Journal of Combinatorial Theory, Series B 161 (2023) 268–300

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series B

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

Coloring polygon visibility graphs and their generalizations ☆

James Davies [a], Tomasz Krawczyk [b,1], Rose McCarty [c], Bartosz Walczak [b,1]

[a] Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Cambridge, UK
[b] Department of Theoretical Computer Science, Faculty of Mathematics and Computer Science, Jagiellonian University, Kraków, Poland
[c] Department of Mathematics, Princeton University, Princeton, NJ, USA

A R T I C L E I N F O

Article history:
Received 12 August 2021
Available online xxxx

Keywords:
Visibility graphs
χ-Boundedness
Pseudoline arrangements
Ordered graphs

A B S T R A C T

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 coloring with the claimed number of colors can be computed in polynomial time.

☆ A preliminary version of this article appeared in: Kevin Buchin, Éric Colin de Verdière (Eds.), 37th International Symposium on Computational Geometry (SoCG 2021), in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 189, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, 2021, 29.

E-mail address: bartosz.walczak@uj.edu.pl (B. Walczak).

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

https://doi.org/10.1016/j.jctb.2023.02.010

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 Fig. 1, top-left). The starting point of and main motivation for this work is the question of Kára, Pór, and Wood [25] 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 $3 \cdot 4^{\omega-1}$.*

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 Fig. 1, bottom-left).

O'Rourke and Streinu [28] studied visibility graphs of pseudo-polygons (polygons on pseudoline arrangements; see Fig. 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 Fig. 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.

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

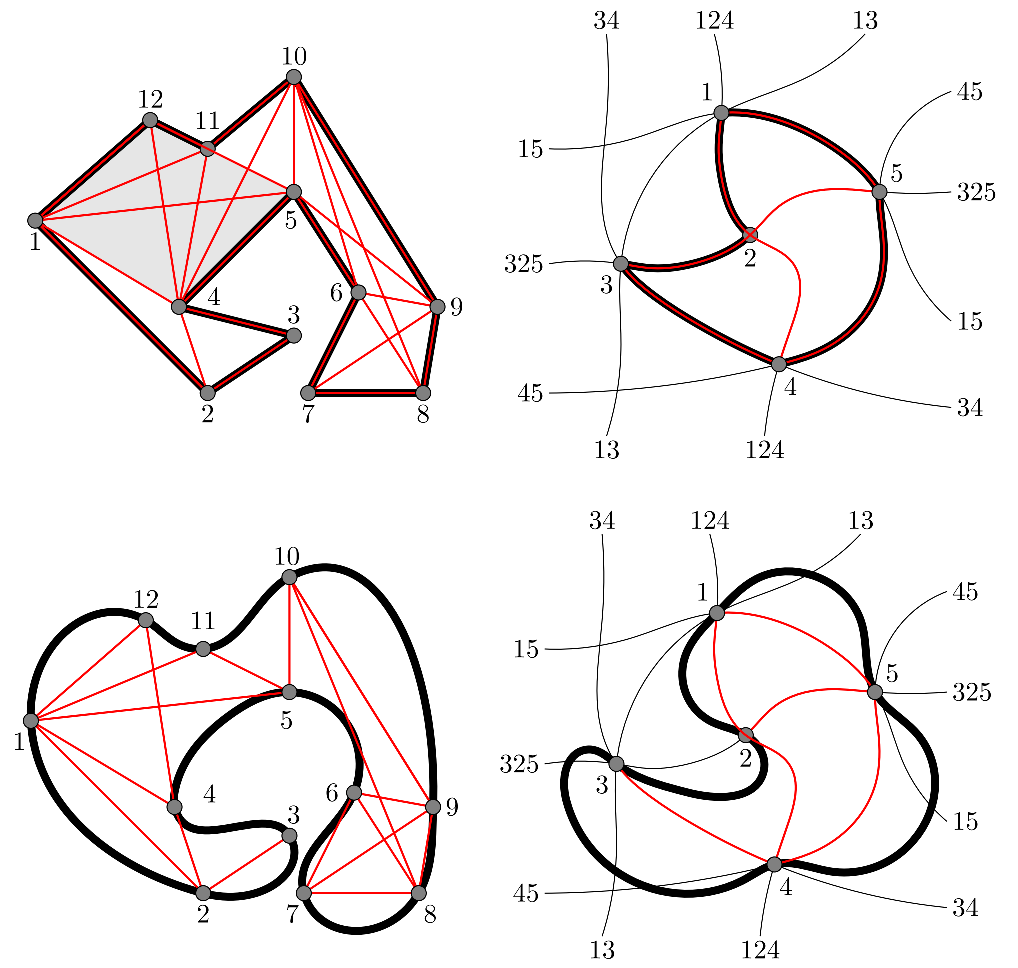

Fig. 1. From left to right: a polygon visibility graph (where the convex hull of a maximum clique is shaded), a pseudo-polygon visibility graph, a curve visibility graph, and a curve pseudo-visibility graph. A "visibility" between each pair of adjacent vertices is drawn with a red (pseudo-)segment. (For interpretation of the colors in the figures, the reader is referred to the web version of this article.)

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. In Section 3, we provide 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 Fig. 2). Excluding these obstructions is equivalent to some conditions previously studied in the context of polygon visibility graphs [2,21] (see Section 3) and is easily verifiable in polynomial time. We prove the following further generalization of Theorem 1.2.

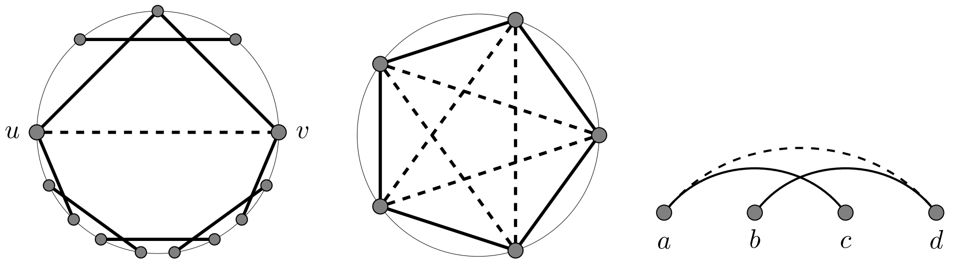

Fig. 2. A graph in $\mathcal{H}$ (left), an ordered hole (middle), and the forbidden configuration for a capped graph (right). Dashed lines indicate non-edges. The pairs of vertices where no lines are drawn can be edges or non-edges.

Theorem 1.3. *Every $\mathcal{H}$-free ordered graph with clique number $\omega \geqslant 2$ 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 coloring with the claimed number of colors.*

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 Fig. 2 (right). This condition has been previously studied for terrain visibility graphs [1,3], where it is called the "X-property". In Section 4, 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 $\omega \geqslant 2$ 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 coloring with the claimed number of colors.*

We prove Theorem 1.4 in Section 5. 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 [19] conjectured that every χ-bounded class of graphs that is hereditary, that is, closed under taking induced subgraphs, is polynomially χ-bounded. This conjecture in general has been disproved [11], but we expect that it holds for capped graphs (and, consequently, for the graphs considered in Theorems 1.1–1.3).

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

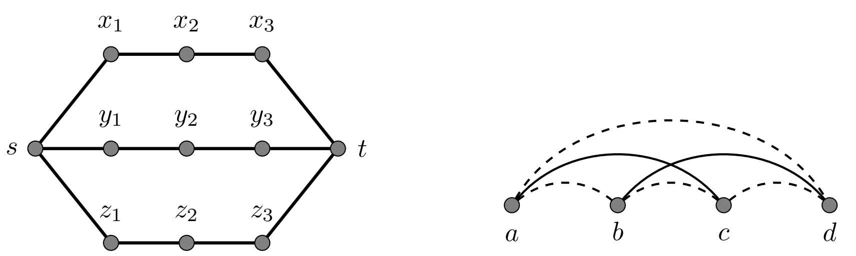

Fig. 3. The banana B_4 (left) and the ordered graph X (right).

While our proof of Theorem 1.4 is direct, we remark that a recent result of Scott and Seymour [34] 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 Fig. 3 (right) as an induced ordered subgraph. In particular, every capped graph is X-free. Tomon [37] 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 [33] 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 [34] 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. Fig. 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. We present more details in Section 6.

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

We conclude the introduction with a brief literature review in order to place Theorems 1.1–1.4 and 1.6 in context.

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 [6] and circle graphs [24] 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 [30], some very general subclasses are [16,32]. There are also very precise results for disjointness graphs of certain kinds of curves in the plane [29].

Less is known about χ-boundedness of visibility graphs, even though various kinds of such graphs have been considered in the literature—see [22] for a survey. Kára, Pór, and Wood [25] conjectured that the class of point visibility graphs is χ-bounded, but this was disproved by Pfender [31]. Some types of bar visibility graphs are related to interval graphs [17] and planar graphs [27] and are therefore known to be χ-bounded.

Axenovich, Rollin, and Ueckerdt [8] 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 [29] 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 [15] and classes of graphs of bounded twin-width [10] are also known to be χ-bounded and have well-understood characterizations as ordered graphs.

Characterizations and algorithms 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-polygon 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 [14,21].

These difficult characterization problems tend to become tractable, and have more natural solutions, in the "pseudo-visibility setting" [1,2,20,28]. This is due to the connection between stretchability of pseudoline arrangements and representability of rank 3 oriented matroids. The pseudo-visibility setting is more combinatorial, because it suffices to find the associated rank 3 oriented matroid without worrying about its representability. For visibility graphs of pseudo-polygons, this approach recently resulted in a polynomial-time recognition algorithm under the assumption that the order of vertices along the pseudo-polygon is given in the input [4]. However, pseudo-polygon visibility graphs are strictly more general than polygon visibility graphs [35].

