

Polygon-Universal Graphs

Tim Ophelders ✉

Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands

Ignaz Rutter ✉ 

Department of Computer Science and Mathematics, Universität Passau, Germany

Bettina Speckmann ✉ 

Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands

Kevin Verbeek ✉

Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands

Abstract

We study a fundamental question from graph drawing: given a pair (G, C) of a graph G and a cycle C in G together with a simple polygon P , is there a straight-line drawing of G inside P which maps C to P ? We say that such a drawing of (G, C) *respects* P . We fully characterize those instances (G, C) which are *polygon-universal*, that is, they have a drawing that respects P for any simple (not necessarily convex) polygon P . Specifically, we identify two necessary conditions for an instance to be polygon-universal. Both conditions are based purely on graph and cycle distances and are easy to check. We show that these two conditions are also sufficient. Furthermore, if an instance (G, C) is planar, that is, if there exists a planar drawing of G with C on the outer face, we show that the same conditions guarantee for every simple polygon P the existence of a planar drawing of (G, C) that respects P . If (G, C) is polygon-universal, then our proofs directly imply a linear-time algorithm to construct a drawing that respects a given polygon P .

2012 ACM Subject Classification Human-centered computing → Graph drawings

Keywords and phrases Graph drawing, partial drawing extension, simple polygon

Digital Object Identifier 10.4230/LIPIcs.SoCG.2021.55

Related Version *Full Version*: <https://arxiv.org/abs/2103.06916>

Funding Research on the topic of this paper was initiated at the 2nd Workshop on Applied Geometric Algorithms in Vierhouten, NL, supported by the Dutch Research Council (NWO); 639.023.208.

Ignaz Rutter: Partially supported by the German Science Foundation (DFG); Ru 1903/3-1.

Bettina Speckmann: Partially supported by the Dutch Research Council (NWO); 639.023.208.

Acknowledgements Ignaz Rutter would like to thank Michael Hoffmann and Vincent Kusters for discussions on conjectures related to this paper.

1 Introduction

Graphs are a convenient way to express relations between entities. To visualize these relations, the corresponding graph needs to be *drawn*, most commonly in the plane and with straight edges. Naturally there are a multitude of different optimization criteria and drawing restrictions that attempt to capture various perceptual requirements or real-world conditions. In this paper we focus on drawings which are constrained to the interiors of simple polygons.

The *polygon-extension problem* asks, whether a given graph admits a (planar) drawing where the outer face is fixed to a given simple polygon P ; see Fig. 1 for examples. Our main focus is the *polygon-universality problem*, which asks whether a given plane graph admits a polygon-extension for every choice of fixing the outer face to a simple polygon. As is often the case with geometric problems, a natural complexity class for the polygon-extension problem is $\exists\mathbb{R}$, the class of problems that can be encoded in polynomial time as an



© Tim Ophelders, Ignaz Rutter, Bettina Speckmann, and Kevin Verbeek;
licensed under Creative Commons License CC-BY 4.0


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

Editors: Kevin Buchin and Éric Colin de Verdière; Article No. 55; pp. 55:1–55:15

Leibniz International Proceedings in Informatics



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



existentially quantified formula of real variables (rather than Boolean variables as for SAT), which was introduced by Schaefer and Štefankovič [12]. The natural complexity class for the polygon-universality problem is $\forall\exists\mathbb{R}$, the universal existential theory of the reals, which has been recently defined by Dobbins et al. [5]. It is known that $\text{NP} \subseteq \exists\mathbb{R} \subseteq \forall\exists\mathbb{R} \subseteq \text{PSPACE}$ [3].


Tutte [13] proved that there is a straight-line planar drawing of a planar graph G inside an arbitrary convex polygon P if one fixes the outer face of (an arbitrary planar embedding of) G to P . This result has been generalized to allow polygons P that are non-strictly convex [4, 6] or even star-shaped polygons [7]. These results have applications in *partial drawing extension* problems. Here, in addition to an input graph G , we are given a subgraph $H \subseteq G$ together with a fixed drawing Γ of H . The question is whether one can extend the given drawing Γ to a planar straight-line drawing of the whole graph G by drawing the vertices and edges of $G - H$ inside the faces of H . If the embedding of G is fixed, the results by Tutte and others allow to reduce the problem by removing vertices of G that are contained in convex or star-shaped faces of Γ . Such reduction rules have led to efficient testing algorithms for special cases, for example, when the drawing of H is convex [10].

Recently, Lubiw et al. [9] showed that it is $\exists\mathbb{R}$ -complete to decide for a given planar graph that is partially fixed to a non-crossing polygon *with holes*, whether the partial drawing can be extended to a planar straight-line drawing that does not intersect the outside of the polygon. That is, the planar polygon extension problem is $\exists\mathbb{R}$ -complete for polygons with holes. They leave the case of simple polygons open.

If we do not insist on straight-line drawings, then other questions arise. Angelini et al. [2] give an $O(mn)$ -time algorithm for testing whether an n -vertex outer-planar graph admits a planar one-bend drawing whose outer face is fixed to a simple polygon on m vertices. Mchedlidze and Urhausen [11] link the number of bends per edge that are necessary for extending a drawing to a convexity measure for the faces of the partial drawing. Angelini et al. [1] present a linear time algorithm to decide if a partially fixed drawing has a planar drawing using Jordan arcs, while Jelínek et al. [8] characterize the solvable instances by forbidden substructures.

Quite recently, Dobbins et al. [5] considered the problem of area-universality for graphs with partial drawings. Let G be a planar graph with a fixed embedding, including the outer face. An assignment of areas to faces of G is *realizable* if G admits a straight-line drawing such that each face has the assigned area. A graph is *area-universal* if every area assignment is realizable. Dobbins et al. prove that it is $\exists\mathbb{R}$ -complete to decide whether an area assignment is realizable for a planar triangulation, which has been partially drawn, and testing area-universality is $\forall\exists\mathbb{R}$ -complete if the planarity condition is dropped (but still parts of the drawing are fixed). They conjecture that the same area-universality problems without a partially fixed drawing are $\exists\mathbb{R}$ - and $\forall\exists\mathbb{R}$ -complete in general.


Notation. Let $G = (V, E)$ be a graph with n vertices. A *drawing* \mathcal{D} of G is a map from each $v \in V$ to points in the plane and from each edge $e \in E$ to a Jordan arc connecting its endpoints. A *straight-line* drawing maps each edge to a straight line segment. A drawing is *planar*, if no two edges intersect, except at common endpoints. A graph G is planar if it has a planar drawing. Let $C = [c_1, \dots, c_t]$ with $c_i \in V$ be a simple cycle in G . An *instance* (G, C) is planar if G has a planar drawing with C as the outer face. Let P be a simple polygon with t vertices $[p_1, \dots, p_t]$ with $p_i \in \mathbb{R}^2$. A drawing \mathcal{D} of (G, C) *respects* P if it is a map $\mathcal{D}: V \rightarrow P$ from vertices to points in P such that $\mathcal{D}(c_i) = p_i$ and for each edge $\{u, v\} \in E$, the line segment between $\mathcal{D}(u)$ and $\mathcal{D}(v)$ lies in P (see Figure 1). That is, \mathcal{D} is a straight-line



■ **Figure 1** Left: an instance (G, C) . Center left: a drawing of (G, C) that respects a polygon P (shaded in grey). Center right: there is no drawing of (G, C) that respects this polygon. Right: A triangulated convex polygon with a drawing that is not triangulation-respecting; moving the vertices along the dashed arrows results in a triangulation-respecting drawing.

drawing of G inside P that fixes the vertices of C to the corresponding vertices of P . An instance (G, C) is *(planar) polygon-universal* if it admits a (planar) straight-line drawing that respects every simple (not necessarily convex) polygon P on t vertices.

