Thickness and Colorability of Geometric Graphs☆

Stephane Durocher[1]

Department of Computer Science, University of Manitoba, Winnipeg, Canada

Ellen Gethner

Department of Computer Science, University of Colorado Denver, Colorado, USA

Debajyoti Mondal[2,*]

Department of Computer Science, University of Manitoba, Winnipeg, Canada

Abstract

The geometric thickness of a graph G is the smallest integer t such that there exist a straight-line drawing Γ of G and a partition of its straight-line edges into t subsets, where each subset induces a planar drawing in Γ. Over a decade ago, Hutchinson, Shermer, and Vince proved that any n-vertex graph with geometric thickness two can have at most $6n-18$ edges, and for every $n \geq 8$ they constructed a geometric thickness-two graph with $6n-20$ edges. In this paper, we construct geometric thickness-two graphs with $6n-19$ edges for every $n \geq 9$, which improves the previously known $6n-20$ lower bound. We then construct a thickness-two graph with 10 vertices that has geometric thickness three, and prove that the problem of recognizing geometric thickness-two graphs is NP-hard, answering two questions posed by Dillencourt, Eppstein and Hirschberg. Finally, we prove the NP-hardness of coloring graphs of geometric thickness t with $4t-1$ colors, which strengthens a result of McGrae and Zito, when $t = 2$.

Keywords: Complete Graph, Geometric Thickness, Coloring.

1. Introduction

The *thickness* $\theta(G)$ of a graph G is the smallest integer t such that the edges of G can be partitioned into t subsets, where each subset induces a planar

☆A preliminary version of the paper appeared in the Proceedings of the 39th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2013) [12].

*Corresponding author

Email addresses: `durocher@cs.umanitoba.ca` (Stephane Durocher), `ellen.gethner@ucdenver.edu` (Ellen Gethner), `jyoti@cs.umanitoba.ca` (Debajyoti Mondal)

[1]Work of the author is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

[2]Work of the author is supported in part by a University of Manitoba Graduate Fellowship.

graph. Since 1963, when Tutte [23] first formally introduced the notion of graph thickness, this property of graphs has been extensively studied for its interest from both the theoretical [2, 5, 7] and practical point of view [19, 21]. A wide range of applications in circuit layout design and network visualization, have motivated the examination of thickness in the geometric setting [7, 13, 16]. The *geometric thickness* $\overline{\theta}(G)$ of a graph G is the smallest integer t such that there exist a *straight-line drawing* (i.e., a drawing on the Euclidean plane, where every vertex is drawn as a point and every edge is drawn as a straight line segment) Γ of G and a partition of its straight-line edges into t subsets, where each subset induces a planar drawing in Γ. If $t = 2$, then G is called a *geometric thickness-two graph* (or, a doubly-linear graph [16]), and Γ is called a *geometric thickness-two representation* of G. While thickness does not impose any restriction on the placement of the vertices in each planar layer, geometric thickness forces the same vertices in different planar layers to share a fixed point in the plane. Eppstein [13] clearly established this difference by constructing thickness-three graphs that have arbitrarily large geometric thickness.

1.1. Structural Properties

Geometric thickness has been broadly examined on several classes of graphs, e.g., complete graphs [7], bounded-degree graphs [4, 11, 13], and graphs with bounded treewidth [8, 10]. Hutchinson, Shermer, and Vince [16] examined properties of graphs with geometric thickness two. They proved that these graphs can have at most $6n - 18$ edges, and for every $n \geq 8$ they constructed a geometric thickness-two graph with $6n - 20$ edges. The graphs that gave the $6n - 20$ lower bound were *rectangle visibility graphs*, i.e., these graphs can be represented such that the vertices are axis-aligned rectangles on the plane with adjacency determined by the horizontal and vertical visibility. Hutchinson et al. [16] proved that a rectangle visibility graph can have at most $6n-20$ edges, therefore, any geometric thickness-two graph with more than $6n - 20$ edges (if exists) cannot be a rectangle visibility graph. Even after several attempts [7, 11] to understand the structural properties of geometric thickness-two graphs, the question whether there exists a geometric thickness two graph with $6n-18$ edges remained open for over a decade. Answering this question is quite challenging since although one can generate many thickness-two graphs with $6n - 18$ or $6n - 19$ edges, no efficient algorithm is known that can determine the geometric thickness of such a graph. However, by examining the point configurations that are likely to support geometric thickness-two graphs with large numbers of edges, we have been able to construct geometric thickness-two graphs with $6n - 19$ edges, which improves the previously known $6n - 20$ lower bound on the maximum number of edges that a graph with geometric thickness two can have. In Section 2 we have shown that the K_9 minus an edge is a thickness-two graph, which has $6n - 19$ edges. We then show that thickness-two graphs that do not contain K_9 minus an edge may also have large number of edges.

Theorem 1 *For each $n \geq 9$, there exists a geometric thickness-two graph with n vertices and $6n-19$ edges that contains K_9 minus an edge as a subgraph. For* *each $n \geq 11$, there exists a geometric thickness-two graph with $6n - 19$ edges that does not contain K_8.*

1.2. Recognition

Although thickness is known for all complete graphs [2] and complete bipartite graphs [5], geometric thickness for these graph classes is not completely characterized. Dillencourt, Eppstein and Hirschberg [7] proved an $\lceil n/4 \rceil$ upper bound on the geometric thickness of K_n, giving a nice construction for representations with geometric thickness $t = \lceil n/4 \rceil$. They also gave a lower bound on the geometric thickness of complete graphs that matches the upper bound for several smaller values of n. Their bounds show that the geometric thickness of K_{15} is greater than its thickness, i.e., $\overline{\theta}(K_{15}) = 4 > \theta(K_{15}) = 3$, which settles the conjecture of [18] on the relation between thickness and geometric thickness. Since the exact values of $\overline{\theta}(K_{13})$ and $\overline{\theta}(K_{14})$ are still unknown, Dillencourt et al. [7] hoped that the relation $\overline{\theta}(G) > \theta(G)$ could be established with a graph of smaller cardinality. In Section 3 we prove that the smallest such graph contains 10 vertices.

Theorem 2 *For every $n \leq 9$ and every graph G on n vertices, $\overline{\theta}(G) = \theta(G)$. For every graph $n > 10$, there exists a graph G' on n vertices such that $\overline{\theta}(G) > \theta(G)$.*

Since determining the thickness of an arbitrary graph is NP-hard [19], Dillencourt et al. [7] suspected that determining geometric thickness might be also NP-hard, and mentioned it as an open problem. The hardness proof of Mansfield [19] relies heavily on the fact that $\theta(K_{6,8}) = 2$. Dillencourt et al. [7] mentioned that this proof cannot be immediately adapted to prove the hardness of the problem of recognizing geometric thickness-two graphs by showing that $\overline{\theta}(K_{6,8}) = 3$. This complexity question has been repeated several times in the literature [8, 13] since 2000, and also appeared as one of the selected open questions in the 11th International Symposium on Graph Drawing (GD 2003) [6]. In Section 4 we settle the question by proving the problem to be NP-hard.

