

Colouring Polygon Visibility Graphs and Their Generalizations

James Davies ✉

Department of Combinatorics and Optimization, School of Mathematics,
University of Waterloo, Canada

Tomasz Krawczyk ✉

Department of Theoretical Computer Science, Faculty of Mathematics and Computer Science,
Jagiellonian University, Kraków, Poland

Rose McCarty ✉

Department of Combinatorics and Optimization, School of Mathematics,
University of Waterloo, Canada

Bartosz Walczak ✉

Department of Theoretical Computer Science, Faculty of Mathematics and Computer Science,
Jagiellonian University, Kraków, Poland

Abstract

Curve pseudo-visibility graphs generalize polygon and pseudo-polygon visibility graphs and form a hereditary class of graphs. We prove that every curve pseudo-visibility graph with clique number ω has chromatic number at most $3 \cdot 4^{\omega-1}$. The proof is carried through in the setting of ordered graphs; we identify two conditions satisfied by every curve pseudo-visibility graph (considered as an ordered graph) and prove that they are sufficient for the claimed bound. The proof is algorithmic: both the clique number and a colouring with the claimed number of colours can be computed in polynomial time.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Visibility graphs, χ -boundedness, pseudoline arrangements, ordered graphs

Digital Object Identifier 10.4230/LIPIcs.SoCG.2021.29

Related Version *Full Version:* <https://arxiv.org/abs/2103.07803>

Funding Tomasz Krawczyk and Bartosz Walczak were partially supported by the National Science Centre of Poland grant 2015/17/D/ST1/00585.

Acknowledgements We thank Bodhayan Roy for sharing the problem of whether polygon visibility graphs are χ -bounded.

1 Introduction

A *polygon* is a Jordan curve made of finitely many line segments. A *polygon visibility graph* is the graph on the set of vertices of a polygon P that has an edge between each pair of *mutually visible* vertices, which means that the line segment connecting them is disjoint from the exterior of P . A class of graphs is χ -*bounded* if there is a function that bounds the chromatic number in terms of the clique number for every graph in the class. A clique in a polygon visibility graph has a natural interpretation – it is the maximum size of a subset of the vertices whose convex hull is disjoint from the exterior of the polygon (see Figure 1, top-left). The starting point of and main motivation for this work is the question of Kára, Pór, and Wood [23] of whether the class of polygon visibility graphs is χ -bounded. We answer it in the affirmative.

► **Theorem 1.1.** *Every polygon visibility graph with clique number ω has chromatic number at most $3 \cdot 4^{\omega-1}$.*



© James Davies, Tomasz Krawczyk, Rose McCarty, and Bartosz Walczak;
licensed under Creative Commons License CC-BY 4.0


37th International Symposium on Computational Geometry (SoCG 2021).

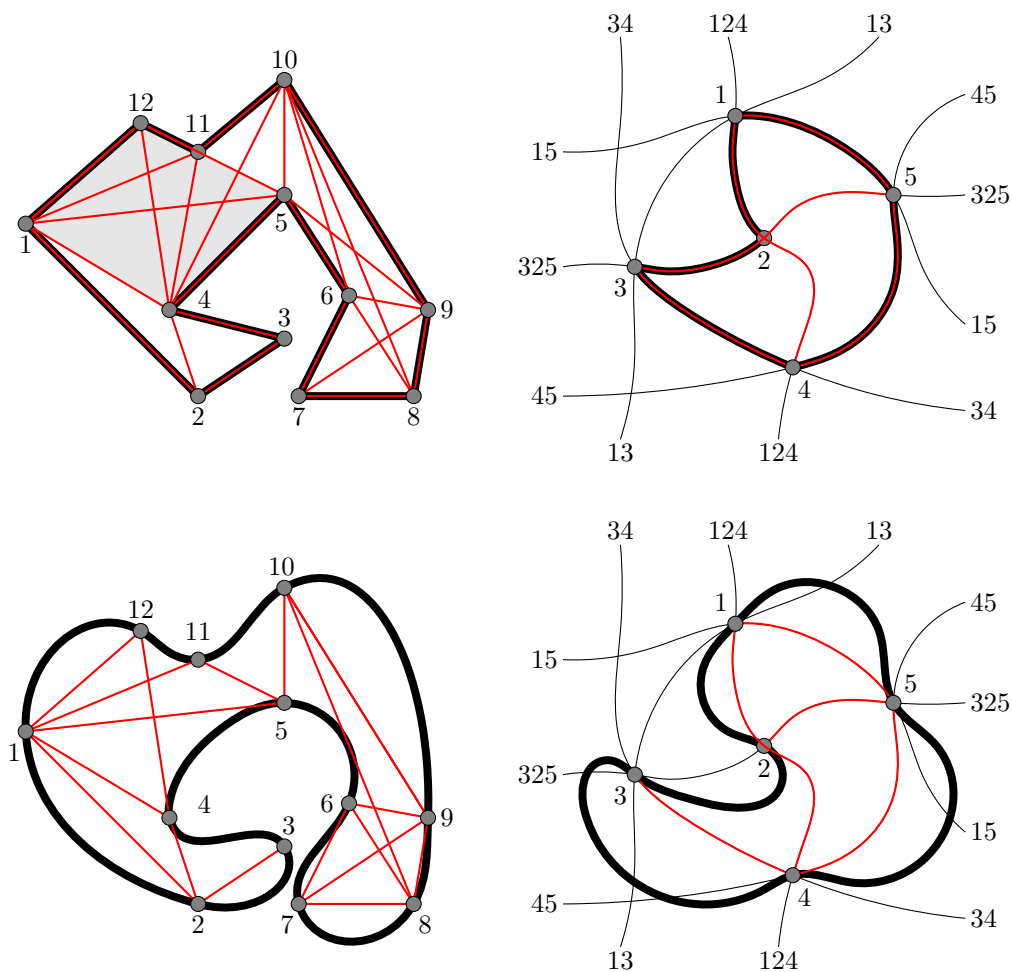
Editors: Kevin Buchin and Éric Colin de Verdière; Article No. 29; pp. 29:1–29:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany






■ **Figure 1** From left to right: a polygon visibility graph (where the convex hull of a maximum clique is shaded), a pseudo-polygon visibility graph, a curve visibility graph, and a curve pseudo-visibility graph. A “visibility” between each pair of adjacent vertices is drawn with a red (pseudo-)segment.

The bound in Theorem 1.1 also holds for all induced subgraphs of polygon visibility graphs. Such graphs can be defined alternatively as *curve visibility graphs*, that is, visibility graphs of points on a Jordan curve, where two points are considered to be mutually visible if the line segment connecting them is disjoint from the exterior of the curve (see Figure 1, bottom-left).

O’Rourke and Streinu [26] studied visibility graphs of pseudo-polygons (polygons on pseudoline arrangements; see Figure 1, top-right), where two vertices of the polygon are considered to be mutually visible if the pseudoline segment connecting them in the arrangement is disjoint from the exterior of the polygon. As a common generalization of these graphs and curve visibility graphs, we define curve pseudo-visibility graphs as follows. For a pseudoline arrangement \mathcal{L} , a Jordan curve K , and a finite set V of points on K any two of which lie on a common pseudoline in \mathcal{L} , the *curve pseudo-visibility graph* $G_{\mathcal{L}}(K, V)$ has vertex set V and has an edge between each pair of vertices such that the pseudoline segment in \mathcal{L} connecting them is disjoint from the exterior of K (see Figure 1, bottom-right). We elaborate on this notion in Section 2; in particular, we show that curve pseudo-visibility graphs are exactly the induced subgraphs of the visibility graphs of pseudo-polygons. With this notion in hand, we provide the following topological generalization of Theorem 1.1.



