Adjacency Graphs of Polyhedral Surfaces

Elena Arseneva
Saint Petersburg State University, Russia

Linda Kleist
Technische Universität Braunschweig, Germany

Boris Klemz
Universität Würzburg, Germany

Maarten Löffler
Utrecht University, The Netherlands

André Schulz
FernUniversität in Hagen, Germany

Birgit Vogtenhuber
Technische Universität Graz, Austria

Alexander Wolff
Universität Würzburg, Germany

Abstract

We study whether a given graph can be realized as an adjacency graph of the polygonal cells of a polyhedral surface in $\mathbb{R}^3$. We show that every graph is realizable as a polyhedral surface with arbitrary polygonal cells, and that this is not true if we require the cells to be convex. In particular, if the given graph contains K_5, $K_{5,81}$, or any nonplanar 3-tree as a subgraph, no such realization exists. On the other hand, all planar graphs, $K_{4,4}$, and $K_{3,5}$ can be realized with convex cells. The same holds for any subdivision of any graph where each edge is subdivided at least once, and, by a result from McMullen et al. (1983), for any hypercube.

Our results have implications on the maximum density of graphs describing polyhedral surfaces with convex cells: The realizability of hypercubes shows that the maximum number of edges over all realizable n-vertex graphs is in $\Omega(n \log n)$. From the non-realizability of $K_{5,81}$, we obtain that any realizable n-vertex graph has $\mathcal{O}(n^{9/5})$ edges. As such, these graphs can be considerably denser than planar graphs, but not arbitrarily dense.

2012 ACM Subject Classification Mathematics of computing → Graphs and surfaces; Mathematics of computing → Combinatoric problems

Keywords and phrases polyhedral complexes, realizability, contact representation

Digital Object Identifier 10.4230/LIPIcs.SoCG.2021.11

Related Version *Full Version*: https://arxiv.org/abs/2103.09803

Funding *Elena Arseneva*: partially supported by RFBR, project 20-01-00488.
Boris Klemz: supported by DFG project WO 758/11-1.
Birgit Vogtenhuber: partially supported by the Austrian Science Fund within the collaborative DACH project *Arrangements and Drawings* as FWF project I 3340-N35.

Acknowledgements We thank the organizers of Dagstuhl Seminar 19352 "Computation in Low-Dimensional Geometry and Topology" for bringing us together. We are particularly indebted to seminar participant Arnaud de Mesmay for asking a question that initiated our research. We also thank the anonymous referees of our EuroCG 2020 submission for their helpful comments.

1 Introduction

A *polyhedral surface* consists of a set of interior-disjoint polygons embedded in $\mathbb{R}^3$, where each edge may be shared by at most two polygons. Polyhedral surfaces have been long studied in computational geometry, and have well-established applications in for instance computer graphics [11] and geographical information science [8].

Inspired by those applications, classic work in this area often focuses on restricted cases, such as surfaces of (genus 0) polyhedra [3, 23], or x, y-monotone surfaces known as *polyhedral terrains* [7]. Such surfaces are, in a sense, 2-dimensional. One elegant way to capture this "essentially 2-dimensional behaviour" is to look at the adjacency graph (see below for a precise definition) of the surface: in both cases described above, this graph is *planar*. In fact, by Steinitz's Theorem the adjacency graphs of surfaces of convex polyhedra are exactly the 3-connected planar graphs [35]. If we allow the surface of a polyhedron to have a *boundary*, then every planar graph has a representation as such a polyhedral surface [12].

Recently, applications in computational topology have intensified the study of polyhedral surfaces of non-trivial topology. In sharp contrast to the simpler case above, where the classification is completely understood, little is known about the class of adjacency graphs that describe general polyhedral surfaces. In this paper we investigate this graph class.

Our model. A polyhedral surface $\mathcal{S} = \{S_1, \ldots, S_n\}$ is a set of n closed polygons embedded in $\mathbb{R}^3$ such that, for all pairwise distinct indices $i, j, k \in \{1, 2, \ldots, n\}$:

- S_i and S_j are interior-disjoint (w.r.t. the 2D relative interior of the objects);
- if $S_i \cap S_j \neq \emptyset$, then $S_i \cap S_j$ is either a single corner or a complete *side* of both S_i and S_j;
- if $S_i \cap S_j \cap S_k \neq \emptyset$ then it is a single corner (i.e., a side is shared by at most two polygons).

To avoid confusion with the corresponding graph elements, we consistently refer to polygon vertices as *corners* and to polygon edges as *sides*.

The *adjacency graph* of a polyhedral surface $\mathcal{S}$, denoted as $\mathcal{G}(\mathcal{S})$, is the graph whose vertices correspond to the polygons of $\mathcal{S}$ and which has an edge between two vertices if and only if the corresponding polygons of $\mathcal{S}$ share a side. Note that a corner–corner contact is allowed in our model but does not induce an edge in the adjacency graph. Further observe that the adjacency graph does not uniquely determine the topology of the surface. Fig. 1 shows an example of a polyhedral surface and its adjacency graph. We say that a polyhedral surface $\mathcal{S}$ *realizes* a graph G if $\mathcal{G}(\mathcal{S})$ is isomorphic to G. In this case, we write $\mathcal{G}(\mathcal{S}) \simeq G$.

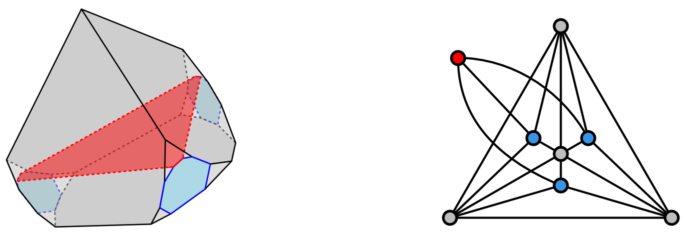

(a) A surface $\mathcal{S}$. **(b)** The graph $\mathcal{G}(\mathcal{S})$.

Figure 1 A convex-polyhedral surface $\mathcal{S}$ and its nonplanar 3-degenerate adjacency graph $\mathcal{G}(\mathcal{S})$.

If every polygon of a polyhedral surface $\mathcal{S}$ is strictly convex, we call $\mathcal{S}$ a *convex-polyhedral* surface. Our paper focuses on convex-polyhedral surfaces; refer to Fig. 2 for an example of a general (nonconvex) polyhedral surface. We emphasize that we do not require that every polygon side has to be shared with another polygon.
Our work relates to two lines of research: Steinitz-type problems and contact representations.

Steinitz-type problems. Steinitz's Theorem gives the positive answer to the *realizability problem* for convex polyhedra. This result is typically stated in terms of the realizabilty of a graph as the 1-skeleton of a convex polyhedron. Our perspective comes from the dual point of view, describing the adjacencies of the faces instead of the adjacencies of the vertices.

Steinitz's Theorem settles the problem raised in this paper for surfaces that are homeomorphic to a sphere. A slightly stronger version of Steinitz's Theorem by Grünbaum and Barnette [5] states that every planar 3-connected graph can be realized as the 1-skeleton of a convex polyhedron with the prescribed shape of one face. Consequently, also in our model we can prescribe the shape of one polygon if the adjacency graph of the surface is planar. For other classes of polyhedra only very few partial results for their graph-theoretic characterizations are known [13, 22]. No generalization for Steinitz's Theorem for surfaces of higher genus is known, and therefore there are also no results for the dual perspective. In higher dimensions, Richter-Gebert's Universality Theorem implies that the realizability problem for abstract 4-polytopes is $\exists\mathbb{R}$-complete [31].

The algorithmic problem of determining whether a given k-dimensional simplicial complex embeds in $\mathbb{R}^d$ is an active field of research [6, 17, 28, 30, 33, 34]. There exist at least three interesting notions of embeddability: linear, piecewise linear, and topological embeddability, which usually are not the same [28]. The case $(k, d) = (1, 2)$, however, corresponds to testing graph planarity, and thus, all three notions coincide, and the problem lies in P.