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 insufficient—see Remark 3.10. The above-mentioned characterizations of pseudo-polygon visibility seem rather unhelpful, as they rely on existence of "blocking vertices" witnessing mutual invisibilities, and such vertices do not necessarily exist in the setting of curve pseudo-visibility. 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 [7,18,23], although our algorithm is certainly slower. We cannot expect to get an exact algorithm for the chromatic number, as Çağırıcı, Hliněný, and Roy [13] proved that it is NP-complete to decide if a polygon visibility graph is 5-colorable, 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 [28] 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 (each very closely on one side of some pseudoline connecting the two components) 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 $|V^*| \geqslant 3$; we first select V^* and then extend $\mathcal{L}$ to $\mathcal{L}^*$ using Levi's extension lemma [26]. 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 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 [26]. 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 $|V^*| \geqslant 3$) 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 $G_{\mathcal{L}}(K', V) = G_{\mathcal{L}}(K, V)$. $\square$

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

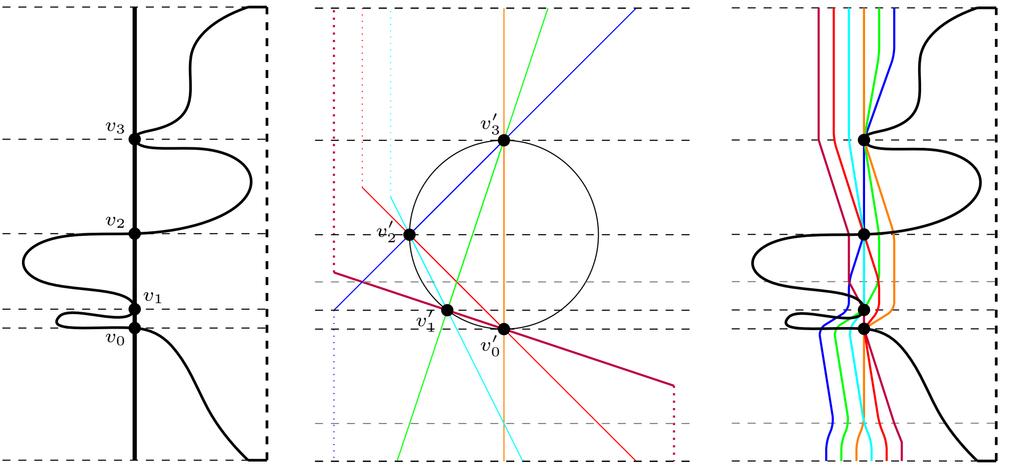

Fig. 4. Replacing L with a bundle of pseudo-lines $\mathcal{B}_L$ in the proof of Proposition 2.2.

small neighborhood 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 neighborhood 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 $v_0, \ldots, v_k$ 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 $0 \leqslant i < j \leqslant k$, 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 $\{L_{i,j}\}_{0 \leqslant i < j \leqslant k}$ horizontally to fit it in a small neighborhood of L and performing local horizontal shifts to guarantee conditions (1) and (2); condition (3) then follows. See Fig. 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 any finite set of points V on a Jordan curve K as being ordered counterclockwise around K, as in Fig. 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* (but not the other way round) and that ac and bd are *crossing edges*. The property that a pair of edges is crossing is preserved under rotation. 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 an ordered 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 not 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 neighborhood 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 [21] observed that these are obstructions for polygon visibility graphs; specifically, ordered-hole-freeness and $\mathcal{H}$-freeness are equivalent to Ghosh's necessary conditions 1 and 2. Ordered graphs with an ordered Hamilton cycle that are ordered-hole-free and $\mathcal{H}$-free (satisfy Ghosh's conditions 1 and 2) have been studied as "quasi-persistent graphs" [2].

If two vertices u and v are non-adjacent in a curve pseudo-visibility graph $G_{\mathcal{L}}(K, V)$, then it is because at least one of two parts of K between u and v "blocks" visibility between them. The intuition behind the next definition is that if there is a "crossing sequence" from u to v, then the part of K from u to v (counterclockwise) cannot "block" visibility between u and v.

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 $e_1, \ldots, e_k$ 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 $1 \leqslant i < k$. Observe that the notion of a crossing sequence is invariant under rotation of $\prec$ as long

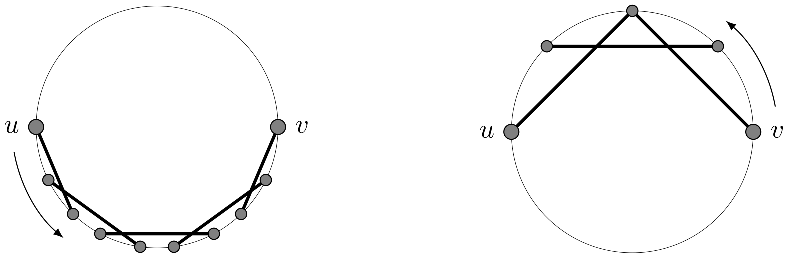

Fig. 5. A crossing sequence from u to v (left) and from v to u (right).

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 Fig. 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.*

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. $\square$

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 Fig. 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 $k \geqslant 4$ 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 Fig. 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 Propositions 3.2 and 3.3 later in this section. Before that, 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 proofs of Propositions 3.2 and 3.3 require 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 Fig. 6. The following lemma was proved by Arroyo, Bensmail, and Richter [5]; 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.*

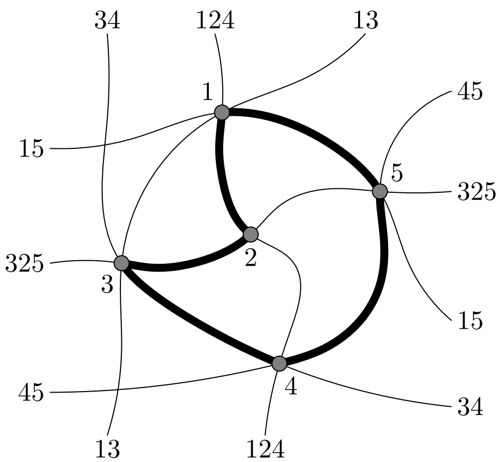

Fig. 6. A pseudo-polygon with five articulation points: $1, 3, 4, 5$ are convex, 2 is concave.

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 showed 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. $\square$

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 $e_1, \ldots, e_k$ from u to v and $f_1, \ldots, f_t$ 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 $1 \leqslant i < j \leqslant k$, the segments ℓ_{e_i} and ℓ_{e_j} intersect if and only if $j = i + 1$, and likewise for the crossing sequence $f_1, \ldots, f_t$.

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 $\mathcal{L}$ disjoint from the exterior of K 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' and thus of K, so $uv \in E(G)$. $\square$

Proof of Proposition 3.3. When K is a Jordan curve and $p, q \in K$, we write $K[p, q]$ and $K(p, q)$ for the closed and the open segment (respectively) of K from p to q in the counterclockwise direction. When L is a pseudoline and $p, q \in L$, we write $L[p, q]$ for the closed segment of L connecting p and q.

Let $G = G_{\mathcal{L}}(K, V)$ be a curve pseudo-visibility graph. By Proposition 2.2, we can assume without loss of generality that $(\mathcal{L}, V)$ is in general position. Suppose for the sake of contradiction that G has an induced cycle of length at least four on vertices $p_1, \ldots, p_k$ with $p_1 \prec \cdots \prec p_k$. We write indices cyclically, so that $p_{k+1} = p_1$, $p_{k+2} = p_2$, and so on. Let K' be the pseudo-polygon on $\mathcal{L}$ formed by the pseudoline segments $p_i p_{i+1}$ with $1 \leqslant i \leqslant k$. It is disjoint from the exterior of K. It follows from Lemma 2.3 that K' is a pseudo-polygon with articulation points $p_1, \ldots, p_k$ in cyclic order (and no other articulation points). By Lemma 3.6, K' has a convex articulation point. Up to rotation, we can assume that p_2 is convex.

By Lemma 3.7, since $p_1 p_3$ is not an edge of G and p_2 is convex, there is a concave articulation point of K' among $p_4, \ldots, p_k$. We now choose a new pseudo-polygon K^* with similar properties, but which is also "minimal". We will then use the existence of a concave articulation point of K^* to reach a contradiction to minimality. For $1 \leqslant i \leqslant k$, let L_i be the pseudoline in $\mathcal{L}$ passing through p_i and p_{i+1}. Choose a pseudo-polygon K^* on $\{L_i\}_{1 \leqslant i \leqslant k}$ so that

(i) $K'[p_1, p_3] \subseteq K^*$ and K^* is disjoint from the exterior of K,
(ii) every concave articulation point of K^* belongs to $\{p_1, \ldots, p_k\}$, and
(iii) subject to conditions (i) and (ii), the region of the plane bounded by K^* is minimal.

Such a pseudo-polygon K^* exists, as K' is a candidate.

With this choice, we still have that $K^*(p_3, p_1)$ contains a concave articulation point of K^*, as otherwise, by Lemma 3.7, the pseudoline segment $p_1 p_3$ in $\mathcal{L}$ would be disjoint from the exterior of K^* and thus from the exterior of K, contradicting the assumption that $p_1 p_3$ is not an edge of G. By (ii), there is an index i with $4 \leqslant i \leqslant k$ such that p_i is a concave articulation point of K^*. Moreover, we have the following.

Claim 3.8. *Neither p_1 nor p_3 is a concave articulation point of K^*.*

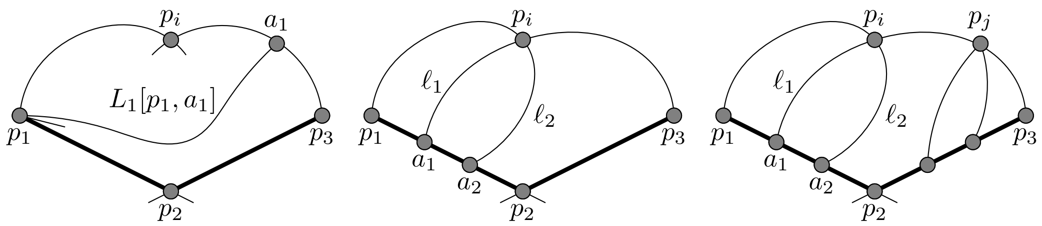

Fig. 7. Illustrations for the proof of Proposition 3.3, with $K^*[p_1, p_3]$ in bold.