■ **Figure 2** A graph in \mathcal{H} (left), an ordered hole (middle), and the forbidden configuration for a capped graph (right). Dashed lines indicate non-edges. The pairs of vertices where no lines are drawn can be edges or non-edges.

► **Theorem 1.2.** *Every curve pseudo-visibility graph with clique number ω has chromatic number at most $3 \cdot 4^{\omega-1}$.*


To prove Theorem 1.2 (and thus Theorem 1.1), we turn our attention to ordered graphs, where an *ordered graph* is a pair (G, \prec) such that G is a graph and \prec is a linear order on the vertices of G . A curve pseudo-visibility graph comes with a natural linear order on the vertices (determined up to rotation), which makes it an ordered graph; it is the order in which the vertices are encountered when following the Jordan curve in the counterclockwise direction starting from an arbitrarily chosen vertex. An ordered graph (H, \prec_H) is an (*induced*) *ordered subgraph* of an ordered graph (G, \prec) if H is a subgraph (an induced subgraph, respectively) of G and \prec_H is the restriction of \prec to the vertices of H . There are two natural families of ordered obstructions to (that is, ordered graphs that cannot occur as induced ordered subgraphs of) curve pseudo-visibility graphs: the family \mathcal{H} that we define in Section 3 and the family of ordered holes (see Figure 2), both easily verifiable in polynomial time. We prove the following further generalization of Theorem 1.2.

► **Theorem 1.3.** *Every \mathcal{H} -free ordered graph with clique number $\omega \geq 2$ has chromatic number at most $3 \cdot 4^\omega(\omega - 1)$ in general and at most $3 \cdot 4^{\omega-1}$ when also ordered-hole-free. Moreover, there is a polynomial-time algorithm that takes in an \mathcal{H} -free ordered graph and computes its clique number ω and a colouring with the claimed number of colours.*

Our proofs of Theorems 1.1–1.3 ultimately lead to the class of capped graphs, which may be of independent interest. A *capped graph* is an ordered graph (G, \prec) such that for any four vertices $a \prec b \prec c \prec d$, if $ac, bd \in E(G)$, then $ad \in E(G)$; see Figure 2 (right). (This condition was previously studied for terrain visibility graphs [1, 3] as the “ X -property”.) We show that the vertices of any \mathcal{H} -free ordered graph can be partitioned into three sets each inducing a capped graph. This way, Theorem 1.3 becomes a corollary to the following.

► **Theorem 1.4.** *Every capped graph with clique number $\omega \geq 2$ has chromatic number at most $4^\omega(\omega - 1)$ in general and at most $4^{\omega-1}$ when also ordered-hole-free. Moreover, there is a polynomial-time algorithm that takes in a capped graph and computes its clique number ω and a colouring with the claimed number of colours.*

Any improvement on the bounds in Theorem 1.4 would immediately imply corresponding improvements in Theorems 1.1–1.3. A major open problem for most known χ -bounded classes of graphs is whether they are polynomially χ -bounded, that is, whether the chromatic number of the graphs in the class is bounded by a polynomial function of their clique number. Esperet [17] conjectured that *every* hereditary class of graphs that is χ -bounded is polynomially χ -bounded. While we have little faith in this conjecture, we do expect that it holds for capped graphs (and, consequently, for the graphs considered in Theorems 1.1–1.3).



■ **Figure 3** The banana B_4 (left) and the ordered graph X (right).

► **Conjecture 1.5.** *There is a polynomial function p such that every capped graph with clique number ω has chromatic number at most $p(\omega)$.*

While our proof of Theorem 1.4 is direct, we remark that a recent result of Scott and Seymour [32] implies χ -boundedness (with a much weaker bound) of the significantly broader class of X -free ordered graphs, that is, ordered graphs excluding the four-vertex ordered graph X illustrated in Figure 3 (right) as an induced ordered subgraph. In particular, every capped graph is X -free. Tomon [35] conjectured that the class of X -free ordered graphs is χ -bounded. This statement implies not only Theorem 1.4 but also the theorem of Rok and Walczak [31] that so-called outerstring graphs are χ -bounded. This is because outerstring graphs (with the natural linear order on the vertices) are easily seen to be X -free. Scott and Seymour [32] proved that for every graph H that is a “banana” (or more generally – a “banana tree”), the class of graphs excluding all subdivisions of H as induced subgraphs is χ -bounded. Figure 3 (left) shows an example of a “banana” B_4 with the property that no subdivision of B_4 can be made X -free under any order of the vertices. This shows that the aforementioned result of Scott and Seymour implies Tomon’s conjecture.

► **Theorem 1.6.** *The class of X -free ordered graphs is χ -bounded.*

We present a detailed proof of Theorem 1.6 in the full version. We conclude the introduction with a brief literature review in order to place Theorems 1.1–1.4 and 1.6 in context.

Geometric graph classes and χ -boundedness

Various classic examples of χ -bounded graph classes are defined in terms of geometric representations. For instance, intersection graphs of axis-parallel rectangles [5] and circle graphs [22] are χ -bounded. Most of the literature in this direction focuses on intersection or disjointness graphs of objects in the plane. While the class of intersection graphs of curves in the plane is not itself χ -bounded [28], some very general subclasses are [14, 30]. There are also very precise results for disjointness graphs of certain kinds of curves in the plane [27].

Less is known about χ -boundedness of visibility graphs, even though various kinds of such graphs have been considered in the literature – see [20] for a survey. Kára, Pór, and Wood [23] conjectured that the class of point visibility graphs is χ -bounded, but this was disproven by Pfender [29]. Some types of bar visibility graphs are related to interval graphs [15] and planar graphs [25] and are therefore known to be χ -bounded.

Axenovich, Rollin, and Ueckerdt [7] considered the problem of whether ordered graphs excluding a fixed ordered graph (H, \prec) as an ordered subgraph (not necessarily induced) have bounded chromatic number; they showed various cases of (H, \prec) for which the answers are positive and negative. In particular, the answer is negative if H contains a cycle (as it is for unordered graphs), but they showed it is also negative for some acyclic ordered graphs

(H, \prec) . Pach and Tomon [27] used some specific classes of forbidden induced ordered graphs as a tool for studying χ -boundedness of disjointness graphs of curves. Max point-tolerance graphs [13] and classes of graphs of bounded twin-width [9] are also known to be χ -bounded and have well-understood characterizations as ordered graphs.

Algorithmic considerations

The class of curve pseudo-visibility graphs is hereditary, whereas most well-known classes of visibility graphs are not, including the classes of point visibility graphs, polygon visibility graphs, and pseudo-visibility graphs. The condition of the class being hereditary is very natural to impose when studying χ -boundedness and implies that curve pseudo-visibility graphs can be characterized by excluded induced (ordered) subgraphs. There has been a good deal of work on the characterization and recognition problems, but for point visibility graphs and polygon visibility graphs the problems appear to be hard (see [12] and [19]).