Contact representations. A realization of a graph as a polyhedral surface can be viewed as a *contact representation* of this graph with polygons in $\mathbb{R}^3$, where a contact between two polygons is realized by sharing an entire polygon side, and each side is shared by at most two polygons. In a general contact representation of a graph, the vertices are represented by interior-disjoint geometric objects, where two objects touch if and only if the corresponding vertices are adjacent. In concrete settings, the object type (disks, lines, polygons, etc.), the type of contact, and the embedding space is specified. Numerous results concerning which graphs admit a contact representation of some type are known; we review some of them.

The well-known Andreev–Koebe–Thurston circle packing theorem [2, 27] states that every planar graph admits a contact representation by touching disks in $\mathbb{R}^2$. A less known but impactful generalization by Schramm [32, Theorem 8.3] guarantees that every triangulation (i.e., maximal planar graph) has a contact representation in $\mathbb{R}^2$ where every inner vertex corresponds to a homothetic copy of a prescribed smooth convex set; the three outer vertices correspond to prescribed smooth arcs whose union is a simple closed curve. If the prototypes and the curve are polygonal, i.e., are not smooth, then there still exists a contact representation, however, the sets representing inner vertices may degenerate to points, which may lead to extra contacts. As observed by Gonçalves et al. [19], Schramm's result implies that every subgraph of a 4-connected triangulation has a contact representation with aligned equilateral triangles and similarly, every inner triangulation of a 4-gon without separating 3- and 4-cycles has a hole-free contact representation with squares [15].

While for the afore-mentioned existence results there are only iterative procedures that compute a series of representations converging to the desired one, there also exist a variety of shapes for which contact representations can be computed efficiently. Allowing for sides of one polygon to be *contained* in the side of adjacent polygons, Duncan et al. [12] showed that, in this model, every planar graph can be realized by hexagons in the plane and that hexagons are sometimes necessary. Assuming side–corner contacts, de Fraysseix et al. [10] showed that every plane graph has a triangle contact representation and how to compute one. Gansner et al. [18] presented linear-time algorithms for triangle side-contact representations for outerplanar graphs, square grid graphs, and hexagonal grid graphs. Kobourov et al. [26] showed that every 3-connected cubic planar graph admits a triangle side-contact representation whose triangles form a tiling of a triangle. For a survey of planar graphs that can be represented by dissections of a rectangle into rectangles, we refer to Felsner [15]. Alam et al. [1] presented a linear-time algorithm for hole-free contact representations of triangulations where each vertex is represented by a 10-sided rectilinear polygon of prescribed area.

Representations with one-dimensional objects in $\mathbb{R}^2$ have also been studied. While every plane bipartite graph has a contact representation with horizontal and vertical segments [9], recognizing segment contact graphs is an NP-complete problem even for planar graphs [20]. Hliněný showed that recognizing curve contact graphs, where no four curves meet in one point, is NP-complete for planar graphs and is solvable in polynomial time for planar triangulations.

Less is known about contact representations in higher dimensions. Every graph is the contact graph of interior-disjoint convex polytopes in $\mathbb{R}^3$ where contacts are shared 2-dimensional facets [37]. Hliněný and Kratochvíl [21] proved that the recognition of unit-ball contact graphs in $\mathbb{R}^d$ is NP-hard for $d = 3, 4$, and 8. Felsner and Francis [16] showed that every planar graph has a contact representation with axis-parallel cubes in $\mathbb{R}^3$. For proper side contacts, Kleist and Rahman [25] proved that every subgraph of an Archimedean grid can be represented with unit cubes, and every subgraph of a d-dimensional grid can be represented with d-cubes. Evans et al. [14] showed that every graph has a contact representation where vertices are represented by convex polygons in $\mathbb{R}^3$ and edges by shared corners of polygons, and gave polynomial-volume representations for bipartite, 1-planar, and cubic graphs.

Contribution and organization. We show that for every graph G there exists a polyhedral surface $\mathcal{S}$ such that G is the adjacency graph of $\mathcal{S}$; see Section 2. For convex-polyhedral surfaces, the situation is more intricate; see Section 3. Every planar graph can be realized by a *flat* convex-polyhedral surface (Proposition 3), i.e., a convex-polyhedral surface in $\mathbb{R}^2$. Some nonplanar graphs cannot be realized by convex-polyhedral surfaces in $\mathbb{R}^3$; in particular this holds for all supergraphs of K_5 (Proposition 5), of $K_{5,81}$ (Theorem 9), and of all nonplanar 3-trees (Theorem 14). Nevertheless, many nonplanar graphs, including $K_{4,4}$ and $K_{3,5}$, have such a realization (Propositions 7 and 8). We remark that all our positive results hold for subgraphs and subdivisions as well (Proposition 2). Similarly, our negative results carry over to supergraphs. For some proofs and additional figures, see the full version of this article [4].

Our results have implications on the maximum density of adjacency graphs of convex-polyhedral surfaces; see Section 4. While the non-realizability of $K_{5,81}$ implies that the number of edges of any realizable n-vertex graph is upperbounded by $\mathcal{O}(n^{9/5})$ edges, the realizability of hypercubes (Section 3.4) implies that it is in $\Omega(n \log n)$. Hence these graphs can be considerably denser than planar graphs, but not arbitrarily dense.

2 General Polyhedral Surfaces

We start with a positive result.

▶ **Proposition 1.** *For every graph G, there exists a polyhedral surface $\mathcal{S}$ such that $\mathcal{G}(\mathcal{S}) \simeq G$.*

Proof. We start our construction with $n = |V(G)|$ interior-disjoint rectangles such that there is a line segment s that acts as a common side of all these rectangles. We then cut away parts of each rectangle thereby turning it into a comb-shaped polygon as illustrated in Fig. 2. These polygons represent the vertices of G. For each pair (P, P') of polygons that are adjacent in G, there is a subsegment $s_{PP'}$ of s such that $s_{PP'}$ is a side of both P and P' that is disjoint from the remaining polygons. In particular, every polygon side is adjacent to at most two polygons. The result is a polyhedral surface whose adjacency graph is G. ◀

If we additionally insist that each polygon shares *all* of its sides with other polygons, the polyhedral surface describe a closed volume. In this model, K_7 can be realized as the Szilassi polyhedron; see Fig. 3. The tetrahedron and the Szilassi polyhedron are the only two known polyhedra in which each face shares a side with every other face [36]. Which other (complete) graphs can be realized in this way remains an open problem.

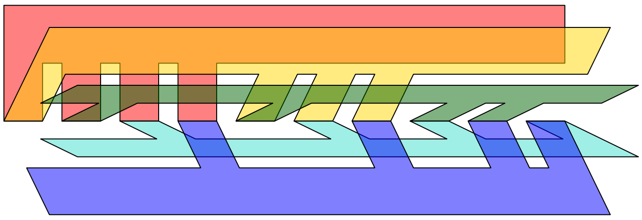

Figure 2 A realization of K_5 by arbitrary polygons with side contacts in $\mathbb{R}^3$.

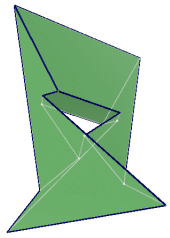

Figure 3 The Szilassi polyhedron realizes K_7 [36].

3 Convex-Polyhedral Surfaces

In this section we investigate which graphs can be realized by *convex*-polyhedral surfaces. First of all, it is always possible to represent a subgraph or a subdivision of an adjacency graph with slight modifications of the corresponding surface. While *trimming* the polygons allows to represent subgraphs, subdivision can be obtained by trimming and inserting chains of polygons. Consequently, we obtain the following result, which we prove formally in the full version of this article [4].

▶ **Proposition 2.** *The set of adjacency graphs of convex-polyhedral surfaces in $\mathbb{R}^3$ is closed under taking subgraphs and subdivisions.*

The existence of a flat surface with the correct adjacencies follows from the Andreev–Koebe–Thurston circle packing theorem; we include a direct proof in the full version [4].

▶ **Proposition 3.** *For every planar graph G, there exists a flat convex-polyhedral surface $\mathcal{S}$ such that $\mathcal{G}(\mathcal{S}) \simeq G$. Moreover, such a surface can be computed in linear time.*

So for planar graphs, corner and side contacts behave similarly. For nonplanar graphs (for which the third dimension is essential), the situation is different. Here, side contacts are more restrictive.

3.1 Complete Graphs

We introduce the following notation. In a polyhedral surface $\mathcal{S}$ with adjacency graph G, we denote by P_v the polygon in $\mathcal{S}$ that represents vertex v of G.