Theorem 3 *It is NP-hard to determine whether the geometric thickness of an arbitrary graph is at most two.*

1.3. Colorability

As a natural generalization of the well-known Four Color Theorem for planar graphs [3], a long-standing open problem in graph theory is to determine the relation between thickness and colorability [17, 22]. For every $t \geq 3$, the best known lower bound on the chromatic number of the graphs with thickness t is $6t-2$, which can be achieved by the largest complete graph of thickness t. On the other hand, every graph with thickness t is $(6t)$-colorable [17]. Recently, McGrae and Zito [20] examined a variation of this problem that given a planar graph and a partition of its vertices into subsets of at most r vertices, asks to assign a color (from a set of s colors) to each subset such that two adjacent vertices in different subsets receive different colors. They proved that the problem is NP-complete when $r = 2$ (respectively, $r > 2$) and $s \leq 6$ (respectively, $s \leq 6r - 4$) colors. In Section 5 we prove the NP-hardness of coloring geometric thickness-t graphs with $4t-1$ colors. As a corollary, we strengthen the result of McGrae and Zito [20] that coloring thickness-($t = r = 2$) graphs with 6 colors is NP-hard. Our hardness result is particularly interesting since no geometric thickness-t graph with chromatic number more than $4t$ is known.

Theorem 4 *It is NP-hard to color an arbitrary geometric thickness-t graph with $4t - 1$ colors.*

2. Geometric Thickness-Two Graphs with $6n - 19$ Edges

Let K_9' be the graph obtained by deleting an edge from K_9. In this section we first construct a geometric thickness-two representation Γ of K_9', which has $6n - 19$ edges. We then show how to add vertices and edges in Γ such that for any $n \geq 9$ one can construct a geometric thickness-two graph with $6n - 19$ edges. Although these graphs contain K_9' as a subgraph, we prove that this is not a necessary condition, i.e., we also construct geometric thickness-two graphs with $6n - 19$ edges that do not contain K_9'. Hutchinson et al. [16, Theorem 6] proved that if any geometric thickness-two graph with $6n - 18$ edges exists, then the convex hull of its geometric thickness-two representation must be a triangle. This representation is equivalent to the union of two plane triangulations that share at least six common edges, i.e., the three outer edges and the other three edges are adjacent to the three outer vertices, as shown in black in Figure 1(a). Since each triangulation contains $3n - 6$ edges, the upper bound of $2(3n - 6) - 6 = 6n - 18$ follows. These properties of an edge maximal geometric thickness-two representation motivated us to examine pairs of triangulations that create many edge crossings when drawn simultaneously. In particular, we first created a set of points interior to the convex hull such that addition of straight line segments from each interior point to the three points on the convex hull creates two plane drawings that, while drawn simultaneously, contain a crossing in all but the six common edges. Figure 1(b) illustrates such a scenario. We then tried to extend each of these two planar drawings to a triangulation by adding new edges such that every new edge crosses at least one initial edge. We found multiple distinct point sets for which all but one newly added edge cross at least one initial edge, resulting in multiple distinct geometric thickness-two representations with $2(3n-6)-7 = 6n-19$ edges. For example, see Figure 1(c), where the underlying graph is K_9', where $K_9' = K_9 \setminus (d, e)$.

We now introduce a few more definitions. Let Γ be a geometric thickness-two representation. A triangle in Γ is *empty* if it contains exactly three vertices on its boundary, but does not contain any vertex in its proper interior, e.g., the triangle Δghi in Figure 1(c). A quadrangle Q in Γ created by the intersection of two empty triangles is called *free* if neither the interior nor the boundary of Q contains any vertex of Γ, as shown in Figure 1(c) in dark-green shade.

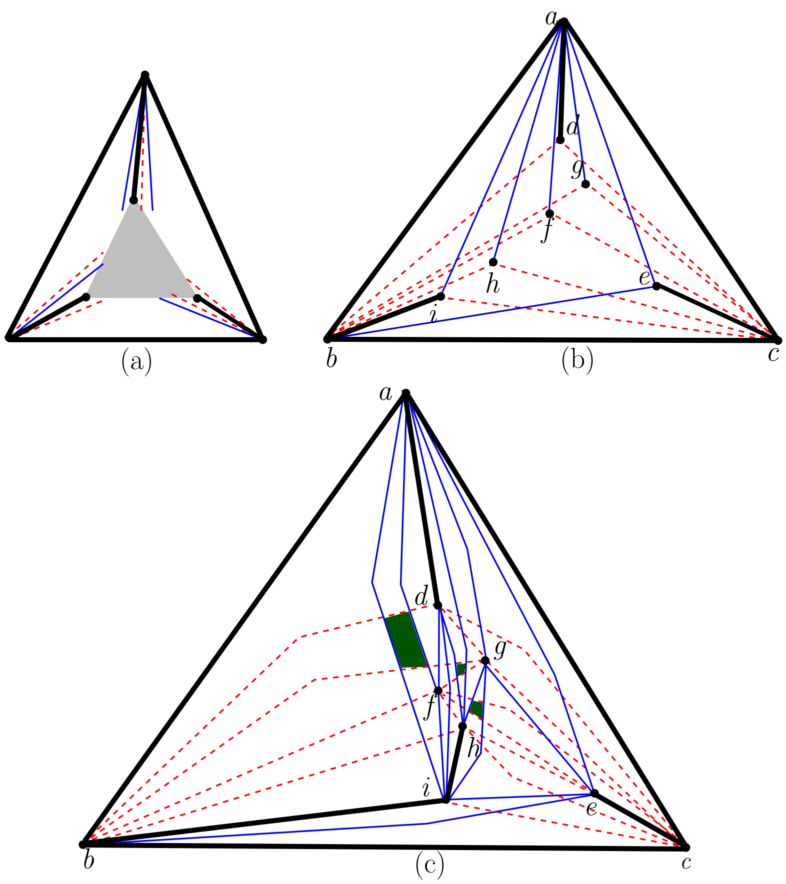

Figure 1: (a) Illustration for the shared edges (bold). (b) Initial point set. (c) A geometric thickness-two representation Γ of K_9', where $K_9' = K_9 \setminus (d, e)$. The planar layers are shown in red (dashed) and blue (thin). Black (bold) edges can belong to either red or blue layer. Free quadrangles are shown in dark-green (shaded). Some edges are drawn with bends for clarity.