These difficult characterization problems tend to become tractable, and have more natural solutions, in the “pseudo-visibility setting” [1, 2, 18, 26]. This is due to the connection between stretchability of pseudoline arrangements and representability of rank-3 oriented matroids. So the pseudo-visibility setting is more combinatorial because it suffices to find the associated rank-3 oriented matroid without worrying about representability. Representability provides a real difficulty; the pseudo-visibility setting is strictly more general for polygon visibility graphs [33], even when certain restrictions are imposed [3].

It is an interesting problem to characterize ordered curve pseudo-visibility graphs by excluded induced ordered subgraphs. The two aforementioned classes of obstructions (\mathcal{H} and the ordered holes) are likely to be insufficient for a full characterization – they roughly correspond to the first two of the four necessary conditions for an ordered graph to be a polygon visibility graph described by Ghosh [19]. Nevertheless, we conjecture the following.

► **Conjecture 1.7.** *Ordered curve pseudo-visibility graphs can be recognized in polynomial time.*

The part of Theorem 1.3 concerning polynomial-time computation of clique number extends well-known results regarding polygon visibility graphs [6, 16, 21], although our algorithm is certainly slower. We cannot expect to get an exact algorithm for the chromatic number, as Çağırıcı, Hliněný, and Roy [11] proved that it is NP-complete to decide if a polygon visibility graph is 5-colourable, even when the polygon is provided as part of the input.

2 Curve pseudo-visibility graphs

A *pseudoline* is a simple curve which separates the plane into two unbounded regions. A *pseudoline arrangement* is a set of pseudolines such that each pair intersects in exactly one point, where they cross. A *pseudo-configuration* is a pair (\mathcal{L}, V) such that \mathcal{L} is a pseudoline arrangement and V is a (finite) set of points on $\bigcup \mathcal{L}$ with the property that any two points in V lie on a common pseudoline in \mathcal{L} (which is therefore unique). A pseudo-configuration (\mathcal{L}, V) is in *general position* if no three points in V lie on a common pseudoline in \mathcal{L} .

Let (\mathcal{L}, V) be a pseudo-configuration and K be a Jordan curve passing through all points in V . The *exterior* of K is the unbounded component of $\mathbb{R}^2 \setminus K$. We say that two points $u, v \in V$ are *mutually visible* in K if the pseudoline segment in \mathcal{L} connecting u and v is disjoint from the exterior of K . The *curve pseudo-visibility graph* $G_{\mathcal{L}}(K, V)$ has vertex set V and has an edge uv for each pair of vertices $u, v \in V$ that are mutually visible in K . The curve K is a *pseudo-polygon* on \mathcal{L} with vertex set V if every segment of K between two consecutive points in V is contained in a single pseudoline in \mathcal{L} . Graphs of the form

$G_{\mathcal{L}}(K, V)$ where K is a pseudo-polygon on \mathcal{L} with vertex set V and (\mathcal{L}, V) is in general position were considered by O'Rourke and Streinu [26] as *pseudo-polygon visibility graphs*. As we will see, the general position assumption is not actually restrictive in this setting.

The following two propositions imply that curve pseudo-visibility graphs are exactly the induced subgraphs of pseudo-polygon visibility graphs. First we find a pseudo-polygon, and then we take care of the general position assumption.

► **Proposition 2.1.** *For every curve pseudo-visibility graph $G = G_{\mathcal{L}}(K, V)$, there exist a pseudo-configuration (\mathcal{L}', V') and a pseudo-polygon K' on \mathcal{L}' with vertex set V' such that $\mathcal{L} \subseteq \mathcal{L}'$, $V \subseteq V'$, the points in V occur in the same cyclic order on K' as on K , and G is the subgraph of $G_{\mathcal{L}'}(K', V')$ induced on V .*


Proof. We can assume that K intersects $\bigcup \mathcal{L}$ only finitely many times. To see this, consider the finite plane graph H with a vertex for each intersection point of two pseudolines in \mathcal{L} (including the points in V) and with an edge for each pseudoline segment in \mathcal{L} connecting two vertices and passing through no other vertex. Let H' be the vertex-spanning subgraph of H obtained by including only the edges whose pseudoline segment is disjoint from the exterior of K . Thus K is contained in the closure of the outer (unbounded) face of H' . By following the boundary of this outer face very closely (and making thin connections between connected components of the boundary if H' is disconnected), we can choose K to intersect $\bigcup \mathcal{L}$ only finitely many times while preserving the graph $G_{\mathcal{L}}(K, V)$ and the order of points on K .

Let (\mathcal{L}^*, V^*) be a pseudo-configuration such that $\mathcal{L} \subset \mathcal{L}^*$, $V \subset V^* \subset K$, every open segment of K connecting two points in $\bigcup \mathcal{L}$ contains a point in $V^* \setminus \bigcup \mathcal{L}$, each point in V^* lies on at least two pseudolines in \mathcal{L}^* , and $|V^*| \geq 3$; we first select V^* and then extend \mathcal{L} to \mathcal{L}^* using Levi's extension lemma [24]. As before, we can assume that K intersects $\bigcup \mathcal{L}^*$ only finitely many times. We further assume that K has the minimum number of intersection points with $\bigcup \mathcal{L}^*$ among all Jordan curves K^* that pass through all of the points in V^* in the same order as K and are such that $G = G_{\mathcal{L}}(K^*, V)$.

Let $V' = K \cap \bigcup \mathcal{L}^*$. In particular, $V^* \subseteq V'$. We extend \mathcal{L}^* to a family of pseudolines \mathcal{L}' such that (\mathcal{L}', V') is a pseudo-configuration, using Levi's extension lemma [24]. For any two points $u, v \in V'$ consecutive on K , let K_{uv} be the segment of K between u and v (which is internally disjoint from $\bigcup \mathcal{L}^*$), let L'_{uv} be the pseudoline in \mathcal{L}' passing through u and v , let K'_{uv} be the segment uv of L'_{uv} , and let E_{uv} be the unbounded component of $\mathbb{R}^2 \setminus (K_{uv} \cup K'_{uv})$. To construct K' , we replace K_{uv} by K'_{uv} for every pair of points $u, v \in V'$ consecutive on K . Since any pseudoline in \mathcal{L}^* intersecting K'_{uv} needs to intersect K_{uv} , every pseudoline in $\mathcal{L}^* \setminus \{L'_{uv}\}$ is fully contained in $E_{uv} \cup \{u, v\}$. Consequently, since each point in V^* lies on at least two pseudolines in \mathcal{L}^* , we have $V^* \subset E_{uv} \cup \{u, v\}$.

We claim that $V' \subset E_{uv} \cup \{u, v\}$ as well. If $L'_{uv} \notin \mathcal{L}^*$, then indeed $V' \subset \bigcup \mathcal{L}^* \subset E_{uv} \cup \{u, v\}$. Now, suppose $L'_{uv} \in \mathcal{L}^*$. We have $u \notin \bigcup \mathcal{L}$ or $v \notin \bigcup \mathcal{L}$ by the choice of V^* , and thus $L'_{uv} \notin \mathcal{L}$. Suppose $K \setminus K_{uv} \not\subset E_{uv}$. Since $\bigcup \mathcal{L} \subseteq \bigcup (\mathcal{L}^* \setminus \{L'_{uv}\}) \subset E_{uv} \cup \{u, v\}$, $V^* \subset E_{uv} \cup \{u, v\}$, and $K \setminus K_{uv}$ is disjoint from K_{uv} , the parts of $K \setminus K_{uv}$ not lying in E_{uv} can be moved into E_{uv} decreasing the number of intersection points with $\bigcup \mathcal{L}^*$ (as $|V^*| \geq 3$) while preserving the graph $G_{\mathcal{L}}(K, V)$, which contradicts the choice of K . Thus $V' \subset (K \setminus K_{uv}) \cup \{u, v\} \subset E_{uv} \cup \{u, v\}$ when $L'_{uv} \in \mathcal{L}^*$.