Proof. Suppose that p_1 is a concave articulation point of K^*. Walk along the pseudoline L_1 from p_1 in the direction opposite to p_2 (towards the interior of K^*) until hitting K^* at a point $a_1 \in L_1 \cap K^*$. Then the pseudo-polygon $K^*[p_1, a_1] \cup L_1[p_1, a_1]$ contradicts the choice of K^*; see Fig. 7 (left). An analogous contradiction is reached when p_3 is concave in K^*. $\square$

Recall that p_i is a concave articulation point of K^* with $4 \leqslant i \leqslant k$. By (i) and (ii), there are open segments ℓ_1 of L_{i-1} and ℓ_2 of L_i that are both contained in the interior of K^*, have p_i as one endpoint and have the other endpoint on K^*; see Fig. 7 (middle). Let a_1 and a_2 be the other endpoints of ℓ_1 and ℓ_2, respectively. We now show that, up to symmetry, the case depicted in Fig. 7 (middle) actually occurs.

Claim 3.9. *Either $a_1, a_2 \in K^*(p_1, p_2)$ or $a_1, a_2 \in K^*(p_2, p_3)$.*

Proof. We have $a_1 \notin K^*[p_i, p_1]$, otherwise the pseudo-polygon $K^*[a_1, p_i] \cup \ell_1$ would contradict the choice of K^*. Furthermore, $a_1 \neq p_2$, because $a_1 \in L_{i-1}$ and L_{i-1} does not pass through p_2 by the general position assumption. By a symmetric argument for a_2, we obtain that $a_1, a_2 \in K^*(p_1, p_3) \setminus \{p_2\}$. Finally, it is not possible that $a_1 \in K^*(p_1, p_2)$ and $a_2 \in K^*(p_2, p_3)$, because then $K^*[a_1, a_2] \cup \ell_1 \cup \ell_2$ would be a pseudo-polygon with no concave articulation points, so the pseudoline segment p_2p_i in $\mathcal{L}$ would be disjoint from the exterior of it (and thus from the exterior of K') by Lemma 3.7, contradicting the assumption that p_2p_i is not an edge of G. $\square$

So, up to symmetry, we can assume that $a_1, a_2 \in K^*(p_1, p_2)$. After possibly changing p_i, we can assume that there is no concave articulation point of K^* on $K^*(p_3, p_i)$ with the same property. By Lemma 3.7, since p_2p_i is not an edge of G, the pseudo-polygon $K^*[a_2, p_i] \cup \ell_2$ has a concave articulation point p_j with $3 \leqslant j < i$. By Claim 3.8, p_3 is not a concave articulation point of K^* so $4 \leqslant j < i$. We can assume that j is maximal with that property, so that all articulation points of K^* on $K^*(p_j, p_i)$ are convex. Repeating the argument from the last claim and by the choice of p_i, there are two segments of L_{j-1} and L_j which have p_j as one endpoint and have the other endpoint on $K'(p_2, p_3)$; see Fig. 7 (right). But then, there is a pseudo-polygon on $\mathcal{L}$ with no concave articulation points that contains p_2 and p_i (and p_j), and therefore, by Lemma 3.7, p_2p_i is an edge of G. This contradiction completes the proof. $\square$

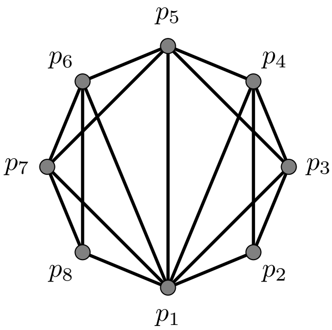

Fig. 8. An ordered graph that is $\mathcal{H}$-free and ordered-hole-free but not a curve pseudo-visibility graph.

Remark 3.10. The ordered graph $(G, \prec)$ depicted in Fig. 8, discovered by Ghosh [21], is $\mathcal{H}$-free and ordered-hole-free but not a curve pseudo-visibility graph. To see the latter, suppose that $G = G_{\mathcal{L}}(K, V)$, where $(\mathcal{L}, V)$ is in general position. Lemma 3.6 applied to the pseudo-triangles $p_2p_3p_4$ and $p_3p_4p_5$ implies that p_3 and p_4 are convex articulation points and therefore, by Lemma 3.7, p_1 is a concave articulation point of the pseudo-pentagon $p_1p_2p_3p_4p_5$. Analogously, p_1 is concave in the pseudo-pentagon $p_5p_6p_7p_8p_1$. However, p_1 cannot be concave in both pseudo-pentagons simultaneously. It can be easily checked that $(G, \prec)$ is $\mathcal{H}$-free and ordered-hole-free.

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 $\prec$. However, capped graphs naturally appear within $\mathcal{H}$-free ordered graphs in the following way.

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

Proof. If $a, b, c, d \in X$, $a \prec b \prec c \prec d$, 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 $(G, \prec)$, which implies $ad \in E(G)$, as $(G, \prec)$ is $\mathcal{H}$-free. $\square$

The clique number of an $\mathcal{H}$-free ordered graph $(G, \prec)$ is the maximum of the clique numbers of the capped graphs $(G[X], \prec|_X)$ defined in Lemma 4.1 taken over all choices of adjacent vertices $u \prec v$. Thus, computing the clique number of $\mathcal{H}$-free ordered graphs is reduced to computing the clique number of capped graphs, and 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.*

Before delving into the full proof of Proposition 4.2, we sketch the proof for the case that $(G, \prec)$ is an ordered polygon visibility graph. The sketch presents some key intuitions behind the full proof. It is based on the "window partition" by Suri [36], which was used in a similar way to approximate chromatic variants of the well-known art gallery problem [9,12].

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 $P_{xy} = (P \cup \operatorname{int} P) \setminus xy$, where $\operatorname{int} P$ is the interior of P. We construct a partition of $V(G)$ into three sets, which we express in terms of a coloring ϕ of $V(G)$ that uses three colors: 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 Fig. 9 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, 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 Fig. 9.

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

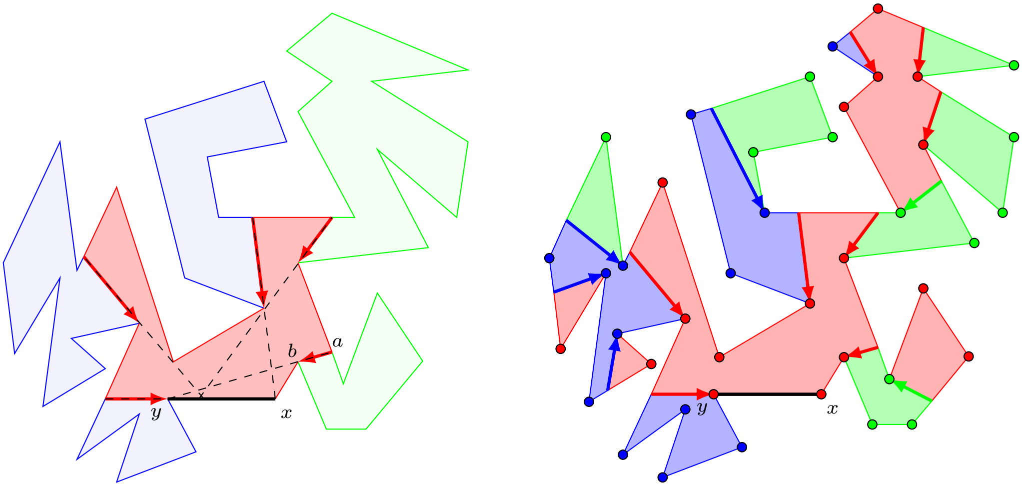

Fig. 9. To the left: a polygon P with window W_{xy} based at xy in red, oriented lines $\overrightarrow{L_{ab}}$ depicted with red arrows and dashed lines, and the left/right-invisible sets I_{ab} in green/blue, respectively. To the right: the final window partition of P_{ab}.

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-coloring ϕ of the vertex set of $(G, \prec)$ such that each color class induces a capped subgraph. The property above allows us to color the windows by three colors (say, red, green, and blue) so that no two points in two different windows of the same color are mutually visible. We color the root window, say, by red. Then we extend this coloring on the remaining windows so that the children of each window W obtain a color different from W and the left children of W are colored with a different color from the right children of W. This way every vertex of P other than x and y is colored. We color x and y arbitrarily; see Fig. 9 (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 [20, 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 $a', d' \in xy$ be such that the segments $a'a$ and $d'd$ are disjoint from the exterior of P; see Fig. 10 for an illustration. 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). To show that ad is disjoint from the exterior of P, it suffices to prove the following two claims.

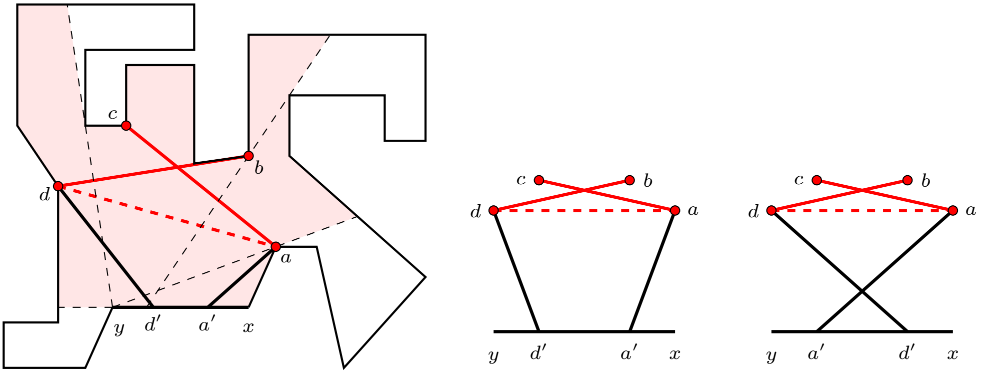

Fig. 10. Possible relations between the segments $a'a, d'd, ac, bd, yx$.