Theorem 1. *For each $n \geq 9$, there exists a geometric thickness-two graph with n vertices and $6n - 19$ edges that contains K_9 minus an edge as a subgraph. For each $n \geq 11$, there exists a geometric thickness-two graph with $6n - 19$ edges that does not contain K_8.*

Proof. Since K_9' has a geometric thickness-two representation, as shown in Figure 1(c), the claim holds for $n = 9$. We now claim that given an n-vertex geometric thickness-two representation with $6n - 19$ edges that contains a free quadrangle, one can construct a geometric thickness-two representation with $n + 1$ vertices and $6(n + 1) - 19$ edges. One can verify this claim as follows. Place a new vertex p interior to the free quadrangle. Since a free quadrangle is the intersection of two empty triangles, one can add three straight line edges from p to the three vertices of each empty triangle such that the new drawing

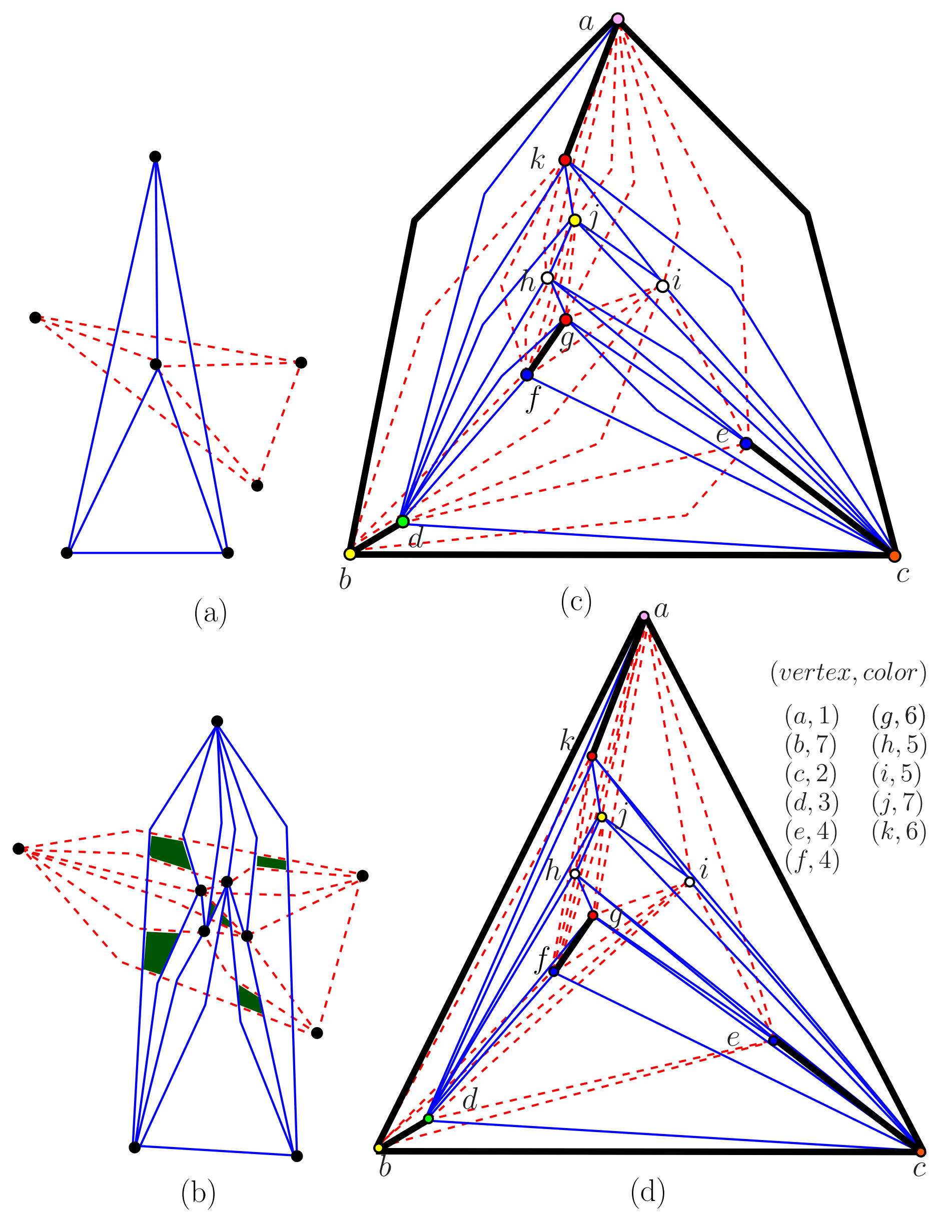

Figure 2: (a)–(b) Adding vertices to a geometric thickness-two drawing. (c)–(d) A graph with 11 vertices, 47 edges and geometric thickness two that does not contain K_8.

in each layer remains planar, as shown in Figure 2(a). Since the number of vertices increases by one, and the number of edges increases by six, the resulting geometric thickness-two representation must have $6n - 19 + 6 = 6(n+1) - 19$ edges.

Observe that there are at least three free quadrangles in the geometric thickness-two representation of K'_9, as shown in Figure 1(c). Furthermore, these quadrangles are *independent* of each other, i.e., the empty triangles corresponding to each quadrangle are different than the empty triangles of the other two quadrangles. Therefore, for each i, where $9 \leq i \leq 12$, we can construct a geometric thickness-two representation Γ_i with i vertices and $6i-19$ edges that contains at least one free quadrangle. We use these four geometric thickness-two representations as the base case, and assume inductively that for each $9 \leq i < n$ there exists a geometric thickness-two representation Γ_i with i vertices and $6i - 19$ edges that contains at least one free quadrangle. We now construct a geometric thickness-two representation with n vertices and $6n - 19$ edges that contains a free quadrangle. By induction, Γ_{n-4} has a free quadrangle. We add four vertices to this quadrangle and complete the triangulation in each planar layer by adding 24 new edges in total, as shown in Figure 2(b). Consequently, the new geometric thickness-two representation Γ_n contains $6(n-4) - 19 + 24 = 6n - 19$ edges. Since the newly added edges create new free quadrangles, Γ_n also contains a free quadrangle.

For each $n \geq 11$, the proof for the claim that there exists a geometric thickness-two graph with $6n - 19$ edges that does not contain any K_8 is similar. Figure 2(c) illustrates a geometric thickness-two representation with 11 vertices and 47 edges that contains at least three independent free quadrangles: $\{(\Delta cik, \Delta agj), (\Delta chj, \Delta aig), (\Delta dhj, \Delta bfi)\}$. The underlying graph does now contain K_8 as a subgraph since it is 7-colorable. To construct a geometric thickness-two representation Γ for a larger value of n, we add a K_8-free subgraph interior to a free quadrangle in the same way as in the earlier part of the proof (see Figures 2(a)–(b)). Hence the resulting graph must also be a K_8-free graph. □

3. Thickness-Two Graphs with $\overline{\theta}(G) \geq 3$

Dillencourt et al. [7] showed that $\overline{\theta}(K_{15}) = 4 > \theta(K_{15}) = 3$ and $\overline{\theta}(K_{6,8}) = 3 > \theta(K_{6,8}) = 2$, and asked to determine the smallest graph G such that $\overline{\theta}(G) > \theta(G)$. In this section we show that the smallest graph G such that $\overline{\theta}(G) > \theta(G)$ contains ten vertices, and construct such a geometric thickness-three graph.

To establish this result we enumerate all possible geometric thickness-two drawings of K'_9 using Aichholzer et al.'s [1] point-set order-type database[3]. Figure 3 illustrates all the three combinatorially different configurations of nine points that support geometric thickness-two drawings of K'_9. It might initially

[3]The code is available online: `http://www.cs.umanitoba.ca/~jyoti/Resources/DrawK9MinusOneEdge.java`

appear that Figures 3(a) and (b) are the same. However, observe that g lies on the left half-plane of (d, e) in Figure 3(a) and on the right-half plane of (d, e) in Figure 3(b).

We enumerated these geometric thickness-two representations by performing the steps S_1 and S_2 below for every point-set order-type P that consists of nine points. We use the concept of intersection graphs of segments: Given a set of straight line segments $\mathcal{L}$, the *proper intersection graph* G of $\mathcal{L}$ consists of $|\mathcal{L}|$ vertices, where each vertex corresponds to a distinct line segment in $\mathcal{L}$, and two vertices of G are adjacent if and only if the corresponding straight line segments properly cross.