▶ **Lemma 4.** *Let $\mathcal{S}$ be a convex-polyhedral surface in $\mathbb{R}^3$ with adjacency graph G. If G contains a triangle uvw, polygons P_v and P_w lie in the same closed halfspace w.r.t. P_u.*

Proof. Due to their convexity, each of P_v and P_w lie entirely in one of the closed halfspaces with respect to the supporting plane of P_u. Moreover, one of the halfspaces contains both P_v and P_w; otherwise they cannot share a side and the edge vw would not be represented. ◀

A graph H is *subisomorphic* to a graph G if G contains a subgraph G' with $H \simeq G'$.

▶ **Proposition 5.** *There exists no convex-polyhedral surface $\mathcal{S}$ in $\mathbb{R}^3$ such that K_5 is subisomorphic to $\mathcal{G}(\mathcal{S})$.*

Proof. Suppose that there is a convex-polyhedral surface $\mathcal{S}$ with $\mathcal{G}(\mathcal{S}) \simeq K_5$. By Lemma 4 and the fact that all vertex triples form a triangle, the surface $\mathcal{S}$ lies in one closed halfspace of the supporting plane of every polygon P of $\mathcal{S}$. In other words, $\mathcal{S}$ is a subcomplex of a (weakly) convex polyhedron, whose adjacency graph must be planar. This yields a contradiction to the nonplanarity of K_5. Together with Proposition 2 this implies the claim. ◀

Evans et al. [14] showed that every bipartite graph has a contact representation by touching polygons on a polynomial-size integer grid in $\mathbb{R}^3$ for the case of corner contacts. As we have seen before, side contacts are less flexible. In particular, in Theorem 9 we show that $K_{5,81}$ cannot be represented. On the positive side, we show in the following that those bipartite graphs that come from subdividing edges of arbitrary graphs can be realized. In our construction, we place the polygons in a cylindrical fashion, which is reminiscent to the realizations created by Evans et al. However, due to the more restrictive nature of side contacts, the details of the two approaches are necessarily quite different.

▶ **Theorem 6.** *Let G be any graph, and let G' be the subdivision of G in which every edge is subdivided with (at least) one vertex. Then there exists a convex-polyhedral surface $\mathcal{S}$ in $\mathbb{R}^3$ such that $\mathcal{G}(\mathcal{S}) \simeq G'$.*

Proof. Let $V(G) = \{v_1, \ldots, v_n\}$, let $E(G) = \{e_1, \ldots, e_m\}$, and let P be a strictly convex polygon with corners $p_1, \ldots, p_{2m}$ in the plane. We assume that $m \geq 2$, that p_1 and p_{2m} lie on the x-axis, and that the rest of the polygon is a convex chain that projects vertically onto the line segment $\overline{p_1 p_{2m}}$, which we call the *long side* of P. We call the other sides *short sides*. We choose P such that no short side is parallel to the long side.

Let Z be a (say, unit-radius) cylinder centered at the z-axis. For each vertex v_i of G, we take a copy P^i of P and place it vertically in $\mathbb{R}^3$ such that its long side lies on the boundary of Z; see Fig. 4a. Each polygon P^i lies inside Z on a distinct halfplane that is bounded by the z-axis. Finally, all polygons are positioned at the same height, implying that for any $j \in \{1, \ldots, 2m\}$, all copies of p_j lie on the same horizontal plane h_j and have the same distance to the z-axis.

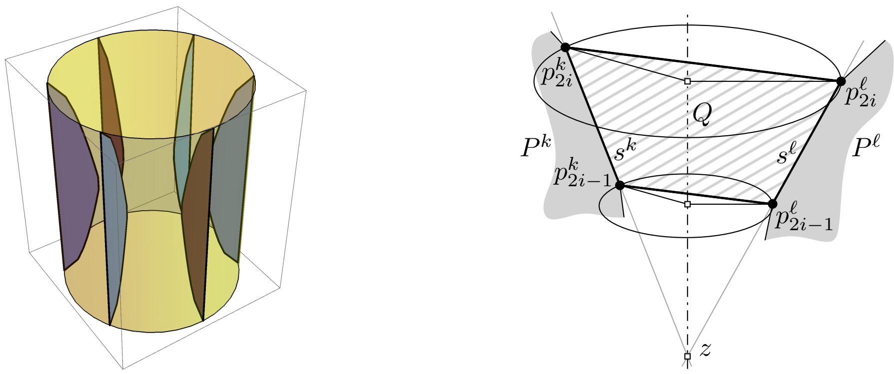

(a) polygons $P^1, \ldots, P^n$.

(b) quadrilateral Q spanned by s^k and s^ℓ.

Figure 4 Illustration for the proof of Theorem 6.

Let $i \in \{1, \ldots, m\}$. Then the side $s = p_{2i-1}p_{2i}$ is a short side of P. For $k = 1, 2, \ldots, n$, we denote by s^k and p^k_i the copies of s and p_i in P^k, respectively. We claim that, for $1 \leq k < \ell \leq n$, the sides s^k and s^ℓ span a convex quadrilateral that does not intersect any P^j with $j \notin \{k, \ell\}$. To prove the claim, we argue as follows; see Fig. 4b.

By the placement of P^k and P^ℓ inside Z, the supporting lines of s^k and s^ℓ intersect at a point z on the z-axis, implying that s^k and s^ℓ are coplanar. Moreover, p^k_{2i-1} and p^ℓ_{2i-1} are at the same distance from z, and the same holds for p^k_{2i} and p^ℓ_{2i}. Hence the triangle spanned by z, p^k_{2i-1}, and p^ℓ_{2i-1} is similar to the triangle spanned by z, p^k_{2i}, and p^ℓ_{2i}, implying that $p^k_{2i-1}p^\ell_{2i-1}$ and $p^k_{2i}p^\ell_{2i}$ are parallel and hence span a convex quadrilateral Q (actually a trapezoid). Finally, no polygon P^j with $j \notin \{k, \ell\}$ can intersect Q as any point in the interior of Q lies closer to the z-axis than any point of P^j at the same z-coordinate, which proves the claim.

We use Q as polygon for the subdivision vertex of the edge e_i of G (in case e_i was subdivided multiple times, we dissect Q accordingly). Let v_a and v_b be the endpoints of e_i. By our claim, Q that does not intersect any P^j with $j \notin \{a, b\}$. The quadrilateral Q lies in the horizontal slice of Z bounded by the horizontal planes h_{2i-1} and h_{2i}. Since any two such slices are vertically separated and hence disjoint, the m quadrilaterals together with the n copies of P constitute a valid representation of G'. ◀

Note that the combination of Proposition 5 and Theorem 6 rules out any Kuratowski-type characterization for adjacency graphs of convex-polyhedral surfaces. This graph class contains a subdivision of K_5, but it does not contain K_5; hence it is not minor-closed.

3.2 Complete Bipartite Graphs

▶ **Proposition 7.** *There exists a convex-polyhedral surface $\mathcal{S}$ such that $\mathcal{G}(\mathcal{S}) \simeq K_{4,4}$.*

Proof sketch. Start with a rectangular box in $\mathbb{R}^3$ and stab it with two rectangles that intersect each other in the center of the box as indicated in Fig. 5 (left). We can now draw polygons on these eight rectangles such that each of the four vertical rectangles (representing the four vertices of one side of the bipartition of $K_{4,4}$) contains a polygon that has a side contact with a polygon on each of the four horizontal or slanted rectangles (representing the other side of the bipartition of $K_{4,4}$). To remove the intersection of the (polygons drawn on the) two slanted rectangles, we shift one corner of the original box; see Fig. 5 (center and right). A description of the resulting polygons via their coordinates and more figures can be found in the full version [4]. Note that, in the resulting representation, there are two additional side contacts between pairs of polygons that are colored blue in Fig. 5 (right). By Proposition 2, these contacts can be removed. ◀

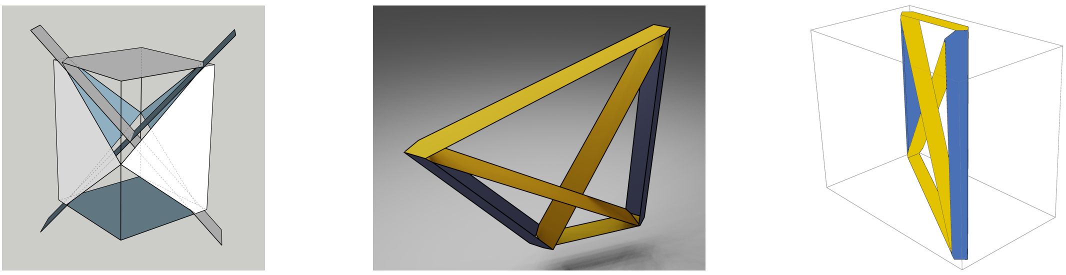

Figure 5 Construction of a convex-polyhedral surface $\mathcal{S}$ with $\mathcal{G}(\mathcal{S}) \simeq K_{4,4}$. The 2-coloring of the polygons in the central figure reflects the bipartition of $K_{4,4}$.