- The polygon P is contained in the closure of the outer face of $\mathcal{F}$ and P does not cross (but may touch) any of the four segments $a'a, d'd, ac, bd$.
- 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 Fig. 10 for an illustration. The second claim can be proved 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. $\square$

Now we proceed to the proof of Proposition 4.2 in the general case. It is convenient to word it in terms of colorings; we will find a 3-coloring such that each color class induces a capped subgraph. The proof will work by extending certain partial colorings "inside a valid segment". To explain this, we need to introduce some notation.

Let $(G, \prec)$ be an ordered graph. For any two vertices x and y of G with $x \prec y$, we define

$$V[x,y] = \{v \in V(G) \colon x \preceq v \preceq y\}, \qquad V[x,y) = \{v \in V(G) \colon x \preceq v \prec y\},$$
$$V(x,y] = \{v \in V(G) \colon x \prec v \preceq y\}, \qquad V(x,y) = \{v \in V(G) \colon x \prec v \prec y\}.$$

A *segment* of $(G, \prec)$ is a set of the form $V[x,y]$, that is, a set of at least two consecutive vertices. In the context of a fixed segment $V[x,y]$, we call the vertices in $V(x,y)$ *internal*, the vertices x and y the *ends*, and the vertices in $V(G) \setminus V[x,y]$ *external*. A segment $V[x,y]$ is *valid* if there is a crossing sequence from y to x (see Fig. 11).

Recall the proof sketch for partitioning the vertex set of a polygon visibility graph into three capped graphs. We had a partial coloring ϕ and a segment $V[x,y]$ such that x and y were colored and connected by an edge. The colored internal vertices were exactly those vertices which were "visible" to the segment xy, and these vertices were all given the same color. We then partitioned the internal vertices into segments which were "left-invisible" and segments which were "right-invisible", and we continued the process.

Now, instead of requiring that x and y are connected by an edge, we just require that the segment $V[x,y]$ is valid. It turns out that even for visibility graphs of points on a

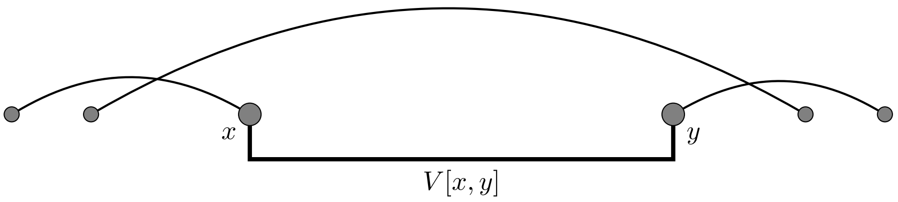

Fig. 11. A valid segment $V[x,y]$.

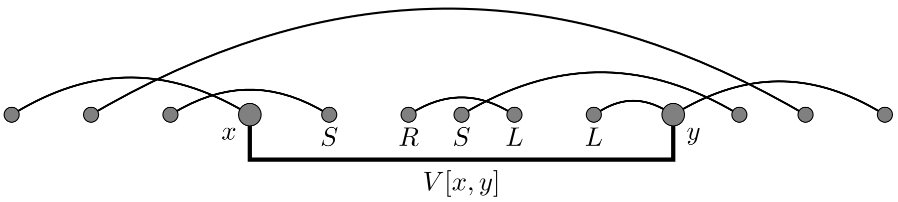

Fig. 12. Internal vertices which are strongly reachable (labelled S) or left- or right-reachable only (labelled L and R respectively).

Jordan curve, it is not always possible to tell from $(G, \prec)$ alone if an "invisible" vertex is "left-invisible" or "right-invisible"; the vertex could, for instance, even be isolated. So instead we will define left-reachable vertices which definitely cannot be "left-invisible" and right-reachable vertices which definitely cannot be "right-invisible". Vertices which are both left- and right-reachable will be called strongly reachable; these are the vertices that are "visible". The proof ends up being rather technical, but these are the key intuitions. For the following, refer to Fig. 12.

Let $V[x,y]$ be a valid segment. A vertex v is *left-reachable* from $V[x,y]$ if $v \in V[x,y)$ and there is a crossing sequence from y to v (so in particular x is left-reachable). A vertex v is *right-reachable* if $v \in V(x,y]$ and there is a crossing sequence from v to x (so y is right-reachable). A vertex is *strongly reachable* if it is both left- and right-reachable (so all strongly reachable vertices are internal). We write $L[x,y)$, $R(x,y]$, and $S(x,y)$ for the set of left-, right-, and strongly reachable vertices, respectively. So $S(x,y) = L[x,y) \cap R(x,y]$. First, we observe the following.

Lemma 4.3. *For any ordered graph $(G, \prec)$ and valid segment $V[x,y]$, every internal vertex which is adjacent to an external vertex is strongly reachable.*

Proof. Let v be an internal vertex which is adjacent to an external vertex u. There are crossing sequences from y to x and from u to v. Thus, by Lemma 3.1, there is a crossing sequence from y to v, and so v is left-reachable. Likewise, since there are crossing sequences from v to u and from y to x, the vertex v is also right-reachable, so it is strongly reachable. □

As further motivation for these definitions, we prove the following.

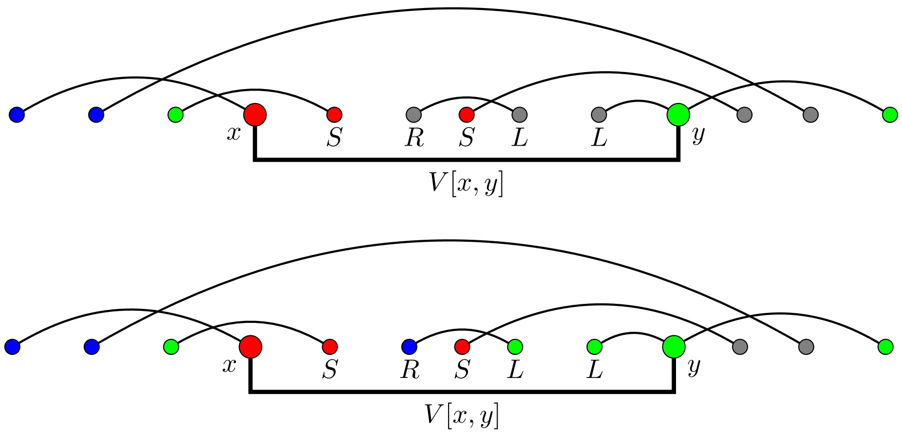

Fig. 13. A $V[x, y]$-precoloring ϕ (top) and a $V[x, y]$-extension of ϕ (bottom).

Lemma 4.4. *For any $\mathcal{H}$-free ordered graph $(G, \prec)$ and valid segment $V[x, y]$, the set $S(x, y) \cup \{x, y\}$ induces a capped subgraph.*

Proof. Otherwise, there are vertices $a, b, c, d \in S(x, y) \cup \{x, y\}$ such that $a \prec b \prec c \prec d$, and $ac, bd \in E(G)$, yet $ad \notin E(G)$. Then d is right-reachable, and thus there is a crossing sequence from d to x. Likewise, a is left-reachable, and there is a crossing sequence from y to a. Thus there is a crossing sequence from d to a—this is trivial if $d = y$ or $a = x$, and otherwise we apply Lemma 3.1. This contradicts the assumption that $(G, \prec)$ is $\mathcal{H}$-free. □

Now we need to define certain partial colorings and what it means to extend them. For the following definition, refer to Fig. 13.

For an ordered graph $(G, \prec)$ and a valid segment $V[x, y]$, a $V[x, y]$-*precoloring* is a partial 3-coloring ϕ such that

(i) every color class of ϕ induces a capped subgraph,
(ii) ϕ colors $S(x, y) \cup \{x, y\}$ and colors no other internal vertices of $V[x, y]$, and
(iii) ϕ colors every vertex in $S(x, y)$ the same color, and no external vertex of $V[x, y]$ which is a neighbor of an internal vertex of $V[x, y]$ is given this color.

A $V[x, y]$-*extension* of ϕ is obtained by taking ϕ and coloring the remaining internal vertices of $V[x, y]$ (which are not assigned any color in ϕ) while maintaining the condition that each color class induces a capped subgraph. By the next lemma, that condition is maintained as long as the set $X \cap V[x, y]$ induces a capped subgraph for each color class X.

Lemma 4.5. *If $(G, \prec)$ is an ordered graph with a segment $V[x, y]$ such that no internal vertex is adjacent to an external vertex and each of the sets $V[x, y]$ and $V(G) \setminus V(x, y)$ induces a capped subgraph, then $(G, \prec)$ is capped.*

Proof. Suppose for the sake of contradiction that there are vertices a, b, c, d such that $a \prec b \prec c \prec d$ and $ac, bd \in E(G)$ while $ad \notin E(G)$. Then at least one of a, b, c, d must be external to $V[x, y]$ and at least one must be internal. If b is external, then either a or d is also external. So, up to symmetry, we can assume that a is external. Then $a \prec x$ and c is not internal. It follows that exactly one of b and d is internal and the other is external—a contradiction. □

We will use the following lemma in order to find valid segments inside a valid segment $V[x, y]$. Conditions (ii) and (iii) are symmetric. As an example, observe that in Fig. 13, condition (i) implies that the second and third vertices in $V[x, y]$ form a valid segment, and condition (ii) implies that the first and second vertices in $V[x, y]$ form a valid segment.