Our algorithms use a triangulation \mathcal{T} of P to construct a drawing or prove the non-universality. We say that a drawing of (G, C) *respects* \mathcal{T} if no edge of G properly crosses an edge of \mathcal{T} . Although every triangulation-respecting drawing is also a drawing, the converse is not true. In fact, there are graphs G , cycles C , and polygons P that have a drawing, but no triangulation of P exists that allows a triangulation-respecting drawing (see Figure 2).



■ **Figure 2** A graph G and a polygon for which there exists a drawing of G , but no triangulation with a triangulation-respecting drawing.

Results and organization. In Section 2 we identify two necessary conditions for an instance (G, C) to be *polygon-universal*. These conditions are purely based on graph and cycle distances and hence easy to check. To show that these two conditions are also sufficient, we use triangulation-respecting drawings: if there is a triangulation \mathcal{T} of a simple polygon P such that (G, C) does not admit a triangulation-respecting drawing for \mathcal{T} , then we can argue that G contains one of two forbidden substructures, violating the necessary conditions (see Section 5). These substructures certify that (G, C) is not polygon-universal.

To arrive at this conclusion, in Section 3 we first present an algorithm that tests in linear time for a given instance (G, C) , a polygon P , and triangulation \mathcal{T} of P , whether there exists a triangulation-respecting drawing of G for \mathcal{T} inside P . If so, we can construct the drawing in linear time. Then, in Section 4, we consider planar instances (G, C) and show that the same algorithm can decide in linear time whether there is a triangulation-respecting drawing that is planar after infinitesimal perturbation. An analysis of this algorithm shows that a planar instance (G, C) is planar polygon-universal if and only if it is polygon-universal.

2 Necessary conditions for polygon-universality


We present two necessary conditions for an instance (G, C) to be polygon-universal. Intuitively, both conditions capture that there need to be “enough” vertices in G between cycle vertices for the drawing not to become “too tight”. The *Pair Condition* captures this for any two vertices on the cycle C . The *Triple Condition* is a bit more involved: even if the Pair Condition is satisfied for any pair of vertices on the cycle, there can still be triples of vertices which together “pull too much” on the graph. Specifically, for an instance (G, C) of a graph G and a cycle $C \subset G$ with t vertices, we denote by $d_G: V \times V \rightarrow \mathbb{N}$ the graph distance in G and by $d_C: V(C) \times V(C) \rightarrow \mathbb{N}$ the distance (number of edges) along the cycle C . The following conditions are necessary for (G, C) to be polygon-universal for all simple polygons P :

Pair For all i and j , we have $d_C(c_i, c_j) \leq d_G(c_i, c_j)$ (and hence $d_C(c_i, c_j) = d_G(c_i, c_j)$).

Triple For all vertices $v \in V$ and distinct i, j, k with $d_C(c_i, c_j) + d_C(c_j, c_k) + d_C(c_i, c_k) \geq t$ (and hence $= t$), we have $d_G(c_i, v) + d_G(c_j, v) + d_G(c_k, v) > t/2$.

To establish that these two conditions are necessary, we use the *link distance* between two points inside certain simple polygons P . Specifically, the link distance of two points q_1 and q_2 with respect to a simple polygon P is the minimum number of segments for a polyline π that lies inside P and connects q_1 and q_2 . If the Pair Condition is violated for two cycle vertices c_i and c_j , we can construct a *Pair Spiral* polygon P (see Figure 3 (left)) such that the link distance between p_i and p_j (the vertices of P to which c_i and c_j are mapped) exceeds $d_G(c_i, c_j)$. Clearly there is no drawing (G, C) that respects P .

If the first condition holds, but the second condition is violated by a vertex v , consider the shortest paths via v that connect c_i, c_j, c_k to each other. By assumption the total length of these three paths is $2d_G(c_i, v) + 2d_G(c_j, v) + 2d_G(c_k, v) \leq t$, while the total length of the paths connecting c_i, c_j, c_k to each other along C is $d_C(c_i, c_j) + d_C(c_j, c_k) + d_C(c_i, c_k) = t$. Since the pair condition holds, the paths via v are not shorter than the paths along C , and therefore the paths via v must be shortest paths connecting the pairs. That is, $d_C(c_i, c_j) = d_G(c_i, v) + d_G(c_j, v)$, $d_C(c_j, c_k) = d_G(c_j, v) + d_G(c_k, v)$ and $d_C(c_i, c_k) = d_G(c_i, v) + d_G(c_k, v)$. In that case, we can construct a *Triple Spiral* polygon P (see Figure 3 (right)) such that there is no point that lies within link-distance $d_G(c_i, v)$ from c_i , link-distance $d_G(c_j, v)$ from c_j , and link-distance $d_G(c_k, v)$ from c_k simultaneously. Hence, there exists no drawing of the aforementioned shortest paths via v that respects P .



■ **Figure 3** Left: Pair Spiral. Points with link-distance greater than $d_G(c_i, c_j)$ from p_i shaded red. Right: Triple Spiral. Points of link-distance $\leq d_G(c_x, v)$ from p_x for one $x \in \{i, j, k\}$ in light gray; for two $x \in \{i, j, k\}$ in dark gray; there is no point q in P with $d_G(c_x, q) \leq d_G(c_x, v)$ for all $x \in \{i, j, k\}$.

3 Triangulation-respecting drawings

In this section we are given the following input: an instance (G, C) consisting of a graph G with n vertices and a cycle C with t vertices, and a simple polygon P with t vertices together with an arbitrary triangulation \mathcal{T} of P . We study the following question: is there a drawing of (G, C) that respects both P and \mathcal{T} ?

We describe a dynamic programming algorithm which can answer this question in linear time. The basic idea is as follows: every edge of \mathcal{T} defines a *pocket* of P . We recursively *sketch* a drawing of G within each pocket. Such a sketch assigns an approximate location, such as an edge or a triangle, to each vertex. Ultimately we combine the location constraints on vertex positions posed by the sketches and decide if they can be satisfied.


We root (the dual tree of) \mathcal{T} at an arbitrary triangle T_{root} . Each edge e of \mathcal{T} partitions P into two regions, one of which contains T_{root} . Let Q be the region not containing T_{root} . We say that Q is a *pocket* with the *lid* $e = e_Q$, and we denote the unique triangle outside Q adjacent to e_Q by T_Q^+ . Since a pocket is uniquely defined by its lid, we will for an edge e also write Q_e to denote the pocket with lid e . We say that a pocket is *trivial* if its lid lies on the boundary of P ; in such case the pocket consists of only that edge. If Q is a non-trivial pocket, then we denote the unique triangle inside Q adjacent to e_Q by T_Q (see Figure 4).