▶ **Proposition 8.** *There exists a convex-polyhedral surface $\mathcal{S}$ such that $\mathcal{G}(\mathcal{S}) \simeq K_{3,5}$.*

Proof. We call the vertices of the smaller bipartition class the gray vertices, and their polygons gray polygons. For the other class we pick a distinct color for every vertex and use the same naming-by-color convention. We start our construction with a triangular prism in which the quadrilateral faces q_1, q_2, q_3 are rectangles of the same size. Each of the faces q_i will contain one gray polygon. All colorful polygons lie inside the prism. We call the lines resulting from the intersection of the supporting planes with the prism the *colorful supporting lines.* Unfolding the faces q_1, q_2, and q_3 yields Fig. 6, which shows the gray polygons and the colorful supporting lines. Note that the vertices of the gray polygons in the figure are actually very small edges that have the slope of the colorful supporting line there are placed on. The colorful polygons are now already determined.

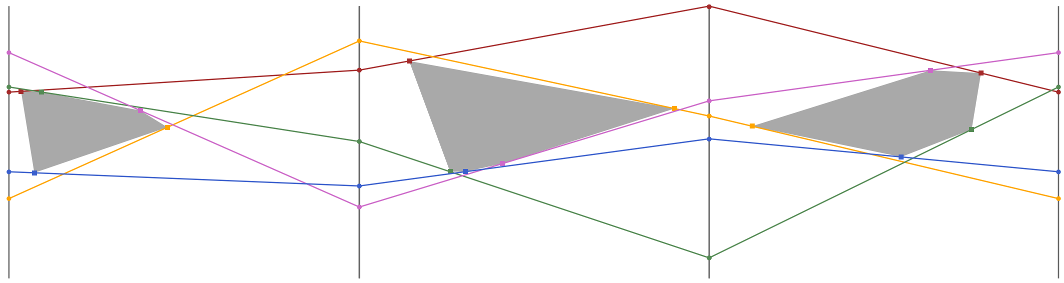

Figure 6 Constructing a convex-polyhedral surface whose adjacency graph is isomorphic to $K_{3,5}$. The prism boundary is unfolded into the plane.

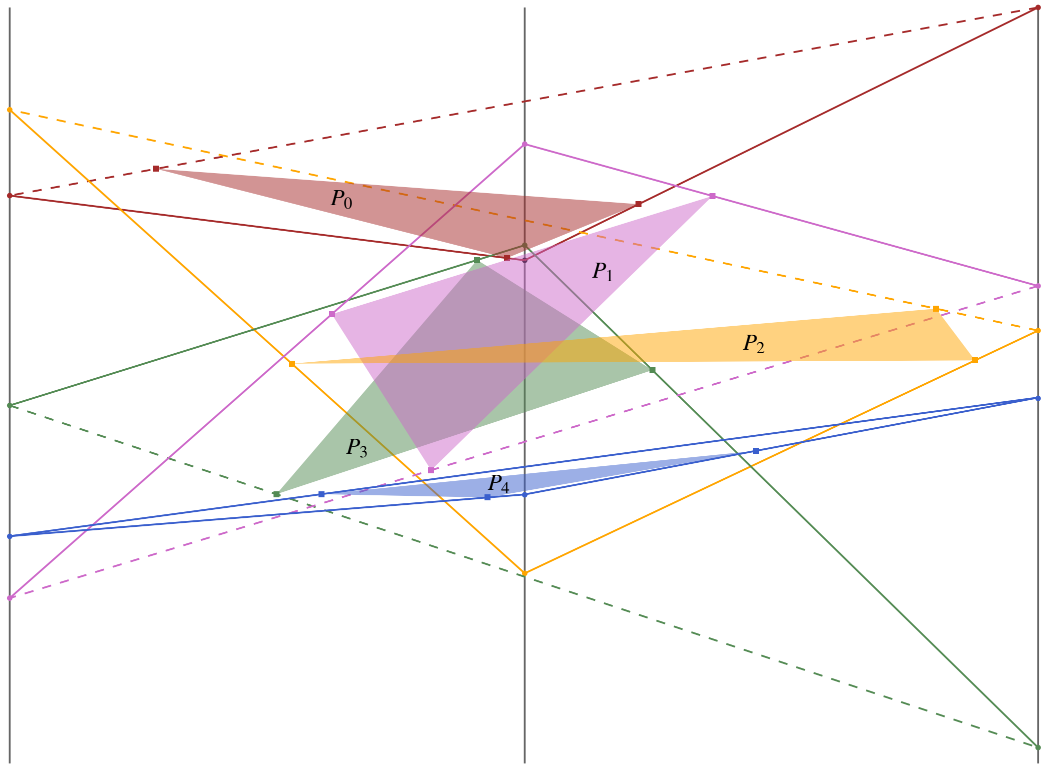

Figure 7 Front view of the prism containing a realization of $K_{3,5}$. Lines on the back are dashed.

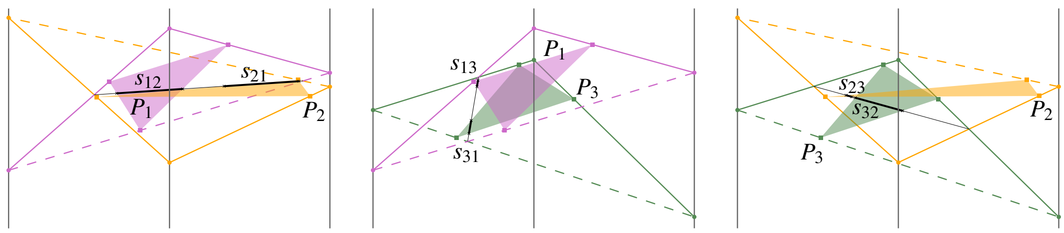

Figure 8 Certificates for the disjointness of colorful polygons P_1, P_2, and P_3. The supporting planes intersect in the thin black lines; thicker segments indicate where the lines intersect polygons.

We are left with checking that the colorful polygons are disjoint. Fig. 7 shows the prism in a view from the side where we dashed all objects on the hidden prism face. The brown polygon P_0 and the blue polygon P_4 avoid all other colorful polygons in this projection and thus they avoid all other polygons in $\mathbb{R}^3$, too.

For the pink polygon P_1, the orange polygon P_2 and the green polygon P_3, we proceed as follows to prove disjointness. Pick two of the polygons and name them P_i and P_j. The line ℓ_{ij} of intersection of the supporting planes of P_i and P_j is determined by the two intersections of the corresponding colorful supporting lines. If the polygons intersect, they have to intersect on this line. Polygon P_i intersects ℓ_{ij} in a segment s_{ij}; polygon P_j intersects ℓ_{ij} in s_{ji}. Fig. 8 shows, however, that s_{ij} and s_{ji} do not overlap in any of the three cases. ◀

In contrast to Propositions 7 and 8, we can show that not every complete bipartite graph can be realized as a convex-polyhedral surface in $\mathbb{R}^3$.