For any two pairs $u, v \in V'$ and $u', v' \in V'$ of consecutive points on K , if the internal parts of K'_{uv} and $K'_{u'v'}$ intersect, then the four points u, u', v, v' occur in this or the reverse order on the boundary of $(E_{uv} \cup \{u, v\}) \cap (E_{u'v'} \cup \{u', v'\})$, so the internal parts of K_{uv} and $K_{u'v'}$ intersect, which is impossible. Thus K' is a Jordan curve – a pseudo-polygon on \mathcal{L}' with vertex set V' . Furthermore, $\bigcup \mathcal{L} \subset \bigcap_{uv} (E_{uv} \cup \{u, v\})$, which implies $G_{\mathcal{L}}(K', V) = G_{\mathcal{L}}(K, V)$. ◀



■ **Figure 4** Replacing L with a bundle of pseudo-lines \mathcal{B}_L in the proof of Proposition 2.2.


► **Proposition 2.2.** *For every curve pseudo-visibility graph $G = G_{\mathcal{L}}(K, V)$, there exist a pseudo-configuration (\mathcal{L}', V') in general position and a pseudo-polygon K' on \mathcal{L}' with vertex set V' such that $V \subseteq V'$, the points in V occur in the same cyclic order on K' as on K , and G is the subgraph of $G_{\mathcal{L}'}(K', V')$ induced on V .*

Proof. By Proposition 2.1, we can assume without loss of generality that K is a pseudo-polygon on \mathcal{L} . Suppose there is a pseudoline L in \mathcal{L} passing through more than two points in V . We show that L can be replaced in \mathcal{L} by a bunch \mathcal{B}_L of pseudolines in a small neighbourhood of L so that the set $(\mathcal{L} \setminus \{L\}) \cup \mathcal{B}_L$ is a pseudoline arrangement and the following conditions hold for any two distinct points $u, v \in V \cap L$.

1. There is a pseudoline $L_{uv} \in \mathcal{B}_L$ passing through u, v , and no other points in V .
 2. If u and v are consecutive points of $V \cap L$ on L , then the segment uv of L_{uv} coincides with the segment uv of L .
 3. If the segment uv of L is disjoint from the exterior of K , then so is the segment uv of L_{uv} .
 4. If the segment uv of L intersects the exterior of K , then so does the segment uv of L_{uv} .
- Condition 4 is automatically satisfied whenever we make \mathcal{B}_L lie in a sufficiently small neighbourhood of L . Applying this replacement repeatedly for every such pseudoline L yields a claimed pseudoline arrangement \mathcal{L}' .

For the replacement step, assume without loss of generality that L is a vertical line (by applying an appropriate homeomorphism of the plane before and the inverse homeomorphism after the step). Enumerate the points in $V \cap L$ as v_0, \dots, v_k from bottom to top. Let C be the circle with vertical diameter v_0v_k . Let $v'_0 = v_0$ and $v'_k = v_k$. For $0 < i < k$, let H_i be the horizontal line through v_i , and let v'_i be the left/the right/any intersection point of C and H_i if the exterior of K touches v_i from the left side/the right side/both sides of the vertical line L (respectively). For $0 \leq i < j \leq k$, let $L'_{i,j}$ be the straight line passing through v'_i and v'_j . The bundle \mathcal{B}_L is obtained by “flattening” the family of lines $\{L'_{i,j}\}_{0 \leq i < j \leq k}$ horizontally to fit it in a small neighbourhood of L and performing local horizontal shifts to guarantee conditions 1 and 2; condition 3 then follows. See Figure 4 for an illustration. ◀

Recall that an *ordered graph* is a tuple (G, \prec) such that G is a graph and \prec is a linear order on its vertex set. While it is more convenient to work with linear orders, the points on a Jordan curve are really ordered cyclically. A *rotation* of a linear order \prec is any linear order obtained from \prec by repeatedly making the largest element the smallest. We think of



■ **Figure 5** A crossing sequence from u to v (left) and from v to u (right).

any finite set of points V on a Jordan curve K as being ordered counterclockwise around K , as in Figure 1 (bottom-left). We call any linear order which begins at an arbitrary point in V and then follows K in the counterclockwise direction a *natural order* of V on K . A curve pseudo-visibility graph $G_{\mathcal{L}}(K, V)$ along with a natural order of V on K forms an *ordered curve pseudo-visibility graph*.

If (G, \prec) is an ordered graph with vertices $a \prec b \prec c \prec d$ and edges ac and bd , we say that ac *crosses* bd , that ac and bd are *crossing edges*, and that bd is *crossed* by ac . Two edges which are not crossing are called *non-crossing*. The property that a pair of edges is crossing/non-crossing is preserved under rotation (since, using this terminology, we do not specify which edge crosses the other). In particular, it is well defined for an ordered curve pseudo-visibility graph regardless of the choice of a natural ordering.

► **Lemma 2.3.** *For a curve pseudo-visibility graph $G_{\mathcal{L}}(K, V)$ with (\mathcal{L}, V) in general position, two distinct edges uv and xy are crossing if and only if the open segments uv and xy of pseudolines in \mathcal{L} intersect.*


Proof. If uv and xy are crossing edges, then the open segments uv and xy must intersect; otherwise K along with uv and xy give an outerplanar drawing of K_4 , which is impossible. If uv and xy are non-crossing while the open segments uv and xy intersect, then we can again obtain an outerplanar drawing of K_4 by re-connecting uv and xy in a sufficiently small neighbourhood of their unique intersection point – a contradiction. ◀

3 Obstructions for curve pseudo-visibility graphs

In this section, we discuss the obstructions mentioned in the introduction: the class \mathcal{H} and the class of ordered holes. Ghosh [19] observed that these are obstructions for polygon visibility graphs, and related obstructions in the pseudo-visibility setting appear in the works of Abello and Kumar [2] and O’Rourke and Streinu [26].

Let u and v be two distinct vertices in an ordered graph (G, \prec) . If $u \prec v$, then a *crossing sequence* from u to v is a sequence of distinct edges e_1, \dots, e_k such that u is the smaller end of e_1 , v is the larger end of e_k , and e_i crosses e_{i+1} for $1 \leq i < k$. The notion of a crossing sequence is invariant under rotation of \prec as long as $u \prec v$. If $v \prec u$, then a *crossing sequence* from u to v is a crossing sequence from u to v in any rotation \prec' of \prec such that $u \prec' v$. These definitions should be thought of cyclically; whichever vertex is smaller, a crossing sequence from u to v begins at u and goes counterclockwise until it hits v (see Figure 5). If u and v are adjacent, then the edge uv is a crossing sequence from u to v and from v to u .

► **Lemma 3.1.** *If (G, \prec) is an ordered graph with vertices $a \prec b \prec c \prec d$ and there are crossing sequences from a to c and from b to d , then there is a crossing sequence from a to d .*



■ **Figure 6** A pseudo-polygon with five articulation points: 1, 3, 4, 5 are convex, 2 is concave.