S_1. Construct a straight-line drawing Γ of K_9 on P.

S_2. For every edge e^* in Γ, execute the following.

- Delete e^* and test whether the proper intersection graph determined by the remaining straight line segments is 2-colorable. If the graph is 2-colorable, then Γ is a geometric thickness-two representation of K'_9.
- Reinsert e^* in Γ.

Let $\Gamma_i, 1 \leq i \leq 3$, be the drawings of K'_9 depicted in Figures 3(a)–(c), respectively. The seven black (bold) edges in each of these drawings do not contain any crossing, i.e., these edges are shared in both triangulations. By E_i and E'_i we denote the set of all edges, and the set of black edges in Γ_i, respectively. Let $E''_i = E_i \setminus E'_i$. We verify that the partition of the edges of E''_i into red and blue is unique by checking that the proper intersection graph G_i of E''_i is connected. Figure 4(a) shows a spanning tree of G_1 of Γ_1.

Fact 1. *Let Γ be a geometric thickness-two representation of K'_9. Then the partition of the straight-line edges of Γ, except the seven edges that do not contain any proper crossing, into two planar layers is unique. Moreover, the unsaturated vertices are connected in each layer of Γ.*

We now categorize the vertices of a K'_9 into two types: *unsaturated* (vertices of degree 7), and *saturated* (vertices of degree 8). The vertices d and e of Figures 3(a)–(c) are unsaturated, and all other vertices are saturated vertices. Take a new vertex and make it adjacent to the two unsaturated vertices and any five saturated vertices of a K'_9. Let G_s denote the resulting graph with 10 vertices. The following theorem shows that $\overline{\theta}(G_s) = 3 > \theta(G_s) = 2$. The idea of the proof is first to show a thickness-two representation of G_s, and then to show that G_s contains a vertex v that is not straight-line visible to all of its neighbors in any geometric thickness-two representation of $G_s \setminus v$. Finally, the proof shows that for every graph G with at most 9 vertices, $\overline{\theta}(G) = \theta(G)$.

Theorem 2. *For every $n \leq 9$ and every graph G on n vertices, $\overline{\theta}(G) = \theta(G)$. For every graph $n > 10$, there exists a graph G' on n vertices such that $\overline{\theta}(G) > \theta(G)$.*

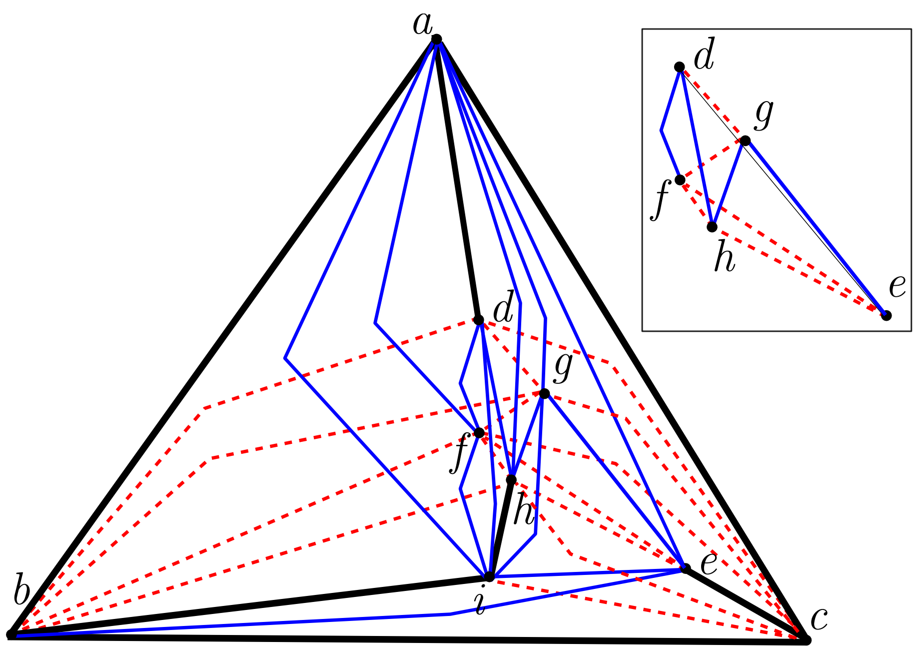

(a)

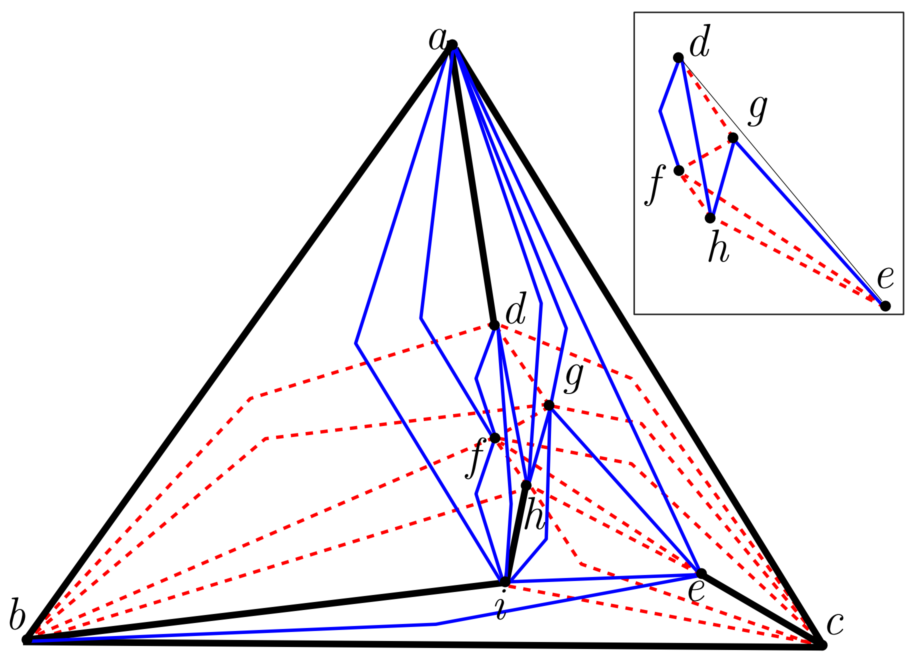

(b)