▶ **Theorem 9.** *There exists no convex-polyhedral surface $\mathcal{S}$ in $\mathbb{R}^3$ such that $K_{5,81}$ is subisomorphic to $\mathcal{G}(\mathcal{S})$.*

To prove the theorem we start with some observations about realizing complete bipartite graphs. We will consider a set R of red polygons, and a set B of blue polygons, so that each red–blue pair must have a side contact. For all $p \in R \cup B$ we denote by $p^=$ the supporting plane of p, by p^- the closed half-space left of $p^=$, and by p^+ the closed half-space right of $p^=$ (orientations can be chosen arbitrarily). We start with a simpler setting where we have an additional constraint. We call B *one-sided w.r.t.* R if all red polygons must be on the same half-space of each blue polygon, i.e., $\forall b \in B\colon ((\forall r \in R\colon r \subseteq b^-) \vee (\forall r \in R\colon r \subseteq b^+))$.

▶ **Lemma 10.** *Let R and B be two sets of convex polygons in $\mathbb{R}^3$ realizing $K_{|R|,|B|}$. If $|R| = 3$ and B is one-sided w.r.t. R, then $|B| \leq 8$.*

Proof. Let $R = \{r_1, r_2, r_3\}$ and let $\mathcal{A}$ be the arrangement of the supporting planes of R. Consider a polygon $b \in B$. For every polygon $r_i \in R$, since b is convex and shares a side with r_i, b is contained in r_i^- or r_i^+. Thus, b is contained in a cell of $\mathcal{A}$. Let $r_* = r_1^= \cap r_2^= \cap r_3^=$ be the intersection of the supporting planes of R. We may assume that no two supporting planes of R coincide; otherwise, by strict convexity, all polygons (in B and finally in $B \cup R$) must lie in the same plane and the non-planarity of $K_{3,3}$ implies that $|B| \leq 2$. Further, if r_* is a line, then every b lies in a cell C, in which one of the red polygons is only present as a subset of r_*. It is not possible to add b such that it has a common segment with r_* and each of the bounding planes of C. Consequently, r_* is either a point or the empty set. We consider two cases: either (1) $r_* = \emptyset$ or no polygon in R contains the point r_*, or (2) one polygon in R does contain the point r_*.

Case (1). Either $r_* = \emptyset$ or no polygon in R contains the point r_*. We claim that at most four cells of $\mathcal{A}$ are incident to all polygons of R, and at most two polygons of B exist per cell.

If $r_* = \emptyset$, $\mathcal{A}$ contains two parallel planes or forms an infinite prism. It is easy to check that $\mathcal{A}$ has at most four cells that are incident to three planes. If $r_* \neq \emptyset$, there are eight cells, called *octants*, of space of the form $Q^{a_0 b_0 c_0} = r_1^{a_0} \cap r_2^{b_0} \cap r_3^{c_0}$, $a_0, b_0, c_0 \in \{+, -\}$. If the point r_* is disjoint from all r_i, then each r_i intersects at most six octants: r_1 rules out the two octants $Q^{\pm b_1 c_1}$, r_2 "rules out" the two octants $Q^{a_2 \pm c_2}$, and r_3 rules out the two octants $Q^{a_3 b_3 \pm}$ for some $a_2, a_3, b_1, b_3, c_1, c_2 \in \{+, -\}$. The maximal number of octants that "remain" is four. For example one could pick $a_2 = a_3 = b_1 = b_3 = c_1 = c_2 = +$. By this choice, $Q^{+--}, Q^{-+-}, Q^{--+}$, and Q^{---} remain. All other choices (a finite set to check) will have at most four octants touching all three polygons.

Now consider one such cell C incident to all r_i. For the argument within this cell, we can truncate r_i to C. Since each r_i lies in a boundary plane of C and we are looking for a polygon b that has all of R on one side, the four polygons r_1, r_2, r_3, b must be in convex position. Assume that we have three such polygons b_1, b_2, and b_3. The cell C has three (unbounded) boundary faces f_1, f_2, and f_3 such that r_j lies in f_j. Let $\ell_{i,j}$ be the segment that is the intersection of the plane $b_i^=$ and the polygon f_j, and let t_i be the triangle formed by $\ell_{i,1}$, $\ell_{i,2}$, $\ell_{i,3}$. Two of these triangles intersect either in two points or they are disjoint. On no face f_j, all three $\ell_{1,j}$, $\ell_{2,j}$, $\ell_{3,j}$ can be disjoint since then the polygon belonging to the "middle segment" intersects the interior of r_j. On the other hand, every pair t_i, $t_{i'}$ has one face f_j without intersection. Since we have three possible pairs, every face f_j contains two segments $\ell_{i,j}$ and $\ell_{i',j}$ that do not intersect. Furthermore, the pair i, i' will be different for every f_j. We call the corresponding parallel segments *upper* and *lower*, depending on whether r_j lies below or above the corresponding plane ($b_i^=$ or $b_{i'}^=$). Now we have a contradiction to B being one-sided with respect to R since there has to be one t_i that belongs to an upper segment on one face f_j and to a lower segment on another face $f_{j'}$. In other words, not all polygons in R lie on the same side of $b_i^=$. Thus, at most two polygons of B lie in C, one corresponding to upper segments and one corresponding to lower segments. This yields at most $2 \cdot 4 = 8$ polygons in B in Case (1).

Case (2). One red polygon contains r_*, w.l.o.g. this is r_1. We claim that at most five cells of $\mathcal{A}$ are incident to all polygons of R, and at most one polygon of B exist per cell.

Similar as in Case (1), now r_2 and r_3 both intersect at most 6 octants, and there are at most five octants that intersect every polygon in R. Consider again some octant Q. Since r_1 contains r_*, now there may be only one polygon $b \in B$ that is contained in Q and has all polygons in R on the same side of $b^=$ (the lower segments in Case (1) are now forbidden). Thus, there are at most $1 \cdot 5 = 5$ polygons in B in Case (2). ◀

With the help of Lemma 10 we can now prove Theorem 9.

Proof of Theorem 9. Assume that $K_{5,81}$ can be realized, and let R be a set of five red polygons. Since every $b \in B$ is adjacent to all polygons in R, b partitions R into two sets: those in b^- and those in b^+. At least one of these subsets must have at least three elements. Arbitrarily charge b to such a set of three polygons. By Lemma 10, each set of three red polygons can be charged at most eight times. There are $\binom{5}{3} = 10$ sets of three red polygons. Therefore, there can be at most $8 \cdot 10 = 80$ blue polygons; a contradiction. Together with Proposition 2 this implies the claim. ◀

3.3 3-Trees

The graph class of 3-trees is recursively defined as follows: K_4 is a 3-tree. A graph obtained from a 3-tree G by adding a new vertex x with exactly three neighbors u, v, w that form a triangle in G is a 3-tree. We say x is *stacked* on the triangle uvw. It follows that for each 3-tree there exists a (not necessarily unique) *construction sequence* of 3-trees $G_4, G_5, \dots, G_n$ such that $G_4 \simeq K_4$, $G_n = G$, and where for $i = 4, 5, \dots, n-1$ the graph G_{i+1} is obtained from G_i by stacking a vertex v_{i+1} on some triangle of G_i.

By Proposition 3, for every planar 3-tree G there is a polyhedral surface $\mathcal{S}$ (even in $\mathbb{R}^2$) with $\mathcal{G}(\mathcal{S}) \simeq G$. On the other hand, we can show that no nonplanar 3-tree has such a realization in $\mathbb{R}^3$. To this end, we observe that a 3-tree is nonplanar if and only if it contains the *triple-stacked triangle* as a subgraph. The triple-stacked triangle is the graph that consists of $K_{3,3}$ plus a cycle that connects the vertices of one part of the bipartition; see Fig. 9. We show that the triple-stacked triangle is not realizable.

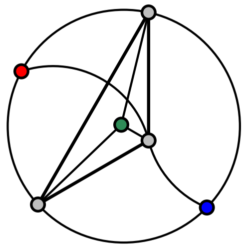

Figure 9 The unique minimal nonplanar 3-tree, which we call triple-stacked triangle.

▶ **Lemma 11.** *Let uvw be a separating triangle in a plane 3-tree $G = (V, E)$. Then there exist vertices $a, b \in V$ that belong to distinct sides of uvw in G such that both $\{a, u, v, w\}$ and $\{b, u, v, w\}$ induce a K_4 in G.*

Proof. Let $G_4, G_5, \dots, G_n$ denote a construction sequence of $G = G_n$, and let k be the largest index in $\{4, 5, \dots, n\}$ such that uvw is nonseparating in G_k. Since uvw is separating in G_{k+1}, it follows that the vertex $v_{k+1} = a$ is stacked on uvw (say, inside uvw) to obtain G_{k+1} and, hence, $\{a, u, v, w\}$ induce a K_4 in G_{k+1} and G.