Proof. Let e_1, \dots, e_k and f_1, \dots, f_t be crossing sequences from a to c and from b to d , respectively. Let e_i be the edge with the smallest index such that its larger end, say v , is greater than b in \prec . Let f_j be the edge with the largest index such that its smaller end is less than v in \prec . Then e_i crosses f_j and $e_1, \dots, e_i, f_j, \dots, f_t$ is a crossing sequence from a to d . ◀

The first family of obstructions, which we denote by \mathcal{H} , is defined as follows: \mathcal{H} is the family of all ordered graphs containing two non-adjacent vertices u and v such that there exist a crossing sequence from u to v and a crossing sequence from v to u . See Figure 2 (left) for an illustration. The second family of obstructions is the family of ordered holes. An *ordered hole* is an ordered graph (H, \prec) on vertex set $V(H) = \{c_1, \dots, c_k\}$, where $k \geq 4$ and $c_1 \prec \dots \prec c_k$, with edge set $E(H) = \{c_1c_2, \dots, c_{k-1}c_k, c_kc_1\}$; see Figure 2 (middle).

► **Proposition 3.2.** *Every ordered curve pseudo-visibility graph is \mathcal{H} -free.*

► **Proposition 3.3.** *Every ordered curve pseudo-visibility graph is ordered-hole-free.*

We prove Proposition 3.2 later in this section, and the proof of Proposition 3.3 is in the full version. Before delving into the proof of the former, we show that we can test in polynomial time whether a given ordered graph is free of the considered obstructions.

► **Proposition 3.4.** *There is a polynomial-time algorithm which takes in an ordered graph (G, \prec) and determines whether (G, \prec) is \mathcal{H} -free.*

Proof. It suffices to test, for any two non-adjacent vertices u and v , whether (G, \prec) has a crossing sequence from u to v . We assume that $u \prec v$ after possibly performing a rotation. We create a directed graph \vec{H} with a vertex for each edge of G and with an arc from e to f for each pair of edges of G such that e crosses f . Then there is a crossing sequence from u to v in (G, \prec) if and only if there is an edge e with smaller end u and an edge f with larger end v such that \vec{H} has a directed path from e to f . ◀

► **Proposition 3.5.** *There is a polynomial-time algorithm which takes in an ordered graph (G, \prec) and determines whether (G, \prec) has an ordered hole.*

Proof. It suffices to test, for any two adjacent vertices $u \prec v$ of G , whether u and v are the first and last vertices of an ordered hole. This can be done by removing all vertices in a triangle with u and v and then testing for a directed path from u to v in the natural digraph. ◀

The proof of Proposition 3.2 requires some preparation. Let K be a pseudo-polygon on \mathcal{L} . A *segment* of K is a part of K that is contained in some pseudoline $L \in \mathcal{L}$ and connects two distinct intersection points of L with other pseudolines in \mathcal{L} . An *articulation point* of K is a point in K that joins two segments of K contained in distinct pseudolines in \mathcal{L} . Such an articulation point of K is *convex* if those two pseudolines extend to the exterior of K at p , and it is *concave* if they extend to the interior of K ; see Figure 6. The following lemma was proven by Arroyo, Bensmail, and Richter [4]; we provide a proof for the reader's convenience.

► **Lemma 3.6.** *Every pseudo-polygon on \mathcal{L} has at least three convex articulation points.*

Proof. Suppose otherwise, and choose a counterexample K with as few articulation points as possible. Since \mathcal{L} is a pseudoline arrangement, K has at least three articulation points. Thus, we can choose consecutive articulation points p_1, p_2 , and p_3 which occur on K in that order counterclockwise so that p_2 is concave and if any articulation point is convex, then p_3 is convex. Now, walk from p_1 towards p_2 along the pseudoline $L \in \mathcal{L}$ passing through p_1 and p_2 , and continue walking on L beyond p_2 (through the interior of K , as p_2 is concave) until hitting K at a point $a \in K \cap L$. Let K' denote the pseudo-polygon formed by the segment p_2a of L and the part of K from a to p_2 counterclockwise. It follows that K' has at most two convex articulation points and has fewer articulation points than K , because at most one articulation point, a , is gained, and the articulation points p_2 and p_3 of K are lost. This is a contradiction, completing the proof. ◀

► **Lemma 3.7.** *Let K be a pseudo-polygon on \mathcal{L} . Let u and v be distinct points on K such that \mathcal{L} contains a pseudoline L passing through u and v . If all articulation points of K other than possibly u and v are convex, then the segment uv of L is disjoint from the exterior of K .*

Proof. First, suppose that neither u nor v is a concave articulation point of K . Suppose for the sake of contradiction that the segment uv of L is not disjoint from the exterior of K , and let xy be a maximal subsegment of it with internal part contained in the exterior of K . Thus $x, y \in K$. The segment xy of L together with one of the parts of K between x and y forms a pseudo-polygon on \mathcal{L} with interior contained in the exterior of K and with at most two convex articulation points: x and y . This contradicts Lemma 3.6.

Now, suppose that u is a concave articulation point of K while v is not. Let L' be a pseudoline containing one of the two segments of K incident to u . Follow L' from u in the other direction (towards the interior of K) until it hits K at some point x . Let K' be a pseudo-polygon formed by the segment ux of L' and the part of K between x and u that contains the point v . Thus x is a convex articulation point of K' , u is no longer a concave articulation point of K' , and every other articulation point of K that lies on K' remains convex on K' . Therefore, as we shown in the first case, the segment uv of L is disjoint from the exterior of K' , so it is disjoint from the exterior of K .

The argument is analogous if v is a concave articulation point of K , except that when u is also a concave articulation point of K , then we apply the same argument as above to reduce to the case that only one of u, v is a concave articulation point of K . ◀

Proof of Proposition 3.2. Let $G = G_{\mathcal{L}}(K, V)$ be a curve pseudo-visibility graph and \prec be a natural order of V on K . By Proposition 2.2, we can assume that (\mathcal{L}, V) is in general position. For an edge $e = uv \in E(G)$, let ℓ_e denote the open segment uv of the pseudoline in \mathcal{L} passing through u and v . Suppose that there are $u, v \in V$ with $u \prec v$ such that there are crossing sequences e_1, \dots, e_k from u to v and f_1, \dots, f_t from v to u . Choose the two crossing sequences so that $k + t$ is minimum. We need to show that uv is an edge of G .

By minimality and Lemmas 2.3 and 3.1, for $1 \leq i < j \leq k$, the segments ℓ_{e_i} and ℓ_{e_j} intersect if and only if $j = i + 1$, and likewise for the crossing sequence f_1, \dots, f_t . Also by Lemma 2.3, each ℓ_{e_i} is disjoint from each ℓ_{f_j} . Therefore, by beginning at u and walking along ℓ_{e_1} until its unique intersection with ℓ_{e_2} is reached, then turning left and walking along ℓ_{e_2} until either v or its unique intersection with ℓ_{e_3} is reached, and so on, we can find an open curve $K_1 \subseteq \bigcup_{i=1}^k \ell_{e_i}$ with ends u and v . Likewise we can find an open curve $K_2 \subseteq \bigcup_{j=1}^t \ell_{f_j}$ with ends v and u . Let $K = K_1 \cup K_2 \cup \{u, v\}$. It follows that K is a pseudo-polygon on \mathcal{L} and all articulation points of K except possibly u and v are convex. Therefore, by Lemma 3.7, the segment ℓ_{uv} is disjoint from the exterior of K , so $uv \in E(G)$. ◀

4 Partitioning into capped graphs

Recall that an ordered graph (G, \prec) is *capped* if the following holds for any four vertices a, b, c, d with $a \prec b \prec c \prec d$: if $ac \in E(G)$ and $bd \in E(G)$, then $ad \in E(G)$. In contrast to previous notions defined in terms of \prec , this one is *not* invariant under rotation of \prec .