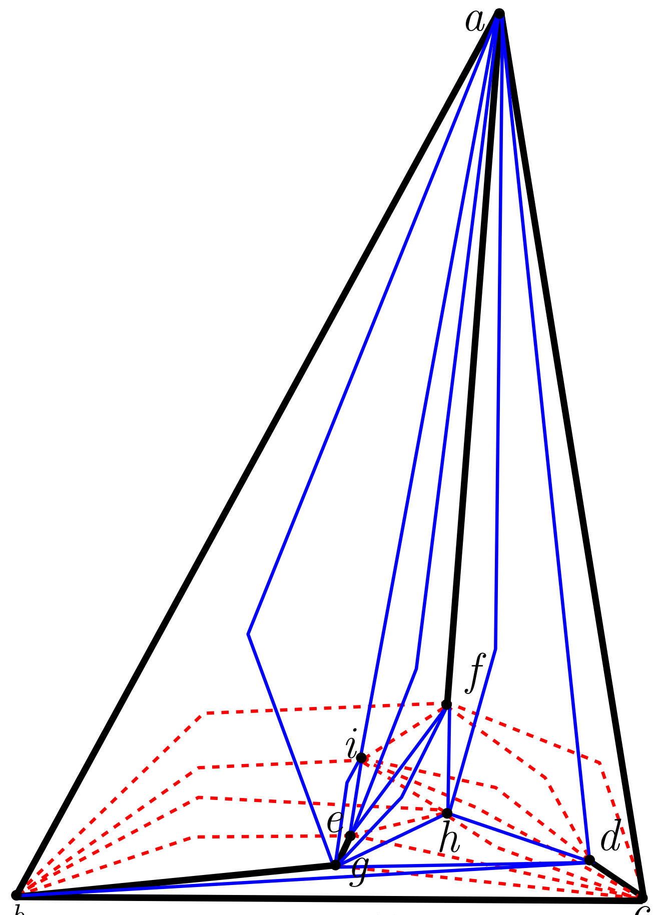

Figure 3: (a)–(c) Geometric thickness-two representations of K_9', where $K_9' = K_9 \setminus (d,e)$. Edges are drawn with polylines for clarity.

Proof. We first prove that the thickness of G_s is two. Let v be the vertex of degree seven in G_s such that $G_s \backslash v$ is a K_9'. Take the geometric thickness-two representation of K_9', as shown in Figure 3(a), and place the vertex v interior to the intersection of the triangles Δbci and Δbce. It is now straightforward to add the edges (with curved lines) from v to the two unsaturated vertices and five saturated vertices of K_9' maintaining the planarity of each layer. See Figure 4(b).

We now prove that the geometric thickness of G_s is three. Since G_s contains a K_9', any geometric thickness-two representation Γ of G_s must contain a geometric thickness-two representation Γ' of K_9' from Figures 3(a)–(c). Since Γ contains Γ', v cannot lie on any straight line segment of Γ'. A vertex u of a planar drawing is *straight-line visible* to a point p if the straight line between

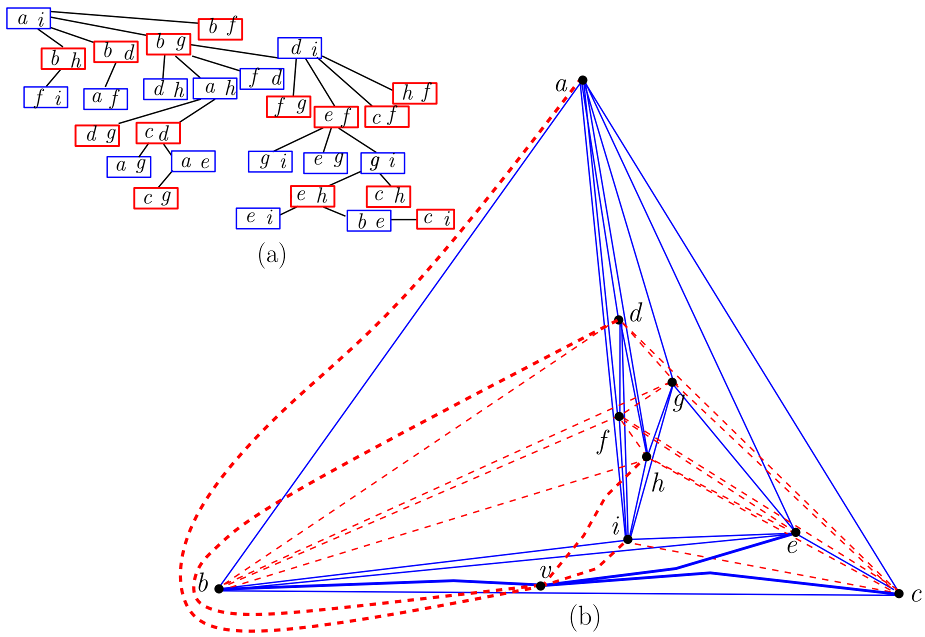

Figure 4: (a) A spanning tree of G_1. (b) The thickness of G_s is two.

p and u does not cross any edge or vertex of the drawing. We now prove that at most six vertices of Γ' can be straight-line visible to a common point, and hence v cannot be adjacent to seven vertices.

Delete all the black edges, i.e., the edges common to both triangulations, from Γ'. Let Γ'' denote the resulting drawing. Figures 5(a) and (b) show the candidate drawings that are obtained from Figures 3(a) and (c), respectively. We do not examine Figure 3(b) separately since its closed regions are similar to that of Figure 3(a). In each planar layer of Γ'', v must lie on some bounded region or on the unbounded region. Observe that the bounded regions in each planar layer are a collection of triangles and quadrangles. If v lies interior to an empty triangle in each planar layer, then it can have at most six adjacent edges. Therefore, v must lie on the unbounded region or on a quadrangle of some planar layer. Each such region Q is unfilled in Figure 5. Examining every such case reveals that no point in these regions can have straight-line visibility to seven distinct vertices of Γ', and hence $\overline{\theta}(G_s) = 3$. The details of the case by case analysis are found in Appendix A. Here we illustrate an example in the context of Figure 5(a), where the quadrangle $Q = idhg$ that includes v lies in the blue layer.

One can partition Q into 5 regions $Q_1, Q_2, \ldots, Q_5$, depending on how the neighboring segments of h intersects Q. Let q be a point in Q. If $q \in Q_1 \cup Q_2$, then since $\{d, g\}$ and $\{i\}$ lie on the opposite sides of the line determined by segment hb, vertex g cannot be visible to q. Therefore, q can see only 3 points