It remains to argue about the existence of the vertex b in the exterior of uvw. If uvw is one of the triangles of the original $G_4 \simeq K_4$, there is nothing to show, so assume otherwise. Let j be the smallest index in $\{5, 6, \dots, n\}$ such that uvw is contained in G_j. It follows that one of u, v, w, say u, is the vertex v_j that was stacked on some triangle xyz of G_{j-1} to obtain G_j. Without loss of generality, we may assume that $\{v, w\} = \{y, z\}$. It follows that $x = b$ forms a K_4 with u, v, w in G_j and G. ◀

▶ **Lemma 12.** *A 3-tree is nonplanar iff it contains the triple-stacked triangle as a subgraph.*

Proof. The triple-stacked triangle is nonplanar because it contains a $K_{3,3}$ (one part of the bipartition is formed by the gray vertices and the other by the colored vertices).

For the other direction, let G be a nonplanar 3-tree. Let $G_4, G_5, \dots, G_n$ be a construction sequence of G. Let k be the smallest index in $\{4, 5, \dots, n\}$ such that G_k is nonplanar. By 3-connectivity, the graph G_{k-1}, which is planar, has a unique combinatorial embedding. Therefore, we may consider G_{k-1} to be a plane graph. Let uvw be the triangle that the vertex v_k was stacked on to obtain G_k from G_{k-1}. Since G_k is nonplanar, the triangle uvw is a separating triangle of G_{k-1}. It follows by Lemma 11 that G_k (and, hence, G) contains the triple-stacked triangle. ◀

▶ **Lemma 13.** *There exists no convex-polyhedral surface $\mathcal{S}$ in $\mathbb{R}^3$ such that triple-stacked triangle is subisomorphic to $\mathcal{G}(\mathcal{S})$.*

Proof. We refer to the vertices of the triple-stacked triangle as the three gray vertices and the three colored (red, green, and blue) vertices; see also Fig. 9. Given the correspondence between vertices and polygons (and their supporting planes), we also refer to the polygons (and the supporting planes) as gray and colored.

Assume that the triple-stacked triangle can be realized. Consider the arrangement of the gray supporting planes. By strict convexity, it follows that if a pair of gray polygons has the same supporting plane, then all their common neighbors lie in the same plane. This implies that all supporting planes coincide – a contradiction to the non-planarity of the triple-stacked triangle. Consequently, the gray supporting planes are pairwise distinct. (Likewise, it holds that no colored and gray supporting plane coincide.)

We now argue that all colored polygons are contained in the same closed cell of the gray arrangement. To see this, fix one gray polygon and observe, by Lemma 4, that all polygons are contained in the same closed half space with respect to its supporting plane.

Note that gray plane arrangement has one of the following two combinatorics: either the three planes have a common point of intersection (cone case) or not (prism-case). In the first case, the planes partition the space into eight cones, one of which contains all polygons; in the second case, the (unbounded) cell containing all polygons forms a (unbounded) prism. For a unified presentation, we transform any occurrence of the first case into the second case. To do so, we move the apex of the cone containing all polygons to the point at infinity by a projective transformation. This turns each face of the cone into a strip that is bounded by two of the extremal rays of the cone, which now has been deformed into a prism.

Consider one of the strips, which we call S. The strip S has to contain one of the gray polygons, which we call P_S. We know that P_S has at least five sides, one for each neighbor. Each of the two *bounding lines* contains a side to realize the adjacency to the other two gray polygons. We call the sides of P_S that realize the adjacencies to the remaining polygons red, green, and blue, in correspondence to the vertex colors. The supporting line of the red side intersects each bounding line of S. We add a red point at each of the intersections. For the blue and green sides we proceed analogously. By convexity of P_S, these points are distinct. This yields a permutation of red, green, blue (see Fig. 10) on each bounding line. The permutations on the boundary of two adjacent strips coincide because the supporting lines are clearly contained in the supporting planes.

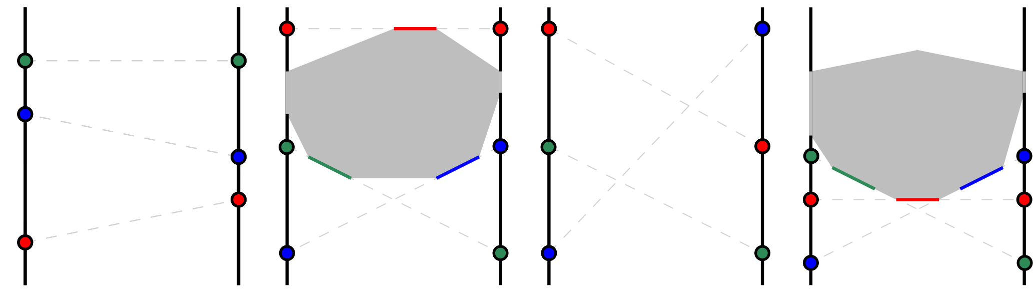

(a) 0 inversions. **(b)** 1 inversion. **(c)** 2 inversions. **(d)** 3 inversions.

Figure 10 The permutations of the intersections with the supporting lines of the red, green, and blue edges as in the proof of Lemma 13. Figures (a) and (c) illustrate possible scenarios. Figures (b) and (d) show impossible scenarios because they do not contain cells of complexity 5.

Consider the line arrangement inside S given by the supporting lines of the red, green, and blue sides. Up to symmetry, Fig. 10 illustrates the different intersection patterns. To realize all contacts, the polygon P_S has to lie inside a cell incident to all five lines, namely the two bounding lines and the supporting lines. It is easy to observe that such a cell exists only if the permutation has exactly one or three inversions, see Figs. 10b and 10d. In particular, the number of inversions is odd.

Following the cyclic order of the bounding lines around the prism, we record three odd numbers of inversions in the permutations before coming back to the start. Since an odd number of inversions does not yield the identity, we obtain the desired contradiction. ◀

Together Lemmas 12 and 13 yield the following theorem.

▶ **Theorem 14.** *Let G be a 3-tree. There exists a convex-polyhedral surface $\mathcal{S}$ in $\mathbb{R}^3$ with $\mathcal{G}(\mathcal{S}) \simeq G$ if and only if G is planar.*

In contrast to Theorem 14, there are nonplanar 3-degenerate graphs that can be realized; see the example in Fig. 1.

3.4 Hypercubes

In a paper from 1983, McMullen, Schulz, and Wills construct a polyhedron for every integer $p \geq 4$ such that all faces are convex p-gons [29, Sect. 4]. In the following, we show and illustrate how their result proves the realizability of any hypercube.

▶ **Proposition 15** ([29]). *For every d-hypercube Q_d, $d \geq 0$, there exists a convex-polyhedral surface $\mathcal{S}$ in $\mathbb{R}^3$ with $\mathcal{G}(\mathcal{S}) \simeq Q_d$ and every polygon of $\mathcal{S}$ is a $(d+4)$-gon.*

The main building block in their construction is a polyhedral surface whose adjacency graph is a $(p-4)$-hypercube. In fact, we observed that the adjacency graph of the polyhedron they finally construct is the Cartesian product of Q_{p-4} and a cycle graph $C_n, n \geq 3$. For the first few steps of their inductive construction, see Fig. 11.

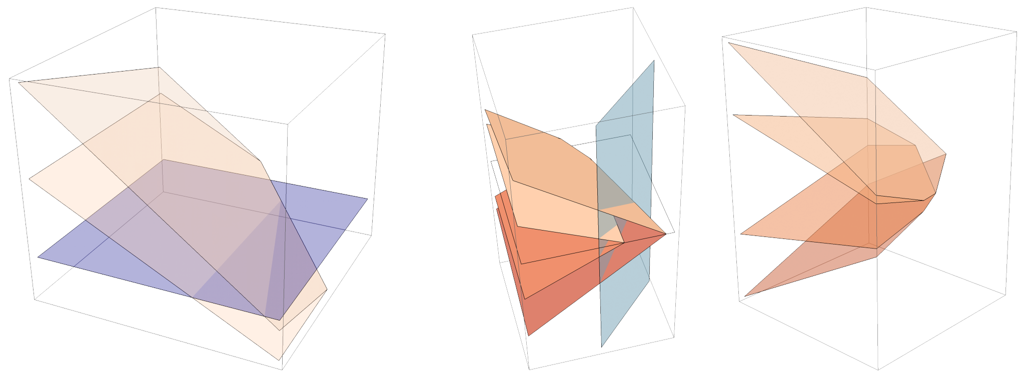

(a) Realization of Q_1 (red), xy-plane (blue). **(b)** Inductive step. **(c)** Realization of Q_2.