► **Lemma 4.1.** *If $u \prec v$ are two adjacent vertices of an \mathcal{H} -free ordered graph (G, \prec) and $X = \{u, v\} \cup \{x \in V(G) : u \prec x \prec v \text{ and } ux, xv \in E(G)\}$, then $(G[X], \prec|_X)$ is a capped graph.*

Proof. If $a, b, c, d \in X$ and $ac, bd \in E(G)$, then ac, bd is a crossing sequence from a to d and du, va (da if $a = u$ or $d = v$) is a crossing sequence from d to a in (G, \prec) , so $ad \in E(G)$. ◀

Lemma 4.1 reduces computing the clique number of an \mathcal{H} -free ordered graph to computing the clique number of the capped graphs $(G[X], \prec|_X)$ for all adjacent pairs of vertices $u \prec v$. Thus, the following proposition allows us to conclude Theorem 1.3 from Theorem 1.4.


► **Proposition 4.2.** *There is a polynomial-time algorithm that takes in an \mathcal{H} -free ordered graph (G, \prec) and partitions its set of vertices into three subsets V_1, V_2 , and V_3 so that for each $i \in \{1, 2, 3\}$, the ordered graph $(G[V_i], \prec|_{V_i})$ is capped.*

The proof of Proposition 4.2 is technically complicated; see the full version of the paper. Below, we sketch the proof for the case that (G, \prec) is an ordered polygon visibility graph. It is based on the “window partition” by Suri [34], which was used in a similar fashion to approximate chromatic variants of the well-known art gallery problem [8, 10].

Proof sketch for polygons. We write pq for the closed line segment connecting points p and q . Let $G = G(P, V)$ be a polygon visibility graph, where P is a polygon with vertex set V . Let \prec be a natural ordering of G . Let x and y be the smallest and the largest vertex in \prec , respectively, so that xy is an edge of P and of G . Let $P_{xy} = (P \cup \text{int } P) \setminus xy$, where $\text{int } P$ is the interior of P . We construct a partition of $V(G)$ into three sets, which we express in terms of a colouring ϕ of $V(G)$ that uses three colours: red, green, and blue. First we describe a procedure that constructs a partition of P_{xy} into “windows”; these windows, as we will see, will be naturally arranged with a tree structure, and the root window will be “based” at xy .

To define the root window, we need to introduce the notion of “visibility from xy ”. We say that a point $p \in P_{xy}$ is *visible from xy* if p lies in the closed half-plane to the left of the line from y to x and there is a point $p' \in xy$ such that $pp' \subseteq P_{xy} \cup xy$. The *window W_{xy} based at xy* consists of all points $p \in P_{xy}$ that are visible from xy . It follows that W_{xy} is a connected subset of P_{xy} ; see Figure 7 for an illustration. This set W_{xy} is the root of the constructed window partition tree.

The points in $P_{xy} \setminus W_{xy}$ form some number (possibly zero) of connected subsets of P_{xy} . It can be shown that each such set is of the form I_{ab} for some polygon I and edge ab of I , where a and b are on P and every point in the segment ab is visible from xy in P . Furthermore,



■ **Figure 7** To the left: a polygon P with window W_{xy} based at xy in red, oriented lines $\overrightarrow{L_{ab}}$ depicted with red arrows and dashed lines, and the left/right-invisible sets I_{ab} in green/blue, respectively. To the right: the final window partition of P_{ab} .

at least one of the points a and b is a vertex of P and the line L_{ab} going through a and b intersects xy ; we direct L_{ab} from a and b towards this intersection point to obtain an oriented line $\overrightarrow{L_{ab}}$. Given this description, we partition the invisible sets I_{ab} into two groups:

- I_{ab} is *left-invisible* if it is “towards the left side” of $\overrightarrow{L_{ab}}$;
- I_{ab} is *right-invisible* if it is “towards the right side” of $\overrightarrow{L_{ab}}$.


We do not give formal definitions, but refer to Figure 7.

Given this partition, it can be shown that there are no mutually visible points in two different left-invisible sets or two different right-invisible sets. Now, for each invisible set I_{ab} , we can recursively obtain a window partition of I_{ab} which is rooted at a window W_{ab} based at ab . The window partition of P_{xy} is then obtained by making each of these windows W_{ab} a *left-child* or a *right-child* of W_{xy} according to whether I_{ab} is left-invisible or right-invisible. The following observation summarizes this construction of the window partition: if two points in different windows W_1 and W_2 are mutually visible, then either W_1 and W_2 are in a parent-child relationship, or there is a window W such that one of W_1, W_2 is a left-child of W and the other is a right-child of W .

Now we show how to obtain the 3-colouring ϕ of the vertex set of (G, \prec) such that each colour class induces a capped subgraph. The property above allows us to colour the windows by three colours (say, red, green, and blue) so that no two points in two different windows of the same colour are mutually visible. We colour the root window, say, by red. Then we extend this colouring on the remaining windows so that the children of each window W obtain a colour different from W and the left children of W are coloured with a different colour from the right children of W . This way every vertex of P other than x and y is coloured. We colour x and y arbitrarily; see Figure 7 (right).

To complete the proof, we need to show that the vertices of P in $W_{xy} \cup \{x, y\}$ induce a capped subgraph of (G, \prec) ; for the other windows we can apply induction. It is well known that the related class of ordered terrain visibility graphs is capped [18, Lemma 1], but we give a proof sketch anyway, because that lemma does not apply directly.

Suppose there are four vertices $a, b, c, d \in W_{xy} \cup \{x, y\}$ such that $a \prec b \prec c \prec d$, $ac \in E(G)$, and $bd \in E(G)$ (it is possible that $a = x$ and/or $d = y$). By the definition of visibility from xy , all four points a, b, c, d are in the closed half-plane to the left of the line from y to x . Let



■ **Figure 8** Possible relations between the segments $a'a, d'd, ac, bd, yx$.

$a', d' \in xy$ be such that the segments $a'a$ and $d'd$ are disjoint from the exterior of P ; see Figure 8 for an illustration. Now, the five segments $a'a, d'd, ac, bd, yx$ divide the plane into a set \mathcal{F} of faces, and exactly one of the faces in \mathcal{F} is unbounded (the outer face). Now, to show that ad is disjoint from the exterior of P , it suffices to prove the following two claims.

- The polygon P is contained in the closure of the outer face of \mathcal{F} and P does not cross (but may touch) any segment in the set $\{a'a, d'd, ac, bd\}$.
- The segment ad is disjoint from the interior of the outer face of \mathcal{F} .

The first claim is quite obvious; see the left side of Figure 8 for an illustration. The second claim can be proven by considering all the cases for how the segments $a'a, d'd, ac$, and bd can be placed with respect to each other. We leave the details to the reader. ◀

5 Colouring capped graphs


This section is devoted to the proof of Theorem 1.4 on colouring capped graphs. The proof relies on decompositions; a *decomposition* of an ordered graph (G, \prec) is a collection of subgraphs such that every edge of G belongs to exactly one subgraph in the collection. For a graph G and a set $F \subseteq E(G)$, we write $G[F]$ and $G - F$ for the graph obtained from G by keeping/deleting (respectively) the edges in F .