Lemma 4.6. *For any ordered graph $(G, \prec)$ and valid segment $V[x, y]$, each segment $V[a, b]$ which satisfies at least one of the following conditions is valid:*

(i) *$a \in L[x, y)$ and $b \in R(x, y]$;*
(ii) *$a, b \in L[x, y)$, no vertex in $V(a, b)$ is left-reachable from $V[x, y]$, and there is a vertex in $V(b, y)$ which is right-reachable from $V[x, y]$;*
(iii) *$a, b \in R(x, y]$, no vertex in $V(a, b)$ is right-reachable from $V[x, y]$, and there is a vertex in $V(x, a)$ which is left-reachable from $V[x, y]$.*

Proof. First suppose that condition (i) holds. Then there are crossing sequences from b to x and from y to a and thus, using Lemma 3.1 if $b \neq y$ and $a \neq x$, also from b to a. So $V[a, b]$ is valid.

Conditions (ii) and (iii) are symmetric via reversing $\prec$. So it suffices to consider the case that condition (ii) holds. Then there are crossing sequences from y to a and from y to b. If the crossing sequence from y to b is a single edge, then, since there is a vertex in $V(b, y)$ which is right-reachable and by Lemma 3.1, there is a crossing sequence from b to x. Then b is right-reachable and $V[a, b]$ is valid by condition (i).

So there is a crossing sequence, say $e_1, \dots, e_k$, from y to b with at least two edges. Let v be the end of e_{k-1} such that $e_1, \dots, e_{k-1}$ is a crossing sequence from y to v (so if $k = 2$, then v is the end of $e_{k-1} = e_1$ which is not y). Then $v \notin V(a, b)$ since then v would be left-reachable. This implies that the end u of e_k which is not b satisfies $u \notin V[a, y]$. Thus, since e_k is a crossing sequence from b to u and there is a crossing sequence from y to a, by Lemma 3.1 there is a crossing sequence from b to a. So $V[a, b]$ is valid. □

We are ready to begin the proof of the main proposition. Afterwards, we will quickly show how to apply it when there is no fixed precoloring.

Proposition 4.7. *There is a polynomial-time algorithm which takes in an $\mathcal{H}$-free ordered graph $(G, \prec)$, vertices x and y such that $V[x, y]$ is a valid segment, and a $V[x, y]$-precoloring ϕ, and returns a $V[x, y]$-extension of ϕ.*

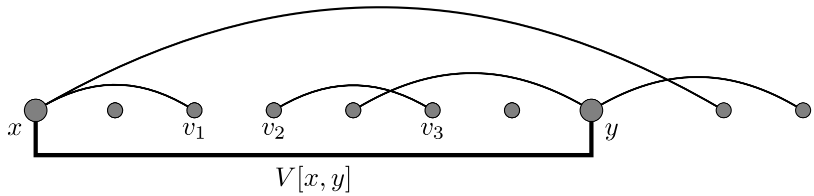

Fig. 14. The vertices $x = v_0 \prec v_1 \prec v_2 \prec v_3 \prec v_4 = y$.

Proof. Throughout the proof we will implicitly use the fact that there is a polynomial-time algorithm to determine whether there is a crossing sequence from a vertex u to a vertex v; see the proof of Proposition 3.4. So in particular we can find the sets $L[x, y)$, $R(x, y]$, and $S(x, y)$ in polynomial time. The algorithm will be called recursively on valid segments $V[a, b] \subset V[x, y]$. To ensure that the overall running time of the algorithm is polynomial, we create a memoization table with an entry for each valid segment $V[a, b]$ of $(G, \prec)$ and each triple of colors c_1, c_2, c_3 such that the coloring that assigns color c_1 to a, color c_2 to b, and color c_3 to all vertices in $S(a, b)$ is a $V[a, b]$-precoloring. Once a $V[a, b]$-extension of any $V[a, b]$-precoloring which assigns these colors to a, b, and $S(a, b)$ (respectively) is computed, its restriction to $V[a, b]$ is stored in the table. Then, whenever a $V[a, b]$-extension of another such $V[a, b]$-precoloring is requested, the colors of all vertices in $V[a, b]$ are copied from the table entry. By Lemma 4.5, this gives a correct $V[a, b]$-extension of the other $V[a, b]$-precoloring.

Now, we describe a single recursion step of the coloring algorithm. If all internal vertices of $V[x, y]$ are colored, then we just return ϕ. So suppose that some internal vertex of $V[x, y]$ is not colored yet. We distinguish two cases depending on whether $S(x, y)$ is empty.

Case 1. The set $S(x, y)$ is empty.

We begin with an informal description of the approach. We will partition $V[x, y]$ into smaller valid segments by greedily constructing something like a crossing sequence from x to y. However, since xy could be an edge, we will have to artificially enforce that these new segments are actually smaller than $V[x, y]$. Moreover, this greedy process could get stuck before reaching y, in which case we will just move over to the next vertex.

Now we describe this procedure in detail. The *successor* of a vertex $a \in V[x, y)$ is the next vertex after a in the order $\prec$. We construct a sequence of vertices $x = v_0 \prec v_1 \prec \cdots \prec v_k \prec v_{k+1} = y$ with $k \geqslant 1$, as follows; see Fig. 14 for an illustration. Set $v_0 := x$. Let v_1 be the largest neighbor of x in $V(x, y)$ or the successor of x if there is no such neighbor. Then, after having defined vertices $x = v_0 \prec v_1 \prec \cdots \prec v_i \prec y$, let v_{i+1} be the largest vertex in $V(v_i, y]$ that has a neighbor in $V(x, v_i]$ or the successor of v_i if no such vertex exists. If $v_{i+1} = y$, then set $k := i$ and stop the process, otherwise continue with the next value of i.

First, we show that $V[v_i, v_{i+1}]$ is a valid segment for each $i \in \{0, \ldots, k\}$ with $V(v_i, v_{i+1})$ non-empty. Going for a contradiction, suppose otherwise. Since $V[v_i, v_{i+1}]$ is not valid, the vertices v_i and v_{i+1} are non-adjacent. Then, the assumption that

$V(v_i, v_{i+1})$ is non-empty and the choice of v_i imply that $i \geqslant 1$ and v_{i+1} has a neighbor $u \in V(v_{i-1}, v_i)$. Now, since $V(v_{i-1}, v_i)$ is non-empty, the vertex v_i has a neighbor $u' \in V[x, v_{i-1}]$. It follows that $v_{i+1}u, u'v_i$ is a crossing sequence from v_{i+1} to v_i, which is a contradiction.

We define a coloring ϕ' obtained from ϕ by coloring some additional vertices, as follows. Color the vertices $v_1, \dots, v_k$ with the color $\phi(y)$. For each $i \in \{0, \dots, k\}$, color the vertices in $S(v_i, v_{i+1})$ using one of the other two colors, alternating the color according to the parity of i in such a way that the color $\phi(x)$ is used for $S(v_0, v_1)$ if $\phi(x) \neq \phi(y)$.

We claim that each color class X of ϕ' induces a capped subgraph. By Lemma 4.5, it suffices to show that $X \cap V[x, y]$ induces a capped subgraph. For the color class of $\phi(y)$, this follows from the fact that the only possible adjacencies among $v_0, \dots, v_{k+1}$ are $v_0v_1, v_1v_2, \dots, v_kv_{k+1}$, and v_0v_{k+1}. So suppose that X is one of the other two color classes. Observe that no vertex in $V(v_j, v_{j+1})$ is adjacent to any vertex in $V(v_i, v_{i+1})$ if $j > i + 1$. Likewise, no vertex in $V(v_j, v_{j+1})$ is adjacent to x if $j \geqslant 1$. Therefore, any two adjacent vertices in $X \cap V[x, y]$ belong to the same set $V[v_i, v_{i+1}]$. Moreover, by Lemma 4.4, the vertices in $\{v_i, v_{i+1}\} \cup S(v_i, v_{i+1})$ induce a capped subgraph for each $i \in \{0, \dots, k\}$. It follows that $X \cap V[x, y]$ indeed induces a capped subgraph.

To construct the requested $V[x, y]$-extension of ϕ, we apply the following procedure. Set $\phi_0 := \phi'$. Then, for each $i \in \{0, \dots, k\}$, recursively find a $V[v_i, v_{i+1}]$-extension ϕ_{i+1} of ϕ_i if $V(v_i, v_{i+1})$ is non-empty, or set $\phi_{i+1} := \phi_i$ otherwise. The partial coloring ϕ_{k+1} is a $V[x, y]$-extension of ϕ. This completes the case that $S(x, y)$ is empty.

Case 2. The set $S(x, y)$ is non-empty.

It is convenient to call two segments *internally disjoint* if they have at most one vertex in common, and that vertex is an end of both segments. Ideally, we would be able to handle this case by finding a collection of internally disjoint valid segments whose union is $V[x, y]$. Each segment would then correspond, informally, to a “left window” or a “right window” as in the polygon case. However, for technical reasons, we will need to allow our “left” and “right” segments to not be internally disjoint.

Formally, we will define sets $\mathcal{L}$ and $\mathcal{R}$ which satisfy the following conditions:

(i) each set in $\mathcal{L} \cup \mathcal{R}$ is a valid segment whose ends are in $L[x, y) \cup R(x, y]$,

(ii) each pair of segments in $\mathcal{L}$ is internally disjoint, so is each pair of segments in $\mathcal{R}$, and if $V[a_1, b_1] \in \mathcal{L}$ and $V[a_2, b_2] \in \mathcal{R}$ are not internally disjoint, then $a_1 \prec a_2 \prec b_1 \prec b_2$;

(iii) each vertex in $V[x, y)$ belongs to $V[a, b)$ for some segment $V[a, b] \in \mathcal{L} \cup \mathcal{R}$, and each vertex in $V(x, y]$ belongs to $V(a, b]$ for some segment $V[a, b] \in \mathcal{L} \cup \mathcal{R}$;

(iv) if $V[a, b] \in \mathcal{L}$, then $b \in R(x, y]$, and if $V[a, b] \in \mathcal{R}$, then $a \in L[x, y)$;

(v) if $V[a, b] \in \mathcal{L}$, then $V(a, b] \cap L[x, y) \neq \emptyset$, and if $V[a, b] \in \mathcal{R}$, then $V[a, b) \cap R(x, y] \neq \emptyset$;

(vi) if $V[a, b] \in \mathcal{L}$, then $V(a, b) \cap R(x, y] = \emptyset$, and if $V[a, b] \in \mathcal{R}$, then $V(a, b) \cap L[x, y) = \emptyset$.