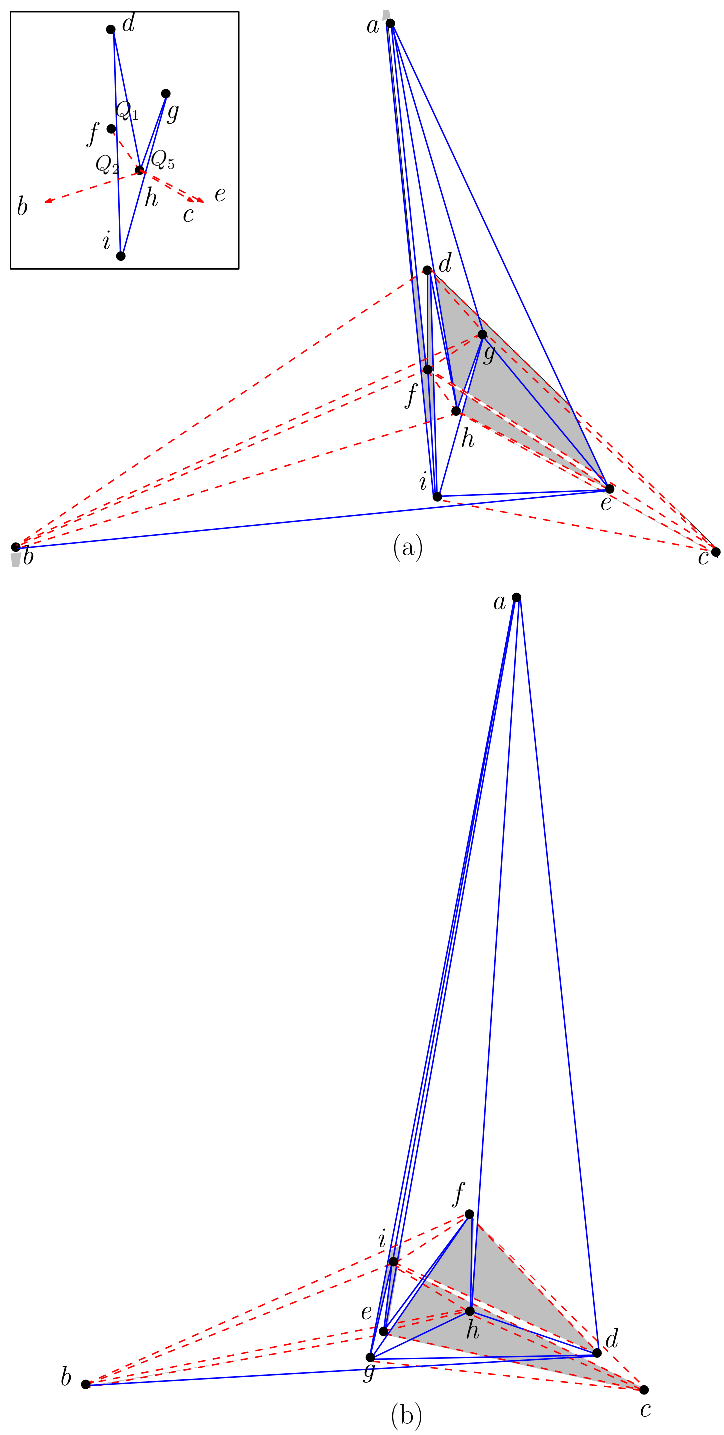

Figure 5: (a)–(b) Illustration for Γ''.

in the blue layer. Observe now that q is either inside a triangle in the red layer, or inside the quadrangle $fche$. If q is inside a triangle in the red layer, then q can see at most six vertices in total, and hence assume that q lies inside $fche$. Since $\{f, h\}$ and $\{e, c\}$ lie on the opposite sides of the line determined by segment ig, vertex h cannot be visible to q. Therefore, also in this case q can see at most six vertices in total. The case when $q \in Q_4 \cup Q_5$ can be analyzed in a similar fashion, where d and g cannot be visible to q simultaneously, and similarly h and f cannot be visible to q simultaneously. In the remaining case we have $q \in Q_3$, where q can see all the vertices $\{i, g, h, d\}$ in the blue layer. Since q is on the outer face in the red layer, only the vertices $\{b, c, i, h\}$ are visible to q. Therefore, q has straight-line visibility to at most six points in total.

Finally, observe that for every graph G with at most nine vertices, $\overline{\theta}(G) = \theta(G)$, as follows. Since $\theta(G) = 1$ if and only if $\overline{\theta}(G) = 1$ [21], we assume that $\theta(G) = 2$. Since $\overline{\theta}(K_9') = 2$, if G is a subgraph of K_9', then $\overline{\theta}(G) = \theta(G) = 2$. Otherwise, $G = K_9$, where $\overline{\theta}(K_9) = \theta(G) = 3$. Therefore, every graph G with $\overline{\theta}(G) > \theta(G)$ must have at least 10 vertices. □

4. Geometric Thickness-Two Graph Recognition

Mansfield [19] showed that the problem of recognizing thickness-two graphs is NP-hard. In this section we prove that determining whether $\overline{\theta}(G) \leq 2$ is NP-hard. For any input graph $\mathcal{G} = (V, E)$, we construct another corresponding graph $\mathcal{G}'$ such that $\mathcal{G}$ is a graph of thickness two if and only if the geometric thickness of $\mathcal{G}'$ is two. We first present some preliminary results, which will be useful to describe the construction of $\mathcal{G}'$. For the clarity of the presentation, proofs of some of the lemmas are given in Appendix B.

Let $G_1, G_2, \ldots, G_k$ be $k \geq 9$ copies of $K_9 - e$. Let d_1, e_1 be the unsaturated vertices of G_1. For each G_j, $j > 1$, make d_1 adjacent to some unsaturated vertex d_j of G_j. Refer to the remaining unsaturated vertex of G_j as e_j. Add a vertex v and make v adjacent to all the unsaturated vertices of $G_1, G_2, \ldots, G_k$. Let $\mathcal{H}_k$ denote the resulting graph, which we refer to as a *rigid graph*. The following lemma describes some properties of $\mathcal{H}_k$, whose proof is included in Appendix B.

Lemma 1. *Let $\mathcal{H}_k$ be a rigid graph. Then in any thickness-two drawing Γ of $\mathcal{H}_k$, the subgraph G' induced by the edges (v, d_1), (v, d_j) and (d_1, d_j), where $1 < j \leq k$, lies in the same layer.*

Observe that G' is a bipartite graph $K_{2,k-1}$ with vertex-partition $\{v, d_1\}$, $\{d_2, \ldots, d_k\}$, plus the edge (v, d_1). We call the graph G' the *core graph* and the vertex v the *pole vertex.*

Let H be a chain of three rigid graphs H_a, H_b, H_c, each a copy of $\mathcal{H}_{17}$. For any H_q, where $q \in \{a, b, c\}$, let $\{d_i^q, e_i^q\}$, where $1 \leq i \leq 17$, be the unsaturated vertices of H_q, and let v^q be the pole of H_q. We now add the edges $(e_2^a, e_2^b), (e_3^a, e_3^b), \ldots, (e_9^a, e_9^b)$ and $(e_2^c, e_{10}^b), (e_3^c, e_{11}^b), \ldots, (e_9^c, e_{17}^b)$. Let H' denote the resulting graph. Figures 6(a)–(b) illustrate a schematic representation of H'.

The following lemma describes some properties of H', whose proof is included in Appendix B.

Lemma 2. *In any thickness-two drawing* Γ *of* H'*, there exists a path from* $v^a, \ldots, v^b, \ldots, v^c$ *that lies on the same layer.*

It is straightforward to extend the above lemma for a chain of more than three rigid graphs, as stated below.

Corollary 1. *Let* H *be a chain of* $q \geq 3$ *rigid graphs* $H_1, H_2, \ldots, H_q$*, each a copy of* $\mathcal{H}_{17}$*. Let* v^i *be the pole of* H_i*, where* $1 \leq i \leq q$*. Then in any thickness-two drawing* Γ *of* H*, there exists a path from* $v^1, \ldots, v^2, \ldots, v^q$ *that lies on the same layer.*

We are now ready to describe the NP-hardness result.

Theorem 3. *It is NP-hard to determine whether the geometric thickness of an arbitrary graph is at most two.*

Proof. We reduce the problem of determining thickness-two graphs, which is NP-complete [19], to recognition of geometric thickness-two graphs. For any input graph $\mathcal{G} = (V, E)$, we construct another corresponding graph $\mathcal{G}'$ such that $\mathcal{G}$ is a graph of thickness two if and only if the geometric thickness of $\mathcal{G}'$ is two.