For ease of explanation we consider all indices on C and P modulo t , that is, we identify c_i and c_{i+t} as well as p_i and p_{i+t} . Moreover, when talking about a non-trivial pocket Q with lid (p_i, p_j) , whose third vertex of T_Q is p_k , we will assume that $i \leq k \leq j$ (otherwise simply shift the indices cyclically). We first define triangulation-respecting drawings for pockets:

► **Definition 1.** A triangulation-respecting drawing for a pocket Q with lid $e_Q = (p_i, p_j)$ is an assignment of the vertices of G to locations inside the polygon P , such that

1. Any vertex c_ℓ with $i \leq \ell \leq j$ is assigned to the polygon vertex p_ℓ .
2. For any edge (u, v) of G , u and v lie on a common triangle (or edges or vertices thereof). We consider the triangles of the triangulation as closed, so that distinct triangles may share a segment (namely an edge of the triangulation) or a point (namely a vertex of P). We define a triangulation-respecting drawing for the entire triangulation analogously, requiring that c_ℓ is assigned to p_ℓ for all ℓ .

A *sketch* is an assignment of the vertices of G to simplices (vertices, edges, or triangles) of the triangulation with the property that, if we draw each vertex anywhere on its assigned simplex, then the result is a triangulation-respecting drawing. We hence interpret a simplex as a closed region of the plane in the remainder of this paper.



■ **Figure 4** A triangulation with labels for the pocket Q (shaded dark) and triangles T_Q and T_Q^+ incident to edge e_Q of the triangulation.

► **Definition 2.** A sketch of the triangulation is a function Γ that assigns vertices of G to simplices of \mathcal{T} , such that (i) for any vertex c_i of the cycle, $\Gamma(c_i) = p_i$, and (ii) for any two adjacent vertices u and v , there exists a triangle of \mathcal{T} that contains both $\Gamma(u)$ and $\Gamma(v)$. A sketch of a pocket is defined similarly, except that vertices p_i of the polygon that lie outside the pocket do not need c_i assigned to them.

We show that a sketch exists (for a pocket or a triangulation) if and only if there is a triangulation-respecting drawing (for that pocket or triangulation). If a pocket admits a sketch, we call a pocket *sketchable*. If a particular pocket is sketchable, then so are all of its subpockets, since any sketch for a pocket is also a sketch for any of its subpockets.

We present an algorithm that for any sketchable pocket constructs a sketch, and for any other pocket reports that it is not sketchable. This algorithm recursively constructs particularly well-behaved sketches for child pockets, and combines these sketches into a new well-behaved sketch. To obtain a sketch for \mathcal{T} , we combine the three well-behaved sketches for the three pockets of the root triangle T_{root} – assuming that all three pockets are sketchable.

Well-behaved sketches. We restrict our attention to *local* sketches for a pocket Q , which assign vertices either to simplices in Q or to the triangle T_Q^+ just outside Q , and *interior* local sketches, which assign vertices to simplices in Q only.

► **Lemma 3.** *If there is a sketch for pocket Q , then there is a local sketch for Q .*

Generally, it is advantageous for a sketch to place its vertices as far “to the outside” as possible, to generate maximum flexibility when combining sketches. Hence, we introduce a preorder \preceq_Q on local sketches of a pocket Q , defined as $\Gamma \preceq_Q \Gamma'$ iff $\Gamma(v) \cap T_Q^+ \subseteq \Gamma'(v)$ for all vertices v . Intuitively, maximal elements with respect to this preorder maximize for each vertex, the intersection of its assigned simplex with T_Q^+ . We call a local sketch Γ of Q *well-behaved* if it is maximal with respect to \preceq_Q , and *interior well-behaved* if it is maximal among all interior local sketches of Q . A similar preorder and notion of well-behaved can be defined for sketches of the entire triangulation, by replacing T_Q^+ by T_{root} in the definition.


The construction. We show in Lemma 5 that for any sketchable pocket Q , we can construct a specific interior well-behaved sketch Λ_Q , and a specific well-behaved sketch Λ_Q^+ . Before we can present the proof, we first need to define Λ_Q and Λ_Q^+ .

If Q is a trivial pocket, that is, it consists of a single edge e_Q of P , we define

$$\Lambda_Q(v) = \begin{cases} p_i & \text{if } v = c_i, \\ p_j & \text{if } v = c_j, \\ e_Q & \text{otherwise.} \end{cases}$$

For non-trivial pockets Q (that do not consist of a single edge), we will define Λ_Q differently. This definition will rely on the definitions of Λ_L^+ and Λ_R^+ for the child pockets L and R of Q (whose outer triangles T_L^+ and T_R^+ equal the inner triangle T_Q of Q). Therefore, we postpone the definition of Λ_Q for non-trivial pockets until after the definition of Λ_Q^+ . Intuitively, Λ_Q^+ pushes those vertices which can be placed anywhere on e_Q out to T_Q^+ if their neighbors allow this and it pushes all remaining vertices as “far out as possible”. Formally, Λ_Q^+ is defined in terms of Λ_Q and will hence be defined if and only if Λ_Q is defined:

$$\Lambda_Q^+(v) = \begin{cases} T_Q^+ & \text{if } e_Q \subseteq \Lambda_Q(v) \text{ and } \forall_{(u,v) \in E} \Lambda_Q(u) \cap e_Q \neq \emptyset, \\ \Lambda_Q(v) \cap e_Q & \text{otherwise, if } \Lambda_Q(v) \cap e_Q \neq \emptyset, \\ \Lambda_Q(v) & \text{otherwise.} \end{cases}$$



■ **Figure 5** Λ_L^+ and Λ_R^+ are merged into Λ_Q which is then transformed into Λ_Q^+ for a subgraph of G . Vertex c is constrained to p_k in Λ_L^+ and can lie anywhere in T_R^+ in Λ_R^+ , hence c is constrained to p_k in both Λ_Q and Λ_Q^+ . Vertex a can lie anywhere in $T_L^+ = T_R^+ = T_Q$ in Λ_Q ; since a is connected to b it is pushed to the edge (p_i, p_j) in Λ_Q^+ (and not further). Vertex d can also lie anywhere in $T_L^+ = T_R^+ = T_Q$ in Λ_Q ; since it has no further restrictions it is pushed all the way to T_Q^+ in Λ_Q^+ .

It remains to define Λ_Q for non-trivial pockets. We will define Λ_Q only if Λ_L^+ and Λ_R^+ are defined for both of its child pockets L and R . We attempt to combine Λ_L^+ and Λ_R^+ into a sketch Λ_Q by taking the more restrictive placement for each vertex; here an assignment to T_L^+ or T_R^+ is interpreted as “no placement restriction” (see Figure 5). Potentially, Λ_L^+ and Λ_R^+ restrict the location of a vertex v in such a way that there is no valid placement for v . In such cases, the following definition assigns that vertex to an “undefined” location.

$$\Lambda_Q(v) = \begin{cases} \Lambda_L^+(v) \cap \Lambda_R^+(v) & \text{if } \Lambda_L^+(v) \cap \Lambda_R^+(v) \neq \emptyset, \\ \Lambda_L^+(v) & \text{otherwise, if } T_Q = \Lambda_R^+(v), \\ \Lambda_R^+(v) & \text{otherwise, if } T_Q = \Lambda_L^+(v), \\ \text{undefined} & \text{otherwise.} \end{cases}$$

If the above equation assigns any vertex to an “undefined” location, we say that Λ_Q is *undefined*. In summary, Λ_Q is defined if and only if Λ_L^+ is defined, Λ_R^+ is defined, and the above equation does not assign any vertex to an undefined location. We inductively show that Λ_Q and Λ_Q^+ are defined if and only if the pocket Q is sketchable. Moreover, if they are defined, then Λ_Q and Λ_Q^+ are interior well-behaved and well-behaved sketches, respectively. Lemma 4 shows that if both Λ_Q and Λ_Q^+ are defined, then they are sketches. Lemma 5 shows that if pocket Q is sketchable, then Λ_Q and Λ_Q^+ are both defined and (interior) well-behaved.