Let $s_1 \prec \cdots \prec s_k$ denote the strongly reachable vertices, with $k \geqslant 1$, and set $s_0 := x$ and $s_{k+1} := y$. For the following definitions, refer to Fig. 15. Throughout the rest of the

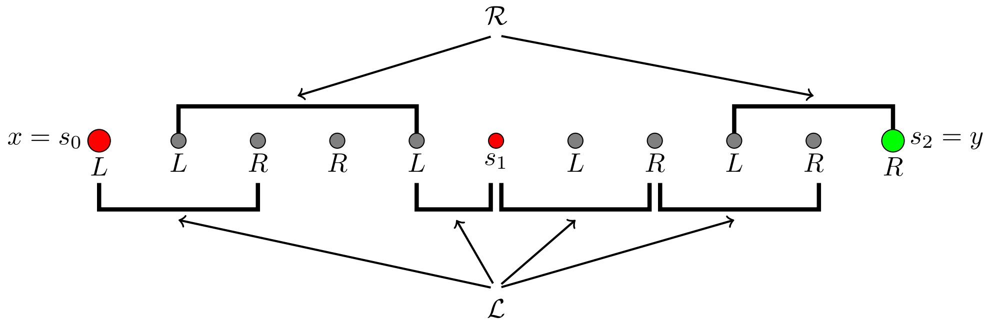

Fig. 15. A depiction of the sets $L[x, y)$, $R(x, y]$, $\mathcal{L}$, and $\mathcal{R}$.

proof, if we call a vertex right-, left-, or strongly reachable without specifying a segment, then we mean with respect to the segment $V[x, y]$.

Fix $i \in \{0, \ldots, k-1\}$. We take all minimal segments $V[\ell_1, \ell_2] \subseteq V[s_i, s_{i+1}]$ with left-reachable ends ℓ_1 and ℓ_2 such that at least one vertex in $V(\ell_1, \ell_2)$ is right-reachable, and we add them to $\mathcal{R}$. Every such segment $V[\ell_1, \ell_2]$ is valid; either ℓ_2 is right-reachable and $V[\ell_1, \ell_2]$ satisfies condition (i) of Lemma 4.6, or $\ell_2 \prec s_{i+1}$ and $V[\ell_1, \ell_2]$ satisfies condition (ii) of Lemma 4.6. Since s_i and s_{i+1} are left-reachable, every right-reachable vertex in $V(s_i, s_{i+1})$ is covered by some segment in $\mathcal{R}$.

Now we add segments to $\mathcal{L}$ to "cover the gaps in $V[s_i, s_{i+1}]$ between segments in $\mathcal{R}$". That is, for each maximal segment $V[\ell_2', \ell_1'] \subseteq V[s_i, s_{i+1}]$ which is internally disjoint from all segments in $\mathcal{R}$ (so that either $\ell_2' = s_i$ or ℓ_2' is the greater end of a segment in $\mathcal{R}$, and similarly for ℓ_1'), we add the segment $V[\ell_2', r']$ to $\mathcal{L}$, where r' is the least right-reachable vertex in $V[s_i, s_{i+1}]$ with $\ell_1' \preceq r'$. The segment $V[\ell_2', r']$ is valid because it satisfies condition (i) of Lemma 4.6.

This defines the segments in $\mathcal{L}$ and $\mathcal{R}$ that are contained in $V[s_i, s_{i+1}]$ for each $i \in \{1, \ldots, k-1\}$. For $V[s_k, s_{k+1}]$, however, this approach may not work, as condition (ii) of Lemma 4.6 need not hold. So within this segment, we switch the roles of $\mathcal{L}$ and $\mathcal{R}$, that is, we first add segments to $\mathcal{L}$ by considering minimal segments that end at right-reachable vertices and have at least one left-reachable internal vertex, and so on. The segments defined this way are valid by conditions (i) and (iii) of Lemma 4.6. This completes the definitions of $\mathcal{L}$ and $\mathcal{R}$; it can be verified that they satisfy conditions (i)–(vi) above.

Now let the two colors other than the common color of $s_1, \ldots, s_k$ be called "L" and "R". (How x and y are colored will not matter.) Our goal is to use the colors "L" and "R" to color the vertices which are strongly reachable from each valid segment in $\mathcal{L}$ and $\mathcal{R}$, respectively. So it is convenient to have a name for the vertices that might potentially receive color "L" (or color "R") in this step. An *R-vertex* is a vertex that belongs to $\{x\} \cup R(x, y] \setminus S(x, y)$ or is internal to some segment in $\mathcal{R}$. An *L-vertex* is a vertex that belongs to $\{y\} \cup L[x, y) \setminus S(x, y)$ or is internal to some segment in $\mathcal{L}$. We have the following claim about where "interesting" edges lie.

Claim 4.8. *For any two adjacent R-vertices (or L-vertices) u and v such that* $\{u, v\} \neq \{x, y\}$*, there is a segment in* $\mathcal{R}$ *(*$\mathcal{L}$*, respectively) that contains both u and v.*

Proof. Up to the symmetry obtained by reversing $\prec$, it suffices to prove the claim for R-vertices u and v. Without loss of generality, $u \prec v$. Since v is an R-vertex, condition (vi) yields $v \notin L[x, y)$. If $v = y$, then the edge yu forms a crossing sequence from y to u witnessing $u \in L[x, y)$, which contradicts condition (vi), as u is an R-vertex and $u \neq x$. Thus $v \notin L[x, y) \cup \{y\}$.

If there is a vertex $w \in L[x, y)$ with $u \prec w \prec v$, then there are crossing sequences from y to w and from u to v (namely, the edge uv), so Lemma 3.1 yields a crossing sequence from y to v witnessing $v \in L[x, y)$, a contradiction. Thus, there is no such vertex w.

Now suppose there is a segment $V[a, b] \in \mathcal{R}$ which contains v. Then $a \in L[x, y)$ by condition (iv). So $a \preceq u$ by the above, and $V[a, b]$ is a segment in $\mathcal{R}$ which contains both u and v, as desired. Thus, we may assume that there is no such segment.

It follows that $v \in R(x, y]$, as v is an R-vertex. By condition (iii), there is a segment $V[a, b] \in \mathcal{L} \cup \mathcal{R}$ such that $v \in V(a, b]$, and the above yields $V[a, b] \in \mathcal{L}$. Thus $b = v$ by condition (vi). By condition (v), there is a vertex $w \in V(a, b] \cap L[x, y)$, which implies $w \preceq u$ by the above, so $a \prec u$. By condition (vi), $u \notin R(x, y]$, so u is an internal vertex of a segment in $\mathcal{R}$ (as u is an R-vertex). Then, by condition (ii), that segment contains both u and v. $\square$

Now we complete the construction of a $V[x, y]$-extension of ϕ. Let $m = |\mathcal{L} \cup \mathcal{R}|$. Enumerate the intervals in $\mathcal{L} \cup \mathcal{R}$ as $V[a_i, b_i]$ for $i \in \{1, \dots, m\}$, in any order. Set $\phi_0 := \phi$. Then, for each $i \in \{1, \dots, m\}$, define ϕ_i from ϕ_{i-1} maintaining the condition that every color class of ϕ_i induces a capped subgraph, as follows. Consider the segment $V[a, b] := V[a_i, b_i] \in \mathcal{L} \cup \mathcal{R}$. If a or b is not colored, then color it "L" if it is in $L[x, y)$ and "R" if it is in $R(x, y]$. Furthermore, color all vertices in $S(a, b)$ color "L" if $V[a, b] \in \mathcal{L}$ and color "R" if $V[a, b] \in \mathcal{R}$. Call the new coloring ϕ'. We will show that ϕ' is a $V[a, b]$-precoloring. Thus, we can recursively compute a $V[a, b]$-extension of ϕ'. Take the extension, uncolor all vertices that are internal to any segment $V[a_j, b_j]$ with $j \in \{i + 1, \dots, m\}$, and call the resulting partial coloring ϕ_i. At the end, by condition (iii) above, the partial coloring ϕ_m is a requested $V[x, y]$-extension of ϕ.

Consider a step in the algorithm for the segment $V[a, b] = V[a_i, b_i] \in \mathcal{L} \cup \mathcal{R}$; we took a partial coloring ϕ_{i-1} with the property that every color class induced a capped subgraph, and we turned it into a partial coloring ϕ'. We just need to show the following.

Claim 4.9. *The partial 3-coloring* ϕ' *is a* $V[a, b]$*-precoloring.*

Proof. Up to the symmetry obtained by reversing $\prec$, we can assume that $V[a, b] \in \mathcal{R}$. By definition, ϕ' colors $S(a, b) \cup \{x, y\}$ and colors no other internal vertices of $V[a, b]$, and ϕ' gives every vertex in $S(a, b)$ color "R". Therefore, by Lemmas 4.4 and 4.5, it suffices to show that

(i) no external vertex u of $V[a,b]$ which is adjacent to an internal vertex v of $V[a,b]$ receives color "R", and
(ii) for each color class X of ϕ', the set $X \setminus V(a,b)$ induces a capped subgraph.