Figure 11 The inductive construction of McMullen et al. [29].

Recall that the d-hypercube has 2^d vertices. The base case for $d = 0$ is given by a single 4-gon, namely by the unit square. What follows is a series of inductive steps. In every step, the value of d increases by one and the number of polygons doubles. Before explaining the step, we state the invariants of the construction. After every step, all polygons have (almost) the same orthogonal projection into the xy-plane. Furthermore, this projection looks like the unit square in which we have replaced the upper right corner with a convex chain as shown in Fig. 12(a). The sides on the convex chain (with negative slopes) have already two incident polygons, the four other sides are currently incident to only one polygon. When projecting the polygons into the xy-plane only the convex chain edges differ.

We explain next how to execute the inductive step. Suppose that we have a polyhedral surface where every polygon is a $(d+4)$-gon fulfilling our invariant. We apply a shear along the z-axis and a vertical shift to the whole surface such that exactly the sides at $x = 1$ lie completely below the xy-plane, see also Fig. 11a. These transformations do not change the projections of the polygons to the xy-plane. We then cut the surface with the xy-plane. By this we slice away one of the sides in all polygons but also add a side that lies in the xy-plane; see Fig. 12(b). Each polygon now has a side that lies in the xy-plane and is disjoint from all other polygons. We now take a copy of the surface at hand and reflect it across the xy-plane. Every polygon of the original (unreflected) surfaces is now glued to its reflected copy via the common side in the xy-plane. With this step, we already have transformed the adjacency

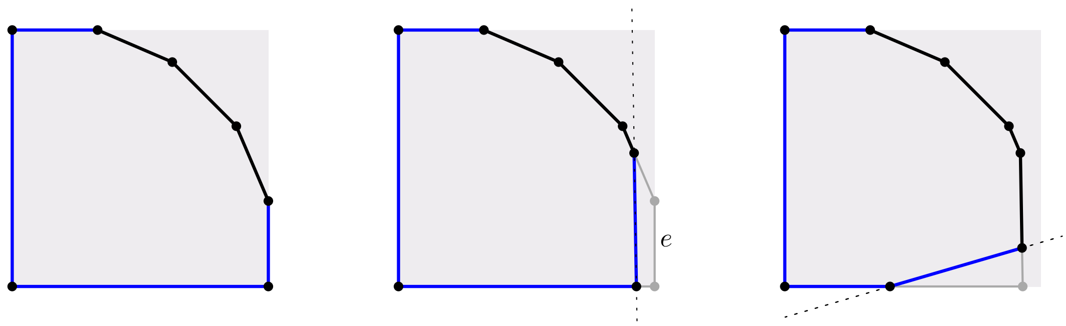

(a) (b) (c)

Figure 12 Projection of a polygon into the xy-plane in the construction of McMullen et al. The gray rectangle depicts the unit square. Edges incident to only one polygon are drawn in blue. (a) The start configuration for $d + 4 = 7$. (b) Cutting with the xy-plane after the shear that puts only e below the xy-plane (before glueing the reflected copy). (c) Slicing off a corner to get the initial situation for $d + 4 = 8$ modulo a projective transformation.

graph from a d-hypercube to a $(d + 1)$-hypercube. We only need to bring the surface back into the shape required by the invariant. To do so, we cut off a corner (see Fig. 12(c)) by slicing the whole construction with an appropriate plane, see also Fig. 11b. This turns all $(d + 4)$-gons into $(d + 5)$-gons. Finally, we apply a projective transformation to restore a required shape. Figures 11a–11c show spatial images of this construction.

4 Bounds on the Density

It is an intriguing question how dense adjacency graphs of convex-polyhedral surfaces can be. In this section, we use realizability and non-realizability results from the previous sections to derive asymptotic bounds on the maximum density of such graphs, which we phrase in terms of the relation between their number of vertices and edges.

Let $\mathcal{G}_n$ be the class of graphs on n vertices with a realization as a convex-polyhedral surface in $\mathbb{R}^3$. Further, let $e_{\max}(n) = \max_{G \in \mathcal{G}_n} |E(G)|$ be the maximum number of edges that a graph in $\mathcal{G}_n$ can have.

▶ **Corollary 16.** *$e_{\max}(n) \in \Omega(n \log n)$ and $e_{\max}(n) \in \mathcal{O}(n^{9/5})$.*

Proof. For the lower bound, note that by Proposition 15, every hypercube is the adjacency graph of a convex-polyhedral surface. As the d-dimensional hypercube has 2^d vertices and $2^d \cdot d/2$ edges, the bound follows.

For the upper bound, we use that, by Theorem 9, the adjacency graph of a convex-polyhedral surface cannot contain $K_{5,81}$ as a subgraph. It remains to apply the Kővari–Sós–Turán Theorem [24], which states that an n-vertex graph that has no $K_{s,t}$ as a subgraph can have at most $\mathcal{O}(n^{2-1/s})$ edges. ◀

Before being aware of the result of McMullen et al. [29], we constructed a family of surfaces with (large, but) constant average degree; see the full version [4]. Our construction is not recursive and therefore easier to understand and visualize; see Fig. 13. Note that some polygons in our construction have polynomial degree.

▶ **Proposition 17.** *There is an unbounded family of convex-polyhedral surfaces in $\mathbb{R}^3$ whose adjacency graphs have average vertex degree $12 - o(1)$.*

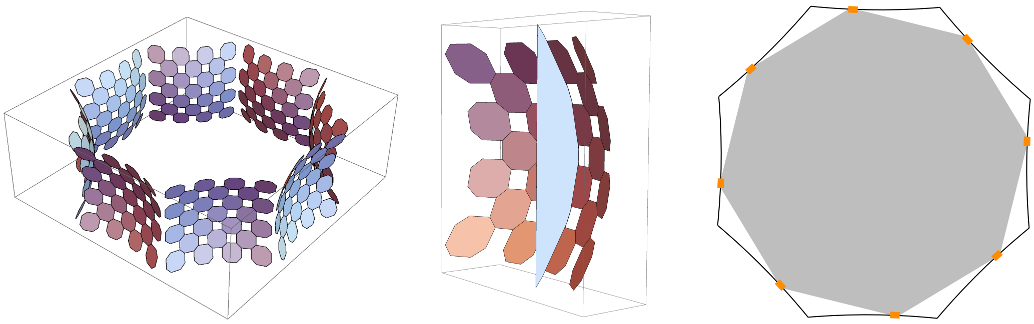

(a) placement of octagon grids. **(b)** a vertical polygon. **(c)** a horizontal polygon from above.

Figure 13 An family of convex-polyhedral surfaces whose adjacency graphs have average vertex degree $12 - o(1)$. The vertical polygons are attached to the "outside" of the grids; the horizontal polygons touch each grid along a single polygon side.

5 Conclusion and Open Problems

In this paper, we have studied the question which graphs can be realized as adjacency graphs of (convex-)polyhedral surfaces. In Corollary 16, we bound the maximum number $e_{\max}(n)$ of edges in realizable graphs on n vertices by $\Omega(n \log n)$ and $\mathcal{O}(n^{9/5})$. It would be interesting to improve upon these bounds. We conjecture that realizability is NP-hard to decide.

References

1 Md. Jawaherul Alam, Therese C. Biedl, Stefan Felsner, Andreas Gerasch, Michael Kaufmann, and Stephen G. Kobourov. Linear-time algorithms for hole-free rectilinear proportional contact graph representations. *Algorithmica*, 67(1):3–22, 2013. doi:10.1007/s00453-013-9764-5.

2 E. M. Andreev. Convex polyhedra in Lobačevskiĭ spaces. *Mat. Sb. (N.S.)*, 81 (123)(3):445–478, 1970. doi:10.1070/SM1970v010n03ABEH001677.