We try to construct a sketch Δ for the root triangle T_{root} , which (similar to Λ_Q for a non-trivial pocket Q) combines well-behaved sketches for its child pockets. Where a non-trivial pocket Q has two child pockets, T_{root} has three child pockets A , B , and C (with $T_A^+ = T_B^+ = T_C^+ = T_{\text{root}}$). The equation for Δ is analogous to that of Λ_Q ; we say that Δ is defined if and only if all of Λ_A^+ , Λ_B^+ , and Λ_C^+ are defined, and the following equation does not assign any vertex to an “undefined” location.

$$\Delta(v) = \begin{cases} \Lambda_A^+(v) \cap \Lambda_B^+(v) \cap \Lambda_C^+(v) & \text{if } \Lambda_A^+(v) \cap \Lambda_B^+(v) \cap \Lambda_C^+(v) \neq \emptyset, \\ \Lambda_A^+(v) & \text{otherwise, if } \Lambda_B^+(v) = \Lambda_C^+(v) = T_{\text{root}}, \\ \Lambda_B^+(v) & \text{otherwise, if } \Lambda_A^+(v) = \Lambda_C^+(v) = T_{\text{root}}, \\ \Lambda_C^+(v) & \text{otherwise, if } \Lambda_A^+(v) = \Lambda_B^+(v) = T_{\text{root}}, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Lemma 6 shows that the triangulation T does not admit a sketch if Δ is undefined. Otherwise, Lemma 4 shows that Δ is a sketch for the triangulation.

► **Lemma 4.** *The functions Λ_Q , Λ_Q^+ and Δ are sketches whenever they are defined.*

Proof. Λ_Q , Λ_Q^+ and Δ are sketches if neighboring vertices are assigned to simplices of a common triangle and c_k is assigned to p_k for all k ; with $i \leq k \leq j$ in the case of Λ_Q and Λ_Q^+ with $Q = Q_{(p_i, p_j)}$. We prove that Λ_Q and Λ_Q^+ are sketches by structural induction.

First consider Λ_Q for a trivial pocket Q . Λ_Q assigns all vertices to simplices of the triangle containing the edge (p_i, p_j) , and assigns $\Lambda_Q(c_i) = p_i$ and $\Lambda_Q(c_j) = p_j$, so Λ_Q is a sketch.

Next, assume that Λ_Q^+ is defined for a pocket Q . By induction, Λ_Q is defined and a sketch. If Λ_Q assigns c_k to p_k , so does Λ_Q^+ . If $\Lambda_Q^+(v) = T_Q^+$, all neighbors of v are by definition assigned to simplices of T_Q^+ . If $\Lambda_Q^+(v) \neq T_Q^+$, then $\Lambda_Q(v) \supseteq \Lambda_Q^+(v) \neq \emptyset$, so if Λ_Q assigns neighboring vertices to simplices of a common triangle, so does Λ_Q^+ . So Λ_Q^+ is a sketch.

Next, assume that Λ_Q is defined for a non-trivial pocket Q . Then Λ_L^+ and Λ_R^+ are sketches for the subpockets L and R of Q with $T_L^+ = T_R^+ = T_Q$. Λ_Q assigns c_k to p_k for all $i \leq k \leq j$. Suppose for a contradiction that Λ_Q assigns two neighboring vertices to simplices that do not share a triangle. Since both Λ_L^+ and Λ_R^+ assign neighboring vertices to simplices of common triangles, there are neighboring vertices u and v such that $\Lambda_Q(u) \not\subseteq \Lambda_L^+(u)$ and $\Lambda_Q(v) \not\subseteq \Lambda_R^+(v)$. So by definition of Λ_Q , we have $\Lambda_L^+(v) \cap \Lambda_R^+(v) = \emptyset$ and $\Lambda_L^+(u) = \Lambda_R^+(v) = T_Q$. Because Λ_L^+ assigns u and v to a common triangle, we have $\Lambda_L^+(u) \cap \Lambda_L^+(v) = T_Q \cap \Lambda_L^+(v) \neq \emptyset$, contradicting that $\Lambda_L^+(v) \cap T_Q = \Lambda_L^+(v) \cap \Lambda_R^+(v) = \emptyset$. Hence Λ_Q is a sketch.

An analogous argument shows that Δ is a sketch if for pockets A , B and C with $T_A^+ = T_B^+ = T_C^+ = T_{\text{root}}$, each of Λ_A^+ , Λ_B^+ and Λ_C^+ are sketches, and for all vertices v , we have $T_{\text{root}} \subseteq (\Lambda_B^+(v) \cap \Lambda_C^+(v)) \cup (\Lambda_A^+(v) \cap \Lambda_C^+(v)) \cup (\Lambda_A^+(v) \cap \Lambda_B^+(v))$ or $\Lambda_A^+(v) \cap \Lambda_B^+(v) \cap \Lambda_C^+(v) \neq \emptyset$. \blacktriangleleft

By the following lemma, Λ_Q and Λ_Q^+ are defined if and only if Q has a sketch.

► Lemma 5. *If a pocket Q is sketchable, then Λ_Q is defined and interior well-behaved, and Λ_Q^+ is defined and well-behaved.*

Proof. We prove this by structural induction along the definitions of Λ_Q and Λ_Q^+ .

For the base case, consider Λ_Q for a trivial pocket Q with lid $e_Q = (p_i, p_{i+1})$. As Λ_Q is defined unconditionally, it is a sketch by Lemma 4. Observe that for any interior local sketch Γ of Q , we have $\Gamma(c_i) = p_i$ and $\Gamma(c_{i+1}) = p_{i+1}$. For all vertices $v \notin \{c_i, c_j\}$, we have $\Gamma(v) \cap e_Q \subseteq \Lambda_Q(v)$, so $\Gamma \preceq_Q \Lambda_Q$. That is, Λ_Q is interior well-behaved.