For the proof of (i), suppose that such vertices u and v exist. We claim that u is an R-vertex. Suppose not. Then u has received its color "R" as an internal vertex of some segment $V[a',b'] \in \mathcal{L}$. If $v \notin V[a',b']$, then $u \in S(a',b')$, so u should have received color "L" instead of "R". Thus $v \in V[a',b']$ and therefore, by condition (ii) on $\mathcal{L}$ and $\mathcal{R}$, we have $a' \prec u \prec a \prec v \preceq b' \prec b$. Since $a \in L[x,y]$ by condition (iv) on $\mathcal{R}$, there are crossing sequences from y to a and from u to v (namely, the edge uv), so Lemma 3.1 yields a crossing sequence from y to v witnessing $v \in L[x,y]$ and contradicting condition (vi) on $\mathcal{R}$. Now, since u and v are R-vertices, Claim 4.8 yields a segment in $\mathcal{R}$ containing both u and v, which contradicts the fact that segments in $\mathcal{R}$ are pairwise internally disjoint (from condition (ii) on $\mathcal{R}$).

Now, we prove (ii). Since we assume that every color class of ϕ_{i-1} induces a capped subgraph, it suffices to prove that the subgraph induced by $X \setminus V(a,b)$ contains no crossing pair of edges involving an edge uv with a vertex $v \in \{a,b\}$ newly colored by ϕ'. So suppose such a crossing pair exists. Then $u \notin \{a,b\}$, so $u \notin V[a,b]$ and u is already colored by ϕ_{i-1}.

Suppose for the sake of contradiction that some segment $V[a',b'] \in \mathcal{L} \cup \mathcal{R}$ contains both u and v. If $V[a',b']$ has been considered before $V[a,b]$, then, by condition (ii) on $\mathcal{L}$ and $\mathcal{R}$, no segment in $\mathcal{L} \cup \mathcal{R}$ other than $V[a,b]$ and $V[a',b']$ contains v, so v must have been colored when $V[a',b']$ has been considered (and not uncolored afterwards). If $V[a',b']$ has not been considered yet, then no vertex in $V(u,v)$ can be colored by ϕ_{i-1} (as it is internal to a segment that has not been considered), so the edge uv cannot participate in a crossing pair of edges induced by $X \setminus V(a,b)$. This shows that no segment in $\mathcal{L} \cup \mathcal{R}$ contains both u and v.

If $u \notin \{x,y\}$, then u has received its color "L" or "R" as $\phi_j(u)$ for some $j \in \{1,\dots,i-1\}$ such that $u \in V[a_j,b_j]$. By the above, $v \notin V[a_j,b_j]$, so $u \in \{a_j,b_j\} \cup S(a_j,b_j)$. Therefore, u and v are L-vertices if $\phi_j(u) = \phi'(v) =$ "L" or R-vertices if $\phi_j(u) = \phi'(v) =$ "R". Now Claim 4.8 yields a segment in $\mathcal{R}$ which contains both u and v. This is a contradiction, which completes the proof of Claim 4.9. □

Claim 4.9 above completes the proof of Proposition 4.7. □

Proof of Proposition 4.2. Let $(G',\prec')$ be the ordered graph obtained from $(G,\prec)$ by adding a new smallest vertex x, a new largest vertex y, and the edge xy. Let ϕ' be any partial coloring which colors only x and y. Then $(G',\prec')$ is $\mathcal{H}$-free, $V[x,y]$ is a valid segment, and ϕ' is a $V[x,y]$-precoloring. Therefore, by Proposition 4.7, we can find (in polynomial time) a 3-coloring of $(G,\prec)$ such that each color class induces a capped subgraph. □

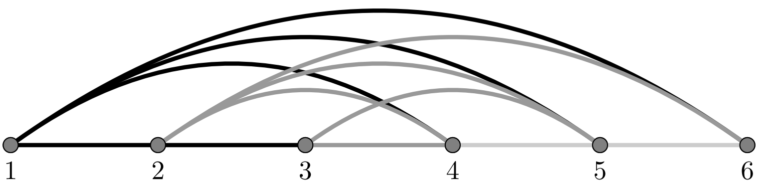

Fig. 16. The decomposition from Proposition 5.1, where darker edges are removed first.

5. Coloring capped graphs

This section is devoted to the proof of Theorem 1.4 on coloring 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 henceforth that G is not triangle-free. We say that an edge $uv \in E(G)$ with $u \prec v$ is *triangle-covered* 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-covered (see Fig. 16).

If a, b, c are vertices of G such that $a \prec b \prec c$, $ac, bc \in E(G)$, and ac is triangle-covered, then bc is triangle-covered. 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-covered. 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 $(G - F, \prec)$.*

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-covered, then ac is triangle-covered, as $(G, \prec)$ is capped. 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-covered. 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, as $(G, \prec)$ is capped. This shows that the clique number of $(G - F, \prec)$ is at most $\omega - 1$, as desired. $\square$

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

Proof of Theorem 1.4. Let $\omega \geqslant 2$ be the clique number of G, and let $\{(G_i, \prec)\}_{i=1}^{\omega-1}$ be a decomposition of $(G, \prec)$ into $\omega - 1$ triangle-free capped subgraphs as in Proposition 5.1. Fix an index i with $1 \leqslant i \leqslant \omega - 1$. Say that an edge bd with $b \prec d$ is *crossed* in $(G_i, \prec)$ if there is an edge $ac \in E(G_i)$ with $a \prec b \prec c \prec d$. 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 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 $c_1 \prec \cdots \prec c_k$ 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 $(G_i, \prec)$ is capped and $c_1, \dots, c_k$ induce an ordered hole in $(G_i, \prec)$. Then, since $(G_i, \prec)$ is capped, the vertices 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. $\square$

Claim 5.4. *There is a 4-coloring 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 colors 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 $bv_1 \cdots v_t c$ with $b \prec v_1 \prec \cdots \prec v_t \prec c$, 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, since $(G_i - F_i, \prec)$ is capped, the set $\{v_1', v_1, \dots, v_t, c\}$ contains a triangle or an ordered hole, which is a contradiction. In the final case that $d \prec v_1'$, since $(G_i - F_i, \prec)$ is capped, the vertices b, v_1, v_1' form a triangle, which is again a contradiction. $\square$

Let ϕ_i be a 4-coloring of $G_i - F_i$ from the last claim, for $1 \leqslant i < \omega$. Let $F = \bigcup_{i=1}^{\omega-1} F_i$. If $(G, \prec)$ is ordered-hole-free, then $F = \emptyset$ and the mapping $v \mapsto (\phi_1(v), \ldots, \phi_{\omega-1}(v))$ is a $4^{\omega-1}$-coloring 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 $n \geqslant 2$. So every non-empty subgraph of $G[F]$ has a vertex of degree less than $4(\omega-1)$, and thus there exists a $4(\omega-1)$-coloring ψ of $G[F]$. Now, the mapping $v \mapsto (\phi_1(v), \ldots, \phi_{\omega-1}(v), \psi(v))$ is a $4^\omega(\omega-1)$-coloring of G. $\square$

We note that the bound in Theorem 1.4 would be substantially improved if one could show that every triangle-free, capped, and ordered-hole-free graph has chromatic number at most 3 (that is, if one could improve the bound in Claim 5.4 from 4 to 3). While these graphs are not necessarily bipartite, we do not know of an example that actually has chromatic number 4.

6. Coloring X-free ordered graphs

In this section, we prove Theorem 1.6 that X-free ordered graphs are χ-bounded, using a surprising connection to a recent theorem of Scott and Seymour [34]. Recall that X is the ordered graph with four vertices and two crossing edges, illustrated in Fig. 3 (right). It is immediate from the definition that capped graphs are X-free. Moreover, various natural classes of ordered intersection graphs are also X-free. Specifically, an *outerstring graph* is the intersection graph of a collection of curves in a half-plane which each have one endpoint on the boundary of that half-plane. It is easy to see that outerstring graphs are X-free when equipped with the ordering of their endpoints along the boundary. Outerstring graphs contain circle graphs and interval filament graphs. Rok and Walczak [33] proved that the class of outerstring graphs is χ-bounded. Theorem 1.6 yields this results as a direct corollary.

We now prove that the underlying (unordered) graph of any X-free ordered graph forbids every induced subdivision of a certain banana. A *banana* is a graph which can be obtained from the disjoint union of paths by identifying one end of each path to a new vertex s, and the other end of each path to a different new vertex t. A *subdivision* of a graph H is a graph obtained from H by replacing edges of H with internally disjoint paths between their ends. Scott and Seymour [34] proved the following.

Theorem 6.1 *(Scott and Seymour [34]). For any banana B, the class of graphs excluding all subdivisions of B as induced subgraphs is χ-bounded.*

Let B_4 denote the banana depicted in Fig. 3 (left); it is obtained from three 4-edge paths. Theorem 6.1 together with the following proposition yields Theorem 1.6 and, consequently, another proof that outerstring graphs and capped graphs are χ-bounded.

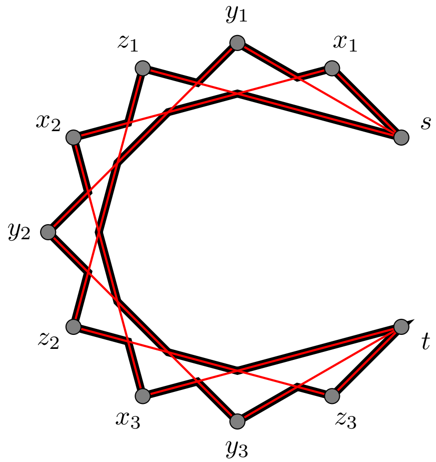

Fig. 17. B_4 as an induced subgraph of a polygon visibility graph.

Proposition 6.2. *For any X-free ordered graph $(H, \prec)$, the graph H contains no subdivision of B_4 as an induced subgraph.*

Proof. First we prove that B_4 itself is not the underlying graph of an X-free ordered graph. Suppose for a contradiction that there is a vertex ordering $\prec$ such that $(B_4, \prec)$ is X-free. Any cyclic reordering of $\prec$ also yields an X-free ordering. So, as none of the edges sx_1, sy_1, sz_1 cross any of the edges tx_3, ty_3, tz_3, after cyclically reordering the vertices, we can assume that each of s, x_1, y_1, z_1 is less than each of t, x_3, y_3, z_3.