▶ **Proposition 5.1.** *There is a polynomial-time algorithm which takes in a capped graph (G, \prec) and returns its clique number ω and a decomposition of (G, \prec) into $\omega - 1$ triangle-free capped graphs. If (G, \prec) is additionally ordered-hole-free, then so is each graph in the decomposition.*

Proof. If G is triangle-free, then we return its clique number (1 or 2) and the decomposition consisting of (G, \prec) itself. Thus assume that G is not triangle-free. We say that an edge $uv \in E(G)$ with $u \prec v$ is *triangle-crossed* if v belongs to a triangle with vertices x and y such that $x, y \preceq u$. Let F be the set of all edges that are not triangle-crossed (see Figure 9).

If a, b, c are vertices of G such that $a \prec b \prec c$, $ac, bc \in E(G)$, and ac is triangle-crossed, then bc is triangle-crossed. This implies that $(G[F], \prec)$ is capped and is ordered-hole-free if (G, \prec) is ordered-hole-free. Furthermore, if vertices a, b, c with $a \prec b \prec c$ form a triangle in G , then bc is triangle-crossed. So $G[F]$ is triangle-free. Moreover, we have the following.

▷ **Claim 5.2.** The graph $(G - F, \prec)$ is capped and has clique number exactly one less than the clique number of (G, \prec) . Furthermore, if (G, \prec) is ordered-hole-free, then so is $(G - F, \prec)$.



■ **Figure 9** The decomposition from Proposition 5.1, where darker edges are removed first.

Proof. Let ω denote the clique number of (G, \prec) . If a, b, c are vertices of G such that $a \prec b \prec c$, $ab, ac \in E(G)$, and ab is triangle-crossed, then ac is triangle-crossed. So $(G - F, \prec)$ is capped and is ordered-hole-free if (G, \prec) is. Furthermore, the clique number of $(G - F, \prec)$ is at least $\omega - 1$, because every edge of a clique in (G, \prec) that is not incident to the smallest vertex of the clique is triangle-crossed. Now, suppose that $Q \subseteq V(G)$ is a clique in $G - F$, and let u and v be the two smallest vertices of Q , with $u \prec v$. Then v is in a triangle of G with vertices x and y such that $x \prec y \preceq u$. It follows that $(Q \setminus \{u\}) \cup \{x, y\}$ is a clique in G . This shows that the clique number of $(G - F, \prec)$ is at most $\omega - 1$, as desired. \triangleleft

To conclude, the algorithm proceeds by continuing with $(G - F, \prec)$. The clique number is the number of subgraphs in the decomposition plus one. \blacktriangleleft

Proof of Theorem 1.4. Let $\omega \geq 2$ be the clique number of G , and let $\{(G_i, \prec)\}_{1 \leq i < \omega}$ be a decomposition of (G, \prec) into $\omega - 1$ triangle-free capped subgraphs as in Proposition 5.1. Fix an index i with $1 \leq i < \omega$. If (G, \prec) is ordered-hole-free, then let $F_i = \emptyset$, and otherwise let F_i be the set of edges of (G_i, \prec) which are not crossed in (G_i, \prec) . An ordered graph is *outerplanar* if it has no crossing pair of edges.

▷ **Claim 5.3.** The ordered graph $(G[F_i], \prec)$ is outerplanar, and $(G_i - F_i, \prec)$ is both capped and ordered-hole-free.

Proof. We can assume that (G, \prec) is not ordered-hole-free. That $(G[F_i], \prec)$ is outerplanar is clear from the definition of F_i . If a, b, c are vertices such that $a \prec b \prec c$, $ab, ac \in E(G_i)$, and ab is crossed, then ac is crossed. This implies that the ordered graph $(G_i - F_i, \prec)$ is capped and every ordered hole in it is an ordered hole in (G_i, \prec) . Suppose for the sake of contradiction that vertices $c_1 \prec \dots \prec c_k$ induce an ordered hole in $(G_i - F_i, \prec)$ and thus in (G_i, \prec) . Since the edge $c_{k-2}c_{k-1}$ is crossed, there is an edge $xy \in E(G_i)$ with $x \prec c_{k-2} \prec y \prec c_{k-1}$. It follows that $x \prec c_1$, as c_1, \dots, c_k induce an ordered hole in (G_i, \prec) . We conclude that x, c_{k-1} , and c_k form a triangle in G_i . This contradiction shows that $(G_i - F_i, \prec)$ is ordered-hole-free. \triangleleft

▷ **Claim 5.4.** There is a 4-colouring of $G_i - F_i$, which can be computed in polynomial time.

Proof. We just use the fact that $(G_i - F_i, \prec)$ is triangle-free, capped, and ordered-hole-free. For each component of $G_i - F_i$, we claim that any level of any breadth-first search tree which is rooted at the smallest vertex according to \prec induces a bipartite subgraph. This suffices to complete the proof, as we can reuse colours at every second level.

Suppose for the sake of contradiction that p is the smallest vertex of a component and C is an induced odd cycle which is contained in a level. Since $(G_i - F_i, \prec)$ is triangle-free and ordered-hole-free, there are $ac, bd \in E(C)$ such that $a \prec b \prec c \prec d$. So $ad \in E(C)$, and none of the edges ac, bd, ad are crossing with any other edge of C . It follows that $V(C) \setminus \{a, d\}$ induces a path $bv_1 \dots v_t c$ with $b \prec v_1 \prec \dots \prec v_t \prec c$, for some positive integer t .

Let P be a shortest path from v_1 to p in $G_i - F_i$, and let v'_1 be the vertex adjacent to v_1 in P . If $a \preceq v'_1 \preceq d$ then we obtain a contradiction by finding a path which is shorter than P from either a or d to p , using the fact that $(G_i - F_i, \prec)$ is capped. Otherwise, if $v'_1 \prec a$, then $\{v'_1, v_1, \dots, v_t, c\}$ contains a triangle or an ordered hole, which is a contradiction. In the final case that $d \prec v'_1$, the vertices b, v_1, v'_1 form a triangle, which is again a contradiction. \triangleleft

Let ϕ_i be a 4-colouring of $G_i - F_i$ from the last claim, for $1 \leq i < \omega$. Let $F = \bigcup_{i=1}^{\omega-1} F_i$. If (G, \prec) is ordered-hole-free, then $F = \emptyset$ and the mapping $v \mapsto (\phi_1(v), \dots, \phi_{\omega-1}(v))$ is a $4^{\omega-1}$ -colouring of G . Otherwise, since every n -vertex outerplanar graph has at most $2n - 3$ edges, every n -vertex subgraph of $G[F]$ has at most $(2n - 3)(\omega - 1)$ edges, for any $n \geq 2$. So every non-empty subgraph of $G[F]$ has a vertex of degree less than $4(\omega - 1)$, and thus there exists a $4(\omega - 1)$ -colouring ψ of $G[F]$. Now, the mapping $v \mapsto (\phi_1(v), \dots, \phi_{\omega-1}(v), \psi(v))$ is a $4^\omega(\omega - 1)$ -colouring of G . \blacktriangleleft