For the inductive step of Λ_Q^+ , consider a (not necessarily trivial) sketchable pocket Q with lid $e_Q = (p_i, p_j)$. By induction we may assume that Λ_Q is defined and interior well-behaved. Therefore Λ_Q^+ is defined and a sketch (by Lemma 4). It remains to show that Λ_Q^+ is well-behaved, so for a contradiction suppose that it is not. Then there exists a local sketch Γ of Q and a vertex v for which $\Gamma(v) \cap T_Q^+ \not\subseteq \Lambda_Q^+(v)$. Because Λ_Q does not assign any vertex to simplices outside Q , we have $\Lambda_Q(v) \subseteq Q$, and by definition of Λ_Q^+ , we have $\Lambda_Q(v) \cap e_Q \subseteq \Lambda_Q^+(v) \cap T_Q^+$. By interior well-behavedness of Λ_Q and assumption that $\Gamma(v) \cap T_Q^+ \not\subseteq \Lambda_Q^+(v)$, we have $\Gamma(v) \not\subseteq Q$. Therefore, $\Gamma(v) = T_Q^+$ and for all $(u, v) \in E$ we have $\Gamma(u) \cap e_Q \neq \emptyset$. By interior well-behavedness of Λ_Q , we have $\Gamma(v) \cap e_Q = e_Q \subseteq \Lambda_Q(v)$, and for all $(u, v) \in E$, that $\emptyset \neq \Gamma(u) \cap e_Q \subseteq \Lambda_Q(u)$ and hence $\Lambda_Q(u) \cap e_Q \neq \emptyset$. So by definition we have $\Lambda_Q^+(v) = T_Q^+$, contradicting that $\Gamma(v) \cap T_Q^+ \not\subseteq \Lambda_Q^+(v)$, so Λ_Q^+ is well-behaved.

For the inductive step of Λ_Q , suppose that Q is a non-trivial sketchable pocket with lid $e_Q = (p_i, p_j)$. Since Q is sketchable, so are the child pockets L and R of Q (with $T_L^+ = T_R^+ = T_Q$). Inductively, we may assume that Λ_L^+ and Λ_R^+ are well-behaved sketches for L and R . Let Γ_L be the local sketch for L obtained from Γ by replacing $\Gamma(v)$ by $T_L^+ = T_Q$ whenever $\Gamma(v) \not\subseteq L$. Define Γ_R to be the analogous local sketch for R . For the sake of contradiction, assume that Λ_Q is not defined. Then there is some vertex v for which $\Lambda_L^+(v) \cap \Lambda_R^+(v) = \emptyset$ and $T_Q \not\subseteq \Lambda_L^+(v) \cup \Lambda_R^+(v)$. Because $T_Q \not\subseteq \Lambda_L^+(v) \cup \Lambda_R^+(v)$, we have $T_Q \not\subseteq \Gamma_L(v) \cup \Gamma_R(v)$, and therefore $\Gamma(v) \subseteq L$ and $\Gamma(v) \subseteq R$. So $\Gamma(v) \subseteq L \cap R = p_m$, where p_m is the vertex of T_Q not

on the lid of Q . Since $\Gamma(v) \neq \emptyset$, well-behavedness of Λ_L^+ and Λ_R^+ tells us that $p_m \in \Lambda_L^+(v)$ and $p_m \in \Lambda_R^+(v)$. But then $\Lambda_L^+(v) \cap \Lambda_R^+(v) \neq \emptyset$, which is a contradiction, so Λ_Q is defined. To show that Λ_Q is also interior well-behaved, we show that $\Gamma(v) \cap e_Q \subseteq \Lambda_Q(v)$ for all $v \in V$. By well-behavedness of Λ_L^+ and Λ_R^+ , we have $\Gamma(v) \cap e_Q \subseteq \Lambda_L^+(v)$ and $\Gamma(v) \cap e_Q \subseteq \Lambda_R^+(v)$. So $\Gamma(v) \cap e_Q \subseteq \Lambda_L^+(v) \cap \Lambda_R^+(v)$, and by definition of Λ_Q we have $\Lambda_L^+(v) \cap \Lambda_R^+(v) \subseteq \Lambda_Q(v)$, and hence $\Gamma(v) \cap e_Q \subseteq \Lambda_Q(v)$. So Λ_Q is interior well-behaved. ◀

Similarly, we can show for a given \mathcal{T} of a polygon P , that if \mathcal{T} has a sketch, then Δ is defined and a well-behaved sketch.

► **Lemma 6.** *If there exists a sketch for a given triangulation \mathcal{T} of a simple polygon P , then Δ is defined and well-behaved.*

The following two corollaries summarize that the existence of a sketch is equivalent to the existence of a well-behaved sketch, both for pockets Q and for a complete triangulation \mathcal{T} .

► **Corollary 7.** *Λ_Q and Λ_Q^+ are defined if and only if Q has a sketch.*

► **Corollary 8.** *Sketch Δ is defined if and only if the triangulation \mathcal{T} has a sketch.*


Computing triangulation-respecting drawings. Any sketch implies a drawing that places vertices anywhere in their assigned simplex. Conversely, any triangulation-respecting drawing implies a sketch that assigns vertices to the corresponding simplex. The definition of Δ hence directly results in a $\mathcal{O}(t|V||E|)$ -time algorithm both to decide the existence of, and to compute a triangulation-respecting drawing. We can improve this running time to linear.

► **Theorem 9.** *There is a linear-time algorithm to decide if (G, C) has a triangulation-respecting drawing for a simple polygon P with fixed triangulation \mathcal{T} ; the same algorithm also constructs a drawing if one exists.*

4 Planar triangulation-respecting drawings

We are given the same input as in Section 3, namely an instance (G, C) consisting of a graph G with n vertices and a cycle C with t vertices, and a simple polygon P with t vertices together with an arbitrary triangulation \mathcal{T} of P . In addition, we assume that the instance (G, C) is planar, that is, G has a planar drawing \mathcal{D} with C on the outer face. Note that \mathcal{D} does not necessarily map vertices of C to vertices of P .

Analogously to Section 3, we can ask the following question: is there a planar drawing of (G, C) that respects both P and \mathcal{T} ? The answer to this question is often “no”, even when both triangulation-respecting drawings and planar polygon-respecting drawings exist. Consider, for example, Figure 6: a planar triangulation-respecting drawing for this combination



■ **Figure 6** Left: A planar instance. Center: a triangulation-respecting drawing in which two vertices coincide. Right: a perturbed drawing that is planar but not triangulation-respecting.


of (G, C) , P , and \mathcal{T} does not exist; any drawing inside P either places two vertices on top of each other, or edges cross edges of the triangulation. Still, triangulation-respecting drawings are a useful tool for our final goal of constructing planar polygon-respecting drawings. The triangulation-respecting drawing of Figure 6 can be perturbed infinitesimally to obtain a planar polygon-respecting drawing (that is not triangulation-respecting). In this section, we show that if a planar instance (G, C) has a triangulation-respecting drawing, then it also has a *weakly-planar triangulation-respecting* one: a triangulation-respecting drawing that is planar and polygon-respecting after infinitesimal perturbation (moving vertices to simplices of \mathcal{T} that contain their original location.) Hence, the algorithm in Section 3 can decide for a planar instance (G, C) whether there is a weakly-planar triangulation-respecting drawing.

Let \mathcal{D} be a planar drawing of (G, C) . We call the triple (G, C, \mathcal{D}) a *plane instance*. We say that a weakly-planar triangulation-respecting drawing \mathcal{W} *accommodates* $(\mathcal{D}, \mathcal{T})$ if there exists a planar polygon-respecting infinitesimal perturbation $\widetilde{\mathcal{W}}$ of \mathcal{W} that is isotopic to \mathcal{D} in the plane. That is, one can be continuously deformed into the other without introducing self-intersections. We now construct a weakly-planar triangulation-respecting drawing that accommodates $(\mathcal{D}, \mathcal{T})$. A plane instance (G, C, \mathcal{D}) is *sketchable* if (G, C) has a sketch (for \mathcal{T}). Recall that a sketch does not have a notion of planarity. However, we show in Theorem 11 that any sketchable plane instance (G, C, \mathcal{D}) has a drawing \mathcal{W} which accommodates $(\mathcal{D}, \mathcal{T})$.