Consider the three paths x_1, x_2, x_3, y_1, y_2, y_3, and z_1, z_2, z_3. No edge of one of these paths may cross an edge of another. So, for instance, we cannot have that $x_1 \prec y_1 \prec x_3 \prec y_3$. Then, up to relabeling these three paths, we can assume that $x_1 \prec y_1 \prec z_1 \prec z_3 \prec y_3 \prec x_3$. Now we examine where the vertex y_2 lies in the ordering. For the same reason, either $x_1 \prec y_2 \prec z_1$ or $z_3 \prec y_2 \prec x_3$; without loss of generality, we can assume the first case. But now one of the two edges sx_1 or sz_1 is crossing with the edge y_2y_3, a contradiction.

Now suppose that $(G, \prec)$ is an X-free ordered graph and v is a degree-2 vertex which is adjacent to non-adjacent vertices u and w. We will show that, where H is the graph obtained from G by deleting v and adding the edge uw, and $\prec_H$ is the restriction of $\prec$ to $V(H)$, the ordered graph $(H, \prec_H)$ is X-free. This will complete the proof. Up to cyclically reordering the vertices and renaming u and w, the only case to consider is that $u \prec v \prec w$. If $(H, \prec_H)$ is not X-free, then there must be a copy of X containing the edge uw. But then either uv or vw is in a copy of X in $(G, \prec)$, which is a contradiction. $\square$

Scott and Seymour [34] actually proved Theorem 6.1 for *banana trees*, which are obtained from trees by replacing each edge with a banana so that the ends of the edge are s and t (the high-degree vertices of the banana). It is natural to ask if this theorem could be applied directly to polygon visibility graphs. However, this is not possible; it is not difficult to see that every banana tree is an induced subgraph of a polygon visibility graph. Fig. 17 illustrates this for the banana B_4.

Data availability

No data was used for the research described in the article.

Acknowledgment

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

References

[1] James Abello, Ömer Egecioglu, Krishna Kumar, Visibility graphs of staircase polygons and the weak Bruhat order, I: from visibility graphs to maximal chains, Discrete Comput. Geom. 14 (3) (1995) 331–358.

[2] James Abello, Krishna Kumar, Visibility graphs and oriented matroids, Discrete Comput. Geom. 28 (4) (2002) 449–465.

[3] Safwa Ameer, Matt Gibson-Lopez, Erik Krohn, Sean Soderman, Qing Wang, Terrain visibility graphs: persistence is not enough, in: Sergio Cabello, Danny Z. Chen (Eds.), 36th International Symposium on Computational Geometry (SoCG 2020), in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 164, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, 2020, 6.

[4] Safwa Ameer, Matt Gibson-Lopez, Erik Krohn, Qing Wang, On the visibility graphs of pseudo-polygons: recognition and reconstruction, in: Artur Czumaj, Qin Xin (Eds.), 18th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2022, in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 227, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, 2022, 7.

[5] Alan Arroyo, Julien Bensmail, R. Bruce Richter, Extending drawings of graphs to arrangements of pseudolines, in: Sergio Cabello, Danny Z. Chen (Eds.), 36th International Symposium on Computational Geometry (SoCG 2020), in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 164, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, 2020, 9.

[6] Edgar Asplund, Branko Grünbaum, On a coloring problem, Math. Scand. 8 (1960) 181–188.

[7] David Avis, David Rappaport, Computing the largest empty convex subset of a set of points, in: Proceedings of the First Annual Symposium on Computational Geometry, Association for Computing Machinery, New York, 1985, pp. 161–167.

[8] Maria Axenovich, Jonathan Rollin, Torsten Ueckerdt, Chromatic number of ordered graphs with forbidden ordered subgraphs, Combinatorica 38 (5) (2018) 1021–1043.

[9] Andreas Bärtschi, Subir K. Ghosh, Matúš Mihalák, Thomas Tschager, Peter Widmayer, Improved bounds for the conflict-free chromatic art gallery problem, in: Proceedings of the Thirtieth Annual Symposium on Computational Geometry, Association for Computing Machinery, New York, 2014, pp. 144–153.

[10] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, Rémi Watrigant, Twin-width III: max independent set, min dominating set, and coloring, in: Nikhil Bansal, Emanuela Merelli, James Worrell (Eds.), 48th International Colloquium on Automata, Languages, and Programming (ICALP 2021), in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 198, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, 2021, 35.

[11] Marcin Briański, James Davies, Bartosz Walczak, Separating polynomial χ-boundedness from χ-boundedness, arXiv:2201.08814, 2022.

[12] Onur Çağırıcı, Subir K. Ghosh, Petr Hliněný, Bodhayan Roy, On conflict-free chromatic guarding of simple polygons, in: Yingshu Li, Mihaela Cardei, Yan Huang (Eds.), Combinatorial Optimization and Applications, in: Lecture Notes in Computer Science, vol. 11949, Springer, Cham, 2019, pp. 601–612.

[13] Onur Çağırıcı, Petr Hliněný, Bodhayan Roy, On colourability of polygon visibility graphs, arXiv: 1906.01904, 2019.
[14] Jean Cardinal, Udo Hoffmann, Recognition and complexity of point visibility graphs, Discrete Comput. Geom. 57 (1) (2017) 164–178.
[15] Daniele Catanzaro, Steven Chaplick, Stefan Felsner, Bjarni V. Halldórsson, Magnús M. Halldórsson, Thomas Hixon, Juraj Stacho, Max point-tolerance graphs, Discrete Appl. Math. 216 (1) (2017) 84–97.
[16] Maria Chudnovsky, Alex Scott, Paul Seymour, Induced subgraphs of graphs with large chromatic number. V. Chandeliers and strings, J. Comb. Theory, Ser. B 150 (2021) 195–243.
[17] Alice M. Dean, William Evans, Ellen Gethner, Joshua D. Laison, Mohammad Ali Safari, William T. Trotter, Bar k-visibility graphs: bounds on the number of edges, chromatic number, and thickness, in: Patrick Healy, Nikola S. Nikolov (Eds.), Graph Drawing, in: Lecture Notes in Computer Science, vol. 3843, Springer, Berlin, Heidelberg, 2006, pp. 73–82.
[18] Stephan Eidenbenz, Christoph Stamm, Maximum clique and minimum clique partition in visibility graphs, in: Jan van Leeuwen, Osamu Watanabe, Masami Hagiya, Peter D. Mosses, Takayasu Ito (Eds.), Theoretical Computer Science: Exploring New Frontiers of Theoretical Informatics, Springer, Berlin, Heidelberg, 2000, pp. 200–212.
[19] Louis Esperet, Graph Colorings, Flows and Perfect Matchings, Habilitation thesis, Université Grenoble Alpes, 2017.
[20] William S. Evans, Noushin Saeedi, On characterizing terrain visibility graphs, J. Comput. Geom. 6 (1) (2015) 108–141.
[21] Subir K. Ghosh, On recognizing and characterizing visibility graphs of simple polygons, Discrete Comput. Geom. 17 (2) (1997) 143–162.
[22] Subir K. Ghosh, Partha P. Goswami, Unsolved problems in visibility graphs of points, segments, and polygons, ACM Comput. Surv. 46 (2) (2013) 22.
[23] Subir K. Ghosh, Thomas C. Shermer, Binay K. Bhattacharya, Partha P. Goswami, Computing the maximum clique in the visibility graph of a simple polygon, J. Discret. Algorithms 5 (3) (2007) 524–532.
[24] András Gyárfás, On the chromatic number of multiple interval graphs and overlap graphs, Discrete Math. 55 (2) (1985) 161–166.
[25] Jan Kára, Attila Pór, David R. Wood, On the chromatic number of the visibility graph of a set of points in the plane, Discrete Comput. Geom. 34 (3) (2005) 497–506.
[26] Friedrich Levi, Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade, Ber. Verh. K. Sächs. Ges. Wiss. Leipz., Math.-Phys. Kl. 78 (1926) 256–267.
[27] Fabrizio Luccio, Silvia Mazzone, Chak Kuen Wong, A note on visibility graphs, Discrete Math. 64 (2) (1987) 209–219.
[28] Joseph O'Rourke, Ileana Streinu, Vertex-edge pseudo-visibility graphs: characterization and recognition, in: Proceedings of the Thirteenth Annual Symposium on Computational Geometry, Association for Computing Machinery, New York, 1997, pp. 119–128.
[29] János Pach, István Tomon, On the chromatic number of disjointness graphs of curves, J. Comb. Theory, Ser. B 144 (2020) 167–190.
[30] Arkadiusz Pawlik, Jakub Kozik, Tomasz Krawczyk, Michał Lasoń, Piotr Micek, William T. Trotter, Bartosz Walczak, Triangle-free intersection graphs of line segments with large chromatic number, J. Comb. Theory, Ser. B 105 (2014) 6–10.
[31] Florian Pfender, Visibility graphs of point sets in the plane, Discrete Comput. Geom. 39 (1–3) (2008) 455–459.
[32] Alexandre Rok, Bartosz Walczak, Coloring curves that cross a fixed curve, Discrete Comput. Geom. 61 (4) (2019) 830–851.
[33] Alexandre Rok, Bartosz Walczak, Outerstring graphs are χ-bounded, SIAM J. Discrete Math. 33 (4) (2019) 2181–2199.
[34] Alex Scott, Paul Seymour, Induced subgraphs of graphs with large chromatic number. VI. Banana trees, J. Comb. Theory, Ser. B 145 (2020) 487–510.
[35] Ileana Streinu, Non-stretchable pseudo-visibility graphs, Comput. Geom. 31 (3) (2005) 195–206.
[36] Subhash Suri, A linear time algorithm for minimum link paths inside a simple polygon, Comput. Vis. Graph. Image Process. 35 (1) (1986) 99–110.
[37] István Tomon, Personal communication.