3 Boris Aronov, Marc J. van Kreveld, René van Oostrum, and Kasturi R. Varadarajan. Facility location on a polyhedral surface. *Discret. Comput. Geom.*, 30(3):357–372, 2003. doi:10.1007/s00454-003-2769-0.

4 Elena Arseneva, Linda Kleist, Boris Klemz, Maarten Löffler, André Schulz, Birgit Vogtenhuber, and Alexander Wolff. Adjacency graphs of polyhedral surfaces. ArXiv report, 2021. arXiv:2103.09803.

5 David W. Barnette and Branko Grünbaum. On Steinitz's theorem concerning convex 3-polytopes and on some properties of planar graphs. In G. Chartrand and S. F. Kapoor, editors, *The Many Facets of Graph Theory*, pages 27–40. Springer Berlin Heidelberg, 1969.

6 Martin Čadek, Marek Krčál, and Lukáš Vokřínek. Algorithmic solvability of the lifting-extension problem. *Discrete Comput. Geom.*, 57(4):915–965, 2017. doi:10.1007/s00454-016-9855-6.

7 Richard Cole and Micha Sharir. Visibility problems for polyhedral terrains. *J. Symb. Comput.*, 7(1):11–30, 1989. doi:10.1016/S0747-7171(89)80003-3.

8 Leila de Floriani, Paola Magillo, and Enrico Puppo. Applications of computational geometry in geographic information systems. In J.R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, chapter 7, pages 333–388. Elsevier, Amsterdam, 1997.

9 Hubert de Fraysseix, Patrice Ossona de Mendez, and János Pach. Representation of planar graphs by segments. *Intuitive Geometry*, 63:109–117, 1991. URL: https://infoscience.epfl.ch/record/129343/files/segments.pdf.

10 Hubert de Fraysseix, Patrice Ossona de Mendez, and Pierre Rosenstiehl. On triangle contact graphs. *Combinatorics, Probability and Computing*, 3:233–246, 1994. doi:10.1017/S0963548300001139.

11 David P. Dobkin. Computational geometry and computer graphics. *Proc. IEEE*, 80:141–1, 1992.

12 Christian A. Duncan, Emden R. Gansner, Yifan Hu, Michael Kaufmann, and Stephen G. Kobourov. Optimal polygonal representation of planar graphs. *Algorithmica*, 63(3):672–691, 2012. doi:10.1007/s00453-011-9525-2.

13 David Eppstein and Elena Mumford. Steinitz theorems for simple orthogonal polyhedra. *J. Comput. Geom.*, 5(1):179–244, 2014. doi:10.20382/jocg.v5i1a10.

14 William Evans, Paweł Rzążewski, Noushin Saeedi, Chan-Su Shin, and Alexander Wolff. Representing graphs and hypergraphs by touching polygons in 3D. In Daniel Archambault and Csaba Tóth, editors, *Proc. Graph Drawing & Network Vis. (GD'19)*, volume 11904 of *LNCS*, pages 18–32. Springer, 2019. doi:10.1007/978-3-030-35802-0_2.

15 Stefan Felsner. Rectangle and square representations of planar graphs. In János Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 213–248. Springer, 2013. doi:10.1007/978-1-4614-0110-0_12.

16 Stefan Felsner and Mathew C. Francis. Contact representations of planar graphs with cubes. In Ferran Hurtado and Marc J. van Kreveld, editors, *Proc. 27th Ann. Symp. Comput. Geom. (SoCG'11)*, pages 315–320. ACM, 2011. doi:10.1145/1998196.1998250.

17 Marek Filakovský, Uli Wagner, and Stephan Zhechev. Embeddability of simplicial complexes is undecidable. In *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA)*, pages 767–785, 2020. doi:10.1137/1.9781611975994.47.

18 Emden R. Gansner, Yifan Hu, and Stephen G. Kobourov. On touching triangle graphs. In Ulrik Brandes and Sabine Cornelsen, editors, *Proc. Graph Drawing (GD'10)*, volume 6502 of *LNCS*, pages 250–261. Springer, 2010. doi:10.1007/978-3-642-18469-7.

19 Daniel Gonçalves, Benjamin Lévêque, and Alexandre Pinlou. Homothetic triangle representations of planar graphs. *J. Graph Alg. Appl.*, 23(4):745–753, 2019. doi:10.7155/jgaa.00509.

20 Petr Hliněný. Contact graphs of line segments are NP-complete. *Discrete Math.*, 235(1):95–106, 2001. doi:10.1016/S0012-365X(00)00263-6.

21 Petr Hliněný and Jan Kratochvíl. Representing graphs by disks and balls (a survey of recognition-complexity results). *Discrete Math.*, 229(1–3):101–124, 2001. doi:10.1016/S0012-365X(00)00204-1.

22 Seok-Hee Hong and Hiroshi Nagamochi. Extending steinitz's theorem to upward star-shaped polyhedra and spherical polyhedra. *Algorithmica*, 61(4):1022–1076, 2011. doi:10.1007/s00453-011-9570-x.

23 Lutz Kettner. Designing a data structure for polyhedral surfaces. In *Proc. 14th Annu. ACM Symp. Comput. Geom. (SoCG'98)*, pages 146–154, 1998. doi:10.1145/276884.276901.

24 Tamás Kővari, Vera T. Sós, and Pál Turán. On a problem of K. Zarankiewicz. *Coll. Math.*, 3(1):50–57, 1954. URL: http://eudml.org/doc/210011.

25 Linda Kleist and Benjamin Rahman. Unit contact representations of grid subgraphs with regular polytopes in 2D and 3D. In Christian Duncan and Antonios Symvonis, editors, *Proc. Graph Drawing (GD'14)*, volume 8871 of *LNCS*, pages 137–148. Springer, 2014. doi:10.1007/978-3-662-45803-7_12.

26 Stephen G. Kobourov, Debajyoti Mondal, and Rahnuma Islam Nishat. Touching triangle representations for 3-connected planar graphs. In Walter Didimo and Maurizio Patrignani, editors, *Proc. Graph Drawing (GD'12)*, volume 7704 of *LNCS*, pages 199–210. Springer, 2013. doi:10.1007/978-3-642-36763-2_18.

27 Paul Koebe. Kontaktprobleme der konformen Abbildung. *Berichte über die Verhandlungen der Sächsischen Akad. der Wissen. zu Leipzig. Math.-Phys. Klasse*, 88:141–164, 1936.

28 Jiří Matoušek, Martin Tancer, and Uli Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. *J. Europ. Math. Soc.*, 13(2):259–295, 2011. doi:10.4171/JEMS/252.

29 P. McMullen, C. Schulz, and J.M. Wills. Polyhedral 2-manifolds in E^3 with unusually large genus. *Israel J. Math.*, 46:127–144, 1983. doi:10.1007/BF02760627.

30 Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. Embeddability in $\mathbb{R}^3$ is NP-hard. *J. ACM*, 67(4):1–29, 2020. doi:10.1145/3396593.

31 Jürgen Richter-Gebert. *Realization spaces of polytopes*, volume 1643 of *Lecture Notes in Mathematics*. Springer, 1996. doi:10.1007/BFb0093761.

32 Oded Schramm. *Combinatorically Prescribed Packings and Applications to Conformal and Quasiconformal Maps*. PhD thesis, Princeton University, 2007. arXiv:0709.0710.

33 Arkadiy Skopenkov. Extendability of simplicial maps is undecidable. ArXiv report, 2020. arXiv:2008.00492.

34 Arkadiy Skopenkov. Invariants of graph drawings in the plane. *Arnold Math. J.*, 6:21–55, 2020. doi:10.1007/s40598-019-00128-5.

35 Ernst Steinitz. Polyeder und Raumeinteilungen. In *Encyclopädie der mathematischen Wissenschaften*, volume 3-1-2 (Geometrie), chapter 12, pages 1–139. B. G. Teubner, Leipzig, 1922.

36 Szilassi polyhedron. Wikipedia entry. Accessed 2019-10-08. URL: https://en.wikipedia.org/wiki/Szilassi_polyhedron.

37 Heinrich Tietze. Über das Problem der Nachbargebiete im Raum. *Monatshefte für Mathematik und Physik*, 16(1):211–216, 1905. doi:10.1007/BF01693778.