Minimal plane instances. We show how to transform any plane instance (G, C, \mathcal{D}) into a *minimal* plane instance, while preserving its sketchability; details can be found in the full version. First of all, we carefully triangulate (G, C, \mathcal{D}) , so as not to influence sketchability. If all faces of \mathcal{D} interior to C are triangles, we call (G, C, \mathcal{D}) a *triangulated instance*. A triangulation of a plane instance (G, C, \mathcal{D}) is a triangulated instance (G', C, \mathcal{D}') such that G' is a supergraph of G (with potentially additional vertices) and \mathcal{D} is the restriction of \mathcal{D}' to G . Second, we remove the interior of all separating triangles. If G does not have any separating triangles, then we contract any edge not on C that preserves sketchability and remove the interior of any separating triangles this edge contraction might create. If no further simplifications are possible, we call (G, C, \mathcal{D}) *minimal*.

► **Lemma 10.** *Every sketchable minimal plane instance (G, C, \mathcal{D}) has a drawing that accommodates $(\mathcal{D}, \mathcal{T})$.*

Proof. Consider a pocket $Q = Q_{(p_i, p_j)}$. We claim that if $\Lambda_Q^+(v) = p_k$ for some $k \in [i, j]$, then $v = c_k$, and moreover that if $\Lambda_Q^+(v) = e_Q$, then Q does not consist of a single edge and v is a neighbor of c_m , where p_m is the third vertex of T_Q . If Q consists of a single edge, then the claim clearly holds. If Q does not consist of a single edge, we have by induction that the claim holds for the subpockets $L = Q_{(p_i, p_m)}$ and $R = Q_{(p_m, p_j)}$. If $\Lambda_L^+(v) = e_L$ for some v , then we claim that the instance is not minimal. Let p_l be the third vertex of T_L , and without loss of generality assume that v is the most counter-clockwise neighbor (according to \mathcal{D}) of c_l for which $\Lambda_L^+(v) = e_L$ (see Figure 7). Then the (triangular) face counter-clockwise of



■ **Figure 7** The vertex u must be assigned to p_i or p_m . In both cases an edge can be contracted while preserving the existence of a sketch.

edge (c_l, v) is a triangle whose third vertex u does not have $\Lambda_L^+(u) = e_L$. Since u has c_l and v as neighbors, $\Lambda_L^+(u)$ is a simplex of T_L , but not e_L , so by definition of Λ_L^+ it is a vertex of T_L . By induction, u is therefore c_m , c_l , or c_i . Because u is a neighbor of c_l , u itself is not c_l . We argue that we can contract an edge while preserving sketchability, contradicting minimality.

First consider the case where $u = c_m$. Then (c_l, c_m) divides G into two subgraphs, to the left and to the right of (c_l, c_m) . We obtain a new sketch by reassigning all vertices of the right subgraph that are placed outside the pocket with lid (p_l, p_m) to that lid. By construction, the triangle c_l, c_m, v lies right of (c_l, c_m) . Then v and its neighbors are reassigned to p_l, p_m , or (p_l, p_m) , so contracting the edge (u, v) maintains sketchability, contradicting minimality.

Now consider the case where $u = c_i$. If $\Lambda_Q^+(u) = \Lambda_Q^+(v)$, it is clear that we can contract the edge and preserve a sketch. So since $\Lambda_Q^+(v)$ is either p_i or p_m , we have $\Lambda_Q^+(v) = p_m$, but then $\Lambda_R^+(v)$ is not T_Q , so either $\Lambda_R^+(v) = p_m$ or $\Lambda_R^+(v) = e_R$. The first case implies $v = c_m$ and hence contradicts $\Lambda_L^+(v) = e_L$. In the second case, the inductive hypothesis implies that v is a neighbor of the third vertex p_r of the triangle T_R . Now consider the subgraph of G that is right of the path consisting of the edges (p_l, v) and (v, p_r) . Any of its vertices that is assigned outside L and R can be assigned to p_m , maintaining a sketch. There exists a path from c_m to v that avoids c_l and c_r , and the last edge of this path can be contracted, contradicting minimality. So $\Lambda_L^+(v) \neq e_L$ and symmetrically $\Lambda_R^+(v) \neq e_R$. By definition, Λ_Q^+ assigns (for $k \in [i, j]$) only c_k to p_k , and only neighbors of c_m to e_Q , so the claim holds.

Therefore, in the sketch Δ , all vertices other than those of C are assigned to edges of T_{root} , or T_{root} itself. If there is such a vertex not on C , then contracting an edge between C and $G \setminus C$ yields an instance with a sketch, contradicting minimality. So all vertices in a minimal instance lie on C . Because G is triangulated, this means that its edges coincide with those of the triangulation of P . So the instance clearly has an accommodating drawing. ◀

The proof of Theorem 11 shows that the accommodating drawing of the minimal instance obtained from the simplification procedure (if that instance is sketchable) can be extended to be an accommodating drawing for the original instance, by undoing the simplification steps.


► **Theorem 11.** *A plane instance (G, C, \mathcal{D}) has a drawing that accommodates $(\mathcal{D}, \mathcal{T})$ if and only if (G, C, \mathcal{D}) is sketchable.*

The algorithm implied by Theorem 9 can check in linear time if a plane instance (G, C, \mathcal{D}) has a sketch and via Theorem 11 the same algorithm can decide in linear time if (G, C, \mathcal{D}) has an accommodating drawing. This drawing can be constructed in polynomial time, following the (polynomial number of) steps in the minimization procedure.

5 Sufficient conditions for polygon-universality

In Section 2 we proved that the Pair and Triple Conditions are necessary for an instance (G, C) to be polygon-universal. Here we show, using triangulation-respecting drawings, that these two conditions are sufficient as well. In Sections 3 and 4 we argued that an instance (G, C) has a triangulation-respecting drawing for a triangulation \mathcal{T} of P if and only if it has a sketch for \mathcal{T} ; we also gave an algorithm that tests whether such a sketch exists. Below we show that if the Pair and Triple Conditions are satisfied for an instance (G, C) then it has a sketch for any triangulation \mathcal{T} . We do so by examining the testing algorithm more closely.

We first show that the Pair Condition alone already implies that each pocket has a sketch. The Triple Condition then allows us to combine sketches at the root T_{root} of \mathcal{T} . Denote by $Q_e = Q_{(p_i, p_j)}$ the pocket with lid $e = (p_i, p_j)$. We want to determine whether an individual vertex can be drawn outside a pocket. Definition 12 relates the position of a vertex in a sketch of a pocket to its distance to points on the cycle, see Figure 8.



■ **Figure 8** Cases for Definition 12. Left: we have $\pi \leq a$ and $\pi \leq b$. Right: we have $\pi_1 \leq a_1 + 1$, $\pi_1 \leq b_1 - 1$, $\pi_2 \leq a_2 - 1$, and $\pi_2 \leq a_2 + 1$.