We construct $\mathcal{G}'$ by replacing each edge (v, w) by a chain H_{vw} of four rigid graphs $H_1, H_2, \ldots, H_4$, each a copy of $\mathcal{H}_{17}$, such that the vertices v and w coincide to poles of H_1 and H_4. Figures 7(a) and (b) depict an input graph G and a schematic representation of the corresponding graph $\mathcal{G}'$, respectively. It is straightforward to construct $\mathcal{G}'$ in polynomial time. We now show that $\mathcal{G}$ is a graph of thickness two if and only if $\mathcal{G}'$ admits a geometric thickness-two representation.

First assume that the thickness of $\mathcal{G}$ is two, and let $\{E_r, E_b\}$ be the corresponding partition of the edges, i.e., $E = (E_r \cup E_b)$ and $E_r \cap E_b = \phi$. We now compute a geometric thickness-two representation of $\mathcal{G}'$. Note that the graphs $\mathcal{G}_r = (V, E_r)$ and $\mathcal{G}_b = (V, E_b)$ are planar. Therefore, we can use the algorithm of Erten and Kobourov [14, 15] to compute a drawing of G on an $O(|V|^3) \times O(|V|^3)$ grid R such that the following properties hold:

P_1. No two edges of E_r or E_b cross.

P_2. Each edge is drawn as polygonal chain with at most two bends, and

P_3. The vertices and bends lie on some integer grid point of R.

Figure 7(c) illustrates an example. We now replace each edge (v, w) of $\mathcal{G}$ with the corresponding chain H_{vw} such that the edges between any pair of rigid graphs lie on the same layer as (v, w). We draw the rigid graphs around the vertices and bend points such that they do not interfere with the other edges of the drawing. Figure 7(d) depicts a schematic geometric thickness-two representation of $\mathcal{G}'$ which corresponds to the drawing of Figure 7(c).

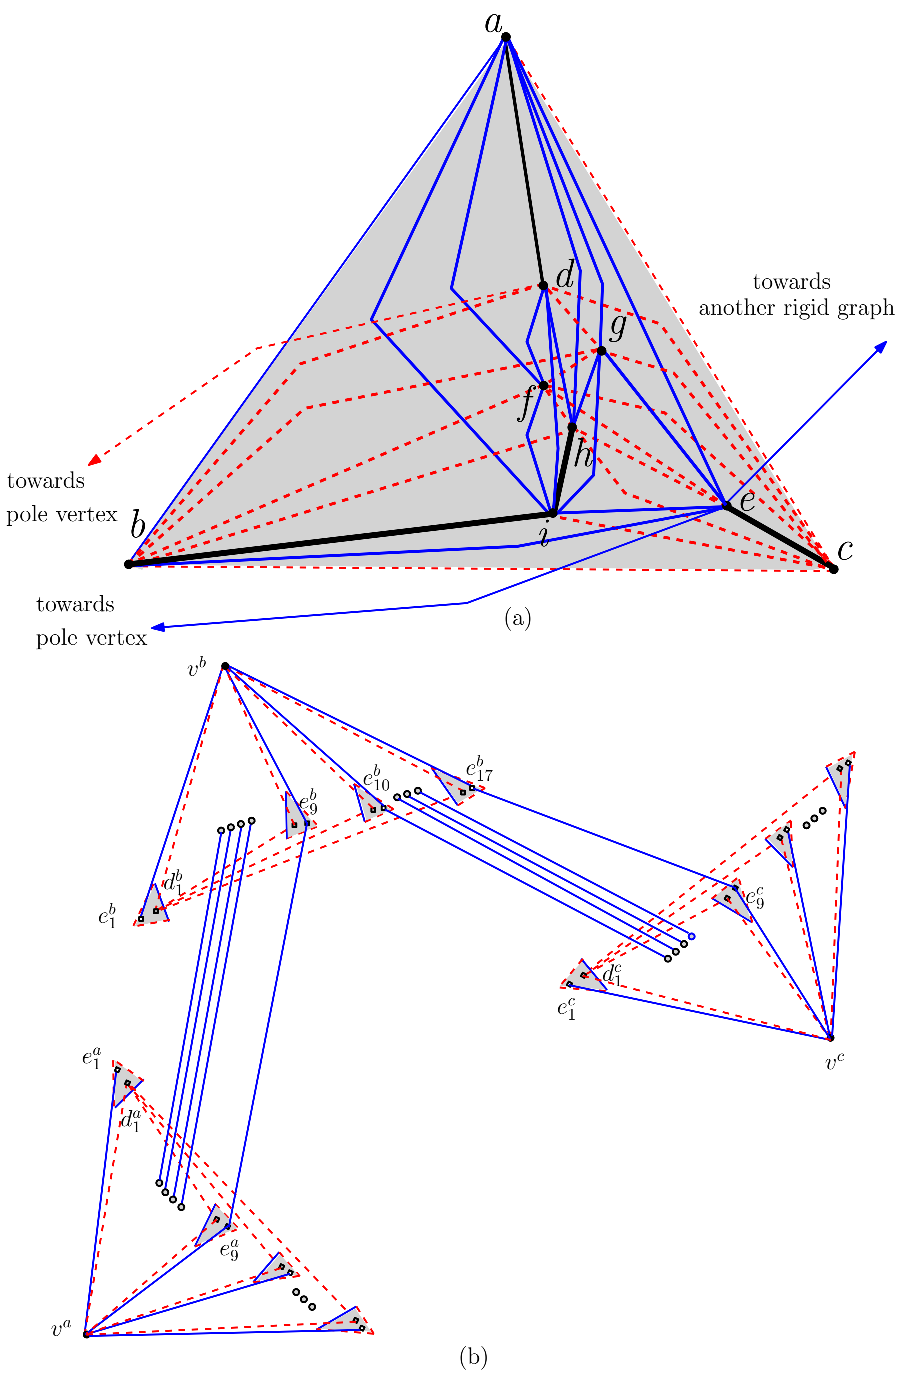

Figure 6: (a) Connections of $K_9 - e$ inside and outside of a rigid graph. (b) A schematic representation of H'.