References

- 1 James Abello, Ömer Egecioglu, and Krishna Kumar. Visibility graphs of staircase polygons and the weak Bruhat order, I: from visibility graphs to maximal chains. *Discrete & Computational Geometry*, 14(3):331–358, 1995.
- 2 James Abello and Krishna Kumar. Visibility graphs and oriented matroids. *Discrete & Computational Geometry*, 28(4):449–465, 2002.
- 3 Safwa Ameer, Matt Gibson-Lopez, Erik Krohn, Sean Soderman, and Qing Wang. Terrain visibility graphs: persistence is not enough. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 6:1–6:13. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, 2020.
- 4 Alan Arroyo, Julien Bensmail, and R. Bruce Richter. Extending drawings of graphs to arrangements of pseudolines. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 9:1–9:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, 2020.
- 5 Edgar Asplund and Branko Grünbaum. On a coloring problem. *Mathematica Scandinavica*, 8:181–188, 1960.
- 6 David Avis and David Rappaport. Computing the largest empty convex subset of a set of points. In *Proceedings of the First Annual Symposium on Computational Geometry*, pages 161–167. Association for Computing Machinery, New York, 1985.
- 7 Maria Axenovich, Jonathan Rollin, and Torsten Ueckerdt. Chromatic number of ordered graphs with forbidden ordered subgraphs. *Combinatorica*, 38(5):1021–1043, 2018.
- 8 Andreas Bärtschi, Subir K. Ghosh, Matúš Mihalák, Thomas Tschager, and Peter Widmayer. Improved bounds for the conflict-free chromatic art gallery problem. In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry*, pages 144–153. Association for Computing Machinery, New York, 2014.
- 9 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: max independent set and coloring. arXiv, 2020. [arXiv:2007.14161](https://arxiv.org/abs/2007.14161).
- 10 Onur Çağırıcı, Petr Hliněný, Subir K. Ghosh, and Bodhayan Roy. On conflict-free chromatic guarding of simple polygons. In Yingshu Li, Mihaela Cardei, and Yan Huang, editors, *Combinatorial Optimization and Applications*, volume 11949 of *Lecture Notes in Computer Science*, pages 601–612. Springer, Cham, 2019.
- 11 Onur Çağırıcı, Petr Hliněný, and Bodhayan Roy. On colourability of polygon visibility graphs. arXiv, 2019. [arXiv:1906.01904](https://arxiv.org/abs/1906.01904).
- 12 Jean Cardinal and Udo Hoffmann. Recognition and complexity of point visibility graphs. *Discrete & Computational Geometry*, 57(1):164–178, 2017.

- 13 Daniele Catanzaro, Steven Chaplick, Stefan Felsner, Bjarni V. Halldórsson, Magnús M. Halldórsson, Thomas Hixon, and Juraj Stacho. Max point-tolerance graphs. *Discrete Applied Mathematics*, 216(1):84–97, 2017.
- 14 Maria Chudnovsky, Alex Scott, and Paul Seymour. Induced subgraphs of graphs with large chromatic number. V. Chandeliers and strings. arXiv, 2016. [arXiv:1609.00314](https://arxiv.org/abs/1609.00314).
- 15 Alice M. Dean, William Evans, Ellen Gethner, Joshua D. Laison, Mohammad Ali Safari, and William T. Trotter. Bar k -visibility graphs: bounds on the number of edges, chromatic number, and thickness. In Patrick Healy and Nikola S. Nikolov, editors, *Graph Drawing*, volume 3843 of *Lecture Notes in Computer Science*, pages 73–82. Springer, Berlin, Heidelberg, 2006.
- 16 Stephan Eidenbenz and Christoph Stamm. Maximum clique and minimum clique partition in visibility graphs. In Jan van Leeuwen, Osamu Watanabe, Masami Hagiya, Peter D. Mosses, and Takayasu Ito, editors, *Theoretical Computer Science: Exploring New Frontiers of Theoretical Informatics*, pages 200–212. Springer, Berlin, Heidelberg, 2000.
- 17 Louis Esperet. *Graph Colorings, Flows and Perfect Matchings*. Habilitation thesis, Université Grenoble Alpes, 2017.
- 18 William S. Evans and Noushin Saeedi. On characterizing terrain visibility graphs. *Journal of Computational Geometry*, 6(1):108–141, 2015.
- 19 Subir K. Ghosh. On recognizing and characterizing visibility graphs of simple polygons. *Discrete & Computational Geometry*, 17(2):143–162, 1997.
- 20 Subir K. Ghosh and Partha P. Goswami. Unsolved problems in visibility graphs of points, segments, and polygons. *ACM Computing Surveys*, 46(2):Article 22, 2013.
- 21 Subir K. Ghosh, Thomas C. Shermer, Binay K. Bhattacharya, and Partha P. Goswami. Computing the maximum clique in the visibility graph of a simple polygon. *Journal of Discrete Algorithms*, 5(3):524–532, 2007.
- 22 András Gyárfás. On the chromatic number of multiple interval graphs and overlap graphs. *Discrete Mathematics*, 55(2):161–166, 1985.
- 23 Jan Kára, Attila Pór, and David R. Wood. On the chromatic number of the visibility graph of a set of points in the plane. *Discrete & Computational Geometry*, 34(3):497–506, 2005.
- 24 Friedrich Levi. Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade. *Berichte über die Verhandlungen der Königlich-Sächsischen Gesellschaft der Wissenschaften zu Leipzig, Mathematisch-Physische Klasse*, 78:256–267, 1926.
- 25 Fabrizio Luccio, Silvia Mazzone, and Chak Kuen Wong. A note on visibility graphs. *Discrete Mathematics*, 64(2):209–219, 1987.
- 26 Joseph O’Rourke and Ileana Streinu. Vertex-edge pseudo-visibility graphs: characterization and recognition. In *Proceedings of the Thirteenth Annual Symposium on Computational Geometry*, pages 119–128. Association for Computing Machinery, New York, 1997.
- 27 János Pach and István Tomon. On the chromatic number of disjointness graphs of curves. *Journal of Combinatorial Theory, Series B*, 144:167–190, 2020.
- 28 Arkadiusz Pawlik, Jakub Kozik, Tomasz Krawczyk, Michał Lasoń, Piotr Micek, William T. Trotter, and Bartosz Walczak. Triangle-free intersection graphs of line segments with large chromatic number. *Journal of Combinatorial Theory, Series B*, 105:6–10, 2014.
- 29 Florian Pfender. Visibility graphs of point sets in the plane. *Discrete & Computational Geometry*, 39(1–3):455–459, 2008.
- 30 Alexandre Rok and Bartosz Walczak. Coloring curves that cross a fixed curve. *Discrete & Computational Geometry*, 61(4):830–851, 2019.
- 31 Alexandre Rok and Bartosz Walczak. Outerstring graphs are χ -bounded. *SIAM Journal on Discrete Mathematics*, 33(4):2181–2199, 2019.
- 32 Alex Scott and Paul Seymour. Induced subgraphs of graphs with large chromatic number. VI. Banana trees. *Journal of Combinatorial Theory, Series B*, 145:487–510, 2020.
- 33 Ileana Streinu. Non-stretchable pseudo-visibility graphs. *Computational Geometry*, 31(3):195–206, 2005.
- 34 Subhash Suri. A linear time algorithm for minimum link paths inside a simple polygon. *Computer Vision, Graphics, and Image Processing*, 35(1):99–110, 1986.
- 35 István Tomon. personal communication.