► **Definition 12.** Vertex v is pulled by pocket $Q_{(p_i, p_j)}$ if and only if either of the following two conditions hold:

1. for some $k \in \{i, \dots, j\}$, we have $d_G(v, c_k) \leq \min(k - i, j - k)$;
2. for some $k, l \in \{i, \dots, j\}$, we have $d_G(v, c_k) \leq \min(k - i + 1, j - k - 1)$ and $d_G(v, c_l) \leq \min(l - i - 1, j - l + 1)$.

Let Q be a pocket and let v be an arbitrary vertex. Either the well-behaved sketch Λ_Q^+ of Q places v outside of Q or Lemma 13 characterizes where in Q vertex v “is stuck”.

► **Lemma 13.** Let v be a vertex and $Q = Q_{(p_i, p_j)}$ with $i < j < i + t$ be a sketchable pocket.

1. If $p_j \notin \Lambda_Q^+(v)$, then $v = c_{j-1}$ or there exists a triangle in Q with vertices p_a, p_b, p_j and $i \leq a < b < j$ such that for pocket $Q' = Q_{(p_a, p_b)}$, $\Lambda_{Q'}^+(v) \neq T_{Q'}^+$.
2. If $p_i \notin \Lambda_Q^+(v)$, then $v = c_{i+1}$ or there exists a triangle in Q with vertices p_i, p_a, p_b and $i < a < b \leq j$ such that for pocket $Q' = Q_{(p_a, p_b)}$, $\Lambda_{Q'}^+(v) \neq T_{Q'}^+$.

Proof. Statements (1) and (2) can be proved using symmetric arguments, so we prove only statement (1). We proceed by induction on the size of Q . If Q consists of the single edge (p_i, p_j) , then $i = j - 1$, and if $p_j \notin \Lambda_Q^+(v)$, then $v = c_i = c_{j-1}$. So assume that Q does not consist of a single edge. Let p_m with $i < m < j$ be the third vertex of T_Q and let $L = Q_{(p_i, p_m)}$ and $R = Q_{(p_m, p_j)}$. If $p_j \notin \Lambda_Q^+(v)$, then $\Lambda_L^+(v)$ or $\Lambda_R^+(v)$ does not contain p_j . If $p_j \notin \Lambda_R^+(v)$, then by induction we are done. So assume that $p_j \in \Lambda_R^+(v)$ and hence that $p_j \notin \Lambda_L^+(v)$. This means that $\Lambda_L^+(v) \neq T_L^+$, completing the proof. ◀

Lemma 14 and 15 relate the characterization in Lemma 13, which uses the well-behaved sketch Λ_Q^+ , to the requirements on graph and cycle distances expressed in Definition 12. Lemma 14 covers the first condition of Definition 12, while Lemma 15 covers the remainder.

► **Lemma 14.** Assume that pocket $Q = Q_{(p_i, p_j)}$ with $i < j < i + t$ admits a sketch. If $\Lambda_Q^+(v) \neq T_Q^+$, then v is pulled by Q .

Proof. If $\Lambda_Q^+(v) \neq T_Q^+$, then either $e_Q \not\subseteq \Lambda_Q(v)$ or for some neighbor u of v , $\Lambda_Q(u)$ does not intersect e_Q . We proceed by induction on the size of Q .

Q is trivial. If Q is trivial, it consists of one edge e_Q , and $\Lambda_Q(u)$ intersects e_Q for all u . By assumption that $\Lambda_Q^+(v) \neq T_Q^+$, we have $e_Q \not\subseteq \Lambda_Q(v)$ by definition of Λ_Q^+ , so $v = c_i$ or $v = c_j$. Hence, $d_G(v, c_k) = 0 \leq \min\{k - i, j - k\} = 0$ for some $k \in \{i, j\}$, so v is pulled by Q .

Q is non-trivial. Next, suppose that Q is non-trivial, and thus contains the triangle T_Q . Let p_m (with $i < m < j$) be the third vertex of T_Q and let $L = Q_{(p_i, p_m)}$ and $R = Q_{(p_m, p_j)}$ be the two subpockets of Q with $T_L^+ = T_R^+ = T_Q$. If $e_Q \not\subseteq \Lambda_Q(v)$, then $T_Q \neq \Lambda_L^+(v)$

or $T_Q \neq \Lambda_R^+(v)$, so by induction v is pulled by L or R , and hence also by Q . So assume that $e_Q \subseteq \Lambda_Q(v)$ and there exists some neighbor u of v for which $\Lambda_Q(u)$ does not intersect e_Q . It follows by construction that $\Lambda_L^+(v) = \Lambda_R^+(v) = T_Q$. Because u is assigned to a simplex of the same triangle as v , we have $\Lambda_L^+(u) \subseteq T_Q$ and $\Lambda_R^+(u) \subseteq T_Q$, so $\Lambda_Q(u) \subseteq T_Q$. Since $\Lambda_Q(u)$ does not intersect e_Q , we have $\Lambda_Q(u) = p_m$ and hence $\Lambda_L^+(u) \cap \Lambda_R^+(u) = p_m$. So either $\Lambda_L^+(u) = e_L$ and $\Lambda_R^+(u) = e_R$, or $\Lambda_L^+(u)$ or $\Lambda_R^+(u)$ is p_m . We consider these cases separately.

Edge case. Suppose that $\Lambda_L^+(u) = e_L$ and $\Lambda_R^+(u) = e_R$. Then $m \in [i + 2, j - 2]$ and by induction u is pulled by both L and R . We distinguish three cases depending on what causes u to be pulled by L and R , and show in each case that v is pulled by Q .

1. If there exists some $l \in [i, m]$ with $d_G(u, c_l) \leq \min(l - i - 1, m - l + 1)$, then

$$\begin{aligned} d_G(v, c_l) &\leq d_G(u, c_l) + 1 \\ &\leq \min(l - i - 1, m - l + 1) + 1 \\ &\leq \min(l - i, j - l), \end{aligned}$$

so v is pulled by Q .

2. Symmetrically, v is pulled by Q if $d_G(u, c_k) \leq \min(k - m + 1, j - k - 1)$ for some $k \in [m, j]$.
3. In the remaining case, there exist $k \in [i, m]$ and $l \in [m, j]$ with $d_G(u, c_k) \leq \min(k - i, m - k)$ and $d_G(u, c_l) \leq \min(l - m, j - l)$. Therefore

$$\begin{aligned} d_G(v, c_k) &\leq d_G(u, c_k) + 1 \\ &\leq \min(k - i, m - k) + 1 \\ &\leq \min(k - i + 1, m - k + 1) \\ &\leq \min(k - i + 1, j - k - 1) \end{aligned}$$

and symmetrically $d_G(v, c_l) \leq \min(l - i - 1, j - l + 1)$, so v is pulled by Q .