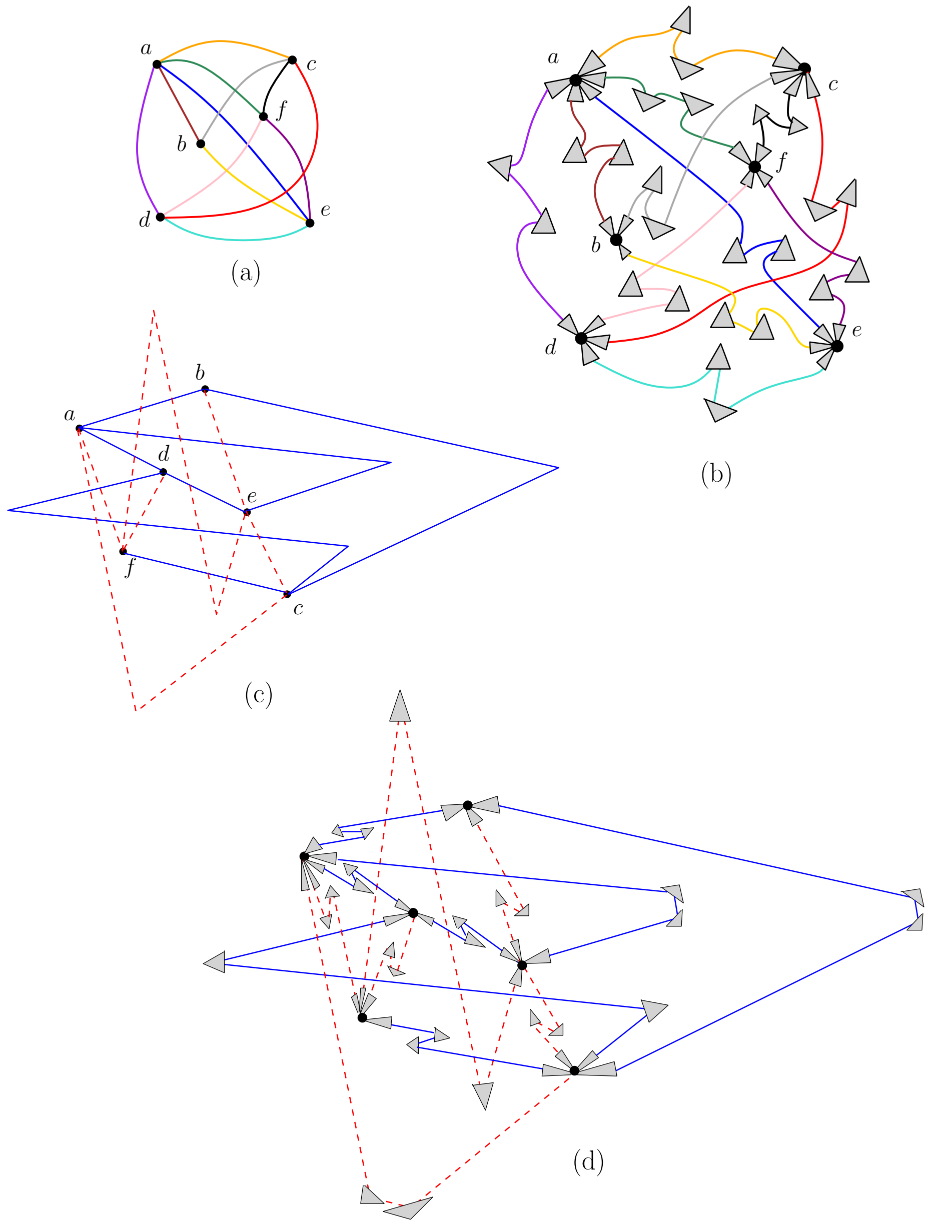

Figure 7: (a) An input graph $\mathcal{G}$. (b) A schematic representation of $\mathcal{G}'$. (c) A thickness-two drawing of $\mathcal{G}$. (d) The corresponding geometric thickness-two representation of $\mathcal{G}'$.

Assume now that $\mathcal{G}'$ admits a geometric thickness-two representation Γ. We show how to compute a thickness-two drawing of $\mathcal{G}$. By Corollary 1, for each edge (v, w) in $\mathcal{G}$, there exists a path $P = (v, \ldots, w)$ in H_{vw} such that all edges of P lie on the same layer of Γ. We delete all edges of H_{vw} except the edges of P, and any resulting isolated vertices. Since deletion of vertices and edges from a geometric thickness-two drawing does not increase geometric thickness, we end up with a thickness-two representation of $\mathcal{G}$. $\square$

Note that we do not yet know whether the problem of finding geometric thickness is in NP. Given a certificate, e.g., a partition of the edges of the given graph G into t sets, it is not clear how to determine whether the graphs induced by each edge set admit a simultaneous geometric embedding (which would imply that the geometric thickness of G is at most t).

5. NP-hardness of Colorability

In this section we show the NP-hardness of coloring a graph with geometric thickness t with $4t - 1$ colors. By $I(G, T, C)$ we denote the problem of coloring a graph G with C colors, where $\overline{\theta}(G) \leq T$. We first introduce a few definitions. A *join* between two graphs is an operation that given two graphs, adds all possible edges that connect the vertices of one graph with the vertices of the other graph. By $\mathcal{G}_t$ we denote a class of thickness-t graphs that satisfies the following conditions:

1. $\mathcal{G}_1$ is the class of planar graphs.
2. If $t>1$, then $\mathcal{G}_t$ consists of the graphs obtained by taking a join of K_2 and G, where $G \in \mathcal{G}_{t-1}$.

We now have the following lemma.

Lemma 3. *It is NP-hard to color an arbitrary graph $G \in \mathcal{G}_t$ with $2t+1$ colors.*

Proof. If $t = 1$, then coloring a planar graph (i.e., $t = 1$) with $2t+1 = 3$ colors is NP-hard [17]. Assume inductively that the claim holds for any thickness less than t. We now prove the hardness for thickness t, where $t > 1$, as follows.

Given an instance $I(G, t-1, 2(t-1)+1)$ where $G \in \mathcal{G}_{t-1}$, we construct a graph H by joining K_2 with a copy of G. Observe that $H \in \mathcal{G}_t$. It is now straightforward to obtain a geometric thickness-t representation for H by adding the vertices of K_2 to a thickness-$(t-1)$ representation of G (one to the left and the other to the right of the representation) such that the edges adjacent to the vertices of K_2 lie in the t-th planar layer but do not create any proper edge crossings. Therefore, $\overline{\theta}(H) \leq t$, and it now suffices to prove that G is $(2(t-1)+1)$-colorable if and only if H is $(2t+1)$-colorable.

If G is $(2(t-1)+1)$-colorable, then we can color the copy of G that is contained in H with $2(t-1)+1$ colors. Finally, we color the vertices of K_2 with two new colors, and thus obtain a $2t+1$ coloring for H.

On the other hand, if H is $(2t+1)$-colorable, then the vertices of K_2 must have two different colors. The vertices of H that belong to the copy of G cannot use these two colors. Therefore, G must be $(2t+1-2)$-colorable, i.e., $(2(t-1)+1)$-colorable. □

Note from the proof of Lemma 3 that $\overline{\theta}(\mathcal{G}_t) \leq t$. We use Lemma 3 to prove the NP-hardness of coloring geometric thickness-t graphs with $4t-1$ colors. We employ induction on t. If $t=1$, then coloring a planar graph (i.e., $t=1$) with $4t-1=3$ colors is NP-hard [17]. We now assume inductively that for any $t' < t$, it is NP-hard to color a geometric thickness-t' graph with $4t'-1$ colors. To prove the hardness of coloring a geometric thickness-t graph with $4t-1$ colors, we reduce the hardness of coloring a geometric thickness-$(t-1)$ graph with $2(t-1)+1$ colors. Given an instance $I(G, t-1, 2(t-1)+1)$, where $G \in \mathcal{G}_{t-1}$, we construct a graph $H(G,t)$ such that $\overline{\theta}(H(G,t)) \leq t$ and $H(G,t)$ is $(4t-1)$-colorable if and only if G is $(2(t-1)+1)$-colorable.

5.1. Construction of $H(G,t)$

Let the number of vertices in G be n. Take n copies $H_1, H_2, \ldots, H_n$ of K_{2t}, and join each vertex of G with a distinct H_i, $1 \leq i \leq n$. Finally, take a copy H' of K_{2t-1} and join it with every H_i. Let the resulting graph be $H(G,t)$. To prove that $\overline{\theta}(H(G,t)) = t$, we first review a construction of Dillