \cap \alpha_i$ is at most $O(n)\varepsilon n^{O(n)} = \varepsilon n^{O(n)}$.

Therefore, an error in one construction step expands by a factor $n^{O(n)}$. Now observe that each point of type (i) is constructed on the next level, and a point of type (ii) is always constructed from points of type (i). Therefore, as there are $O(n)$ levels, an error is propagated at most $O(n)$ times, and therefore, the total error propagated from moving each constructed point by a distance $< \varepsilon$ is at most $\varepsilon n^{O(n^2)}$.

By the previous lemma, there is a constant c such that any origin in the construction of Lemma 2.4 is separated at least by a distance $A = 1/n^{cn}$ from any other origin and any ray not incident to it. Therefore, by choosing $\varepsilon = n^{-c'n^3}$ for c' large enough, the total propagation will never exceed A , and therefore perturbing the basic construction points of the reference frame and each further constructed point by a distance $< \varepsilon$ will not change the intersection graph of the modified rays.

Therefore, to construct the required set of rays in polynomial time, we multiply every coordinate by the smallest power of 2 larger than $1/\varepsilon$ and snap every constructed point to the nearest integer while following the construction of Lemma 2.4. Each coordinate can then be represented by $O(n^3 \log n)$ bits. \square

Let $\alpha(G)$ be the size of the largest independent set in a graph G . The following simple lemma can be deduced from the observation that subdividing an edge twice increases the independence number by exactly one.

Lemma 2.9 *If G' is an even subdivision of G where each edge $e \in E(G)$ is subdivided $2k_e$ times, then $\alpha(G') = \alpha(G) + \sum_e k_e$.* \square

By combining Lemmas 2.5, 2.8, and 2.9, we obtain a polynomial-time reduction from the maximum independent set problem in planar graphs to the maximum clique problem in ray intersection graphs.

Theorem 2.10 *Finding a maximum clique in a ray intersection graph is NP-hard, even when the input is given by a geometric representation as a set of rays.* \square

Acknowledgments Research was supported in part by the Slovenian Research Agency, program P1-0297 and projects J1-4106, L7-4119, BI-BE/11-12-F-002, and within the EUROCORES Programme EUROGIGA (projects GReGAS and ComPoSe) of the European Science Foundation. Large part of the work was done while Sergio was visiting ULB.

References

1. Agarwal, P.K., Mustafa, N.H.: Independent set of intersection graphs of convex objects in 2D. *Comput. Geom.* **34**(2), 83–95 (2006)
2. Ambühl, C., Wagner, U.: The clique problem in intersection graphs of ellipses and triangles. *Theory Comput. Syst.* **38**(3), 279–292 (2005)
3. Bang-Jensen, J., Reed, B., Schacht, M., Šámal, R., Toft, B., Wagner, U.: On six problems posed by Jarik Nešetřil. In: Klazar, M., Kratochvíl, J., Loeb, M., Matoušek, J., Thomas, R., Valtr, P. (eds.) *Topics in Discrete Mathematics. Algorithms and Combinatorics*, vol. 26, pp. 613–627. Springer, Heidelberg (2006)
4. Chalopin, J., Gonçalves, D.: Every planar graph is the intersection graph of segments in the plane. In: 41st Annual ACM Symposium on Theory of Computing, STOC, pp. 631–638. ACM, New York (2009)
5. Engebretsen, L., Holmerin, J.: Clique is hard to approximate within $n^{1-o(1)}$. In: *Proceedings of the 27th International Colloquium on Automata, Languages and Programming (ICALP'00)*, pp. 2–12. Springer, Heidelberg (2000)
6. Fox, J., Pach, J.: Computing the independence number of intersection graphs. In: *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pp. 1161–1165. SIAM, Philadelphia, PA (2011)
7. Francis, M.C., Kratochvíl, J., Vyskočil, T.: Segment representation of a subclass of co-planar graphs. *Discrete Math.* **312**(10), 1815–1818 (2012)
8. Garey, M.R., Johnson, D.S.: The rectilinear Steiner tree problem is NP-complete. *SIAM J. Appl. Math.* **32**, 826–834 (1977)
9. Hastad, J.: Clique is hard to approximate within $n^{1-\epsilon}$. In: *Proceedings of the 37th Annual Symposium on Foundations of Computer Science (FOCS'96)*, pp. 627–636. IEEE Computer Society, Washington, DC (1996)
10. Kratochvíl, J., Kuběna, A.: On intersection representations of co-planar graphs. *Discrete Math.* **178**(1–3), 251–255 (1998)
11. Kratochvíl, J., Matoušek, J.: Intersection graphs of segments. *J. Comb. Theory Ser. B* **62**(2), 289–315 (1994)
12. Kratochvíl, J., Nešetřil, J.: Independent set and clique problems in intersection-defined classes of graphs. *Comment. Math. Univ. Carolinae.* **31**(1), 85–93 (1990)
13. Middendorf, M., Pfeiffer, F.: The max clique problem in classes of string-graphs. *Discrete Math.* **108**(1–3), 365–372 (1992)