Corner case. Assume that $\Lambda_L^+(u) = p_m$ (the case $\Lambda_R^+(u) = p_m$ is symmetric). Then $p_i \notin \Lambda_L^+(u)$, so by Lemma 13, we have either $u = c_{i+1}$ or there exists some pocket $Q' = Q_{(p_a, p_b)}$ with $i < a < b \leq m$ such that $\Lambda_{Q'}^+(u) \neq T_{Q'}$. If $u = c_{i+1}$, then for $k = i + 1$ we have $d_G(v, c_k) \leq d_G(u, c_k) + 1 \leq 1 \leq \min(k - i, j - k)$ in which case v is pulled by Q . Otherwise, u is by induction pulled by some pocket $Q_{(p_a, p_b)}$ with $i < a < b \leq m < j$, and since $d_G(v, u) \leq 1$, the triangle inequality shows that v is pulled by $Q_{(p_i, p_j)}$.

By induction, v is pulled by Q whenever $\Lambda_Q^+(v) \neq T_Q^+$. ◀

► **Lemma 15.** Assume that pocket $Q = Q_{(p_i, p_j)}$ with $i < j < i + t$ admits a sketch. If v is pulled by Q but there exists no $k \in [i, j]$ such that $d_G(v, c_k) \leq \min(k - i, j - k)$, then there exist $k, o, l \in [i, j]$ such that $k < o < l$ and a vertex $x \neq v$ with $d_G(x, c_k) \leq \min(k - i, o - k) \leq \min(k - i + 1, j - k - 1) - d_G(x, v)$ and $d_G(x, c_l) \leq \min(l - o, j - l) \leq \min(l - i - 1, j - l + 1) - d_G(x, v)$.

Lemma 16 ties together our preceding arguments to show that if a pocket has no sketch, then the Pair Condition is violated. This directly implies Corollary 17.

► **Lemma 16.** If pocket $Q = Q_{(p_i, p_j)}$ with $i < j < i + t$ has no sketch, then there exist $i \leq k \leq l \leq j$ such that $d_G(c_k, c_l) < d_C(c_k, c_l)$.

► **Corollary 17.** If the Pair Condition is satisfied, then all pockets have a sketch.

We now established that the Pair Condition implies that each pocket has a sketch. If additionally the Triple Condition is satisfied, then Theorem 18 shows that we can combine the sketches for the three pockets, whose lids are the edges of the root triangle T_{root} , to obtain a sketch, and hence a triangulation-respecting drawing, for (G, C) .

► **Theorem 18.** *If an instance (G, C) satisfies the Pair and Triple Conditions, then (G, C) has a triangulation-respecting drawing for any triangulation \mathcal{T} of any simple polygon P .*

► **Corollary 19.** *If a plane instance (G, C, \mathcal{D}) satisfies the Pair and Triple Conditions, then (G, C) has a drawing that accommodates $(\mathcal{D}, \mathcal{T})$ for any triangulation \mathcal{T} of any simple polygon P .*

6 Discussion and conclusion

We have characterized the (planar) polygon-universal graphs (G, C) by means of simple combinatorial conditions involving (graph-theoretic) distances along the cycle C and in the graph G . In particular, this shows that, even though the recognition of polygon-universal graphs most naturally lies in $\forall\exists\mathbb{R}$, it can in fact be tested in polynomial time, by explicitly checking the Pair and the Triple Conditions. Our main open question concerns the restriction to simple polygons without holes. Can a similar characterization be achieved in the presence of holes? Or is the polygon-universality problem for simple polygons with holes $\forall\exists\mathbb{R}$ -complete?

Another interesting question concerns the running time for recognizing polygon-universal graphs. Testing the Pair and Triple Conditions naively requires at least $\Omega(n^3)$ time. On the other hand, at least in the non-planar case, given (G, C) and a polygon P with arbitrary triangulation T , we can in linear time either find an extension or a violation of the Pair/Triple Condition, which shows that the instance is not polygon-universal. (Recall that a polygon-extension for P might exist, though not one that respects T , see Figure 2). For planar instances, the contraction to minimal instances causes an additional linear factor in the running time. Can (planar) polygon-universality be tested in $o(n^2)$ time?

References

- 1 Patrizio Angelini, Giuseppe Di Battista, Fabrizio Frati, Vít Jelínek, Jan Kratochvíl, Maurizio Patrignani, and Ignaz Rutter. Testing planarity of partially embedded graphs. *ACM Transactions on Algorithms*, 11(4):32:1–32:42, 2015. doi:10.1145/2629341.
- 2 Patrizio Angelini, Philipp Kindermann, Andre Löffler, Lena Schlipf, and Antonios Symvonis. One-bend drawings of outerplanar graphs inside simple polygons. In *Abstr. 36th European Workshop on Computational Geometry*, pages 70:1–70:6, 2020.
- 3 John Canny. Some algebraic and geometric computations in PSPACE. In *Proc. 20th Symposium on Theory of Computing*, page 460–467, 1988. doi:10.1145/62212.62257.
- 4 Erin W. Chambers, David Eppstein, Michael T. Goodrich, and Maarten Löffler. Drawing graphs in the plane with a prescribed outer face and polynomial area. *Journal of Graph Algorithms and Applications*, 16(2):243–259, 2012. doi:10.7155/jgaa.00257.
- 5 Michael G. Dobbins, Linda Kleist, Tillmann Miltzow, and Paweł Rzażewski. $\forall\exists\mathbb{R}$ -completeness and area-universality. In *Proc. 44th International Workshop on Graph-Theoretic Concepts in Computer Science*, LNCS 11159, pages 164–175, 2018. doi:10.1007/978-3-030-00256-5_14.
- 6 Christian A. Duncan, Michael T. Goodrich, and Stephen G. Kobourov. Planar drawings of higher-genus graphs. *Journal of Graph Algorithms and Applications*, 15(1):7–32, 2011. doi:10.1007/978-3-642-11805-0_7.
- 7 Seok-Hee Hong and Hiroshi Nagamochi. Convex drawings of graphs with non-convex boundary constraints. *Discrete Applied Mathematics*, 156(12):2368–2380, 2008. doi:10.1016/j.dam.2007.10.012.

- 8 Vít Jelínek, Jan Kratochvíl, and Ignaz Rutter. A Kuratowski-type theorem for planarity of partially embedded graphs. *Computational Geometry*, 46(4):466–492, 2013. doi:10.1016/j.comgeo.2012.07.005.
- 9 Anna Lubiw, Tillmann Miltzow, and Debajyoti Mondal. The complexity of drawing a graph in a polygonal region. In *Proc. 26th International Symposium on Graph Drawing and Network Visualization*, LNCS 11282, pages 387–401, 2018. doi:10.1007/978-3-030-04414-5_28.
- 10 Tamara Mchedlidze, Martin Nöllenburg, and Ignaz Rutter. Extending convex partial drawings of graphs. *Algorithmica*, 76(1):47–67, 2016. doi:10.1007/s00453-015-0018-6.
- 11 Tamara Mchedlidze and Jérôme Urhausen. β -stars or on extending a drawing of a connected subgraph. In *Proc. 26th International Symposium on Graph Drawing and Network Visualization*, LNCS 11282, pages 416–429, 2018. doi:10.1007/978-3-030-04414-5_30.
- 12 Marcus Schaefer and Daniel Štefankovič. Fixed points, nash equilibria, and the existential theory of the reals. *Theory of Computing Systems*, 60:172–193, 2017. doi:10.1007/s00224-015-9662-0.
- 13 William T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, 13(3):743–768, 1963. doi:10.1112/plms/s3-13.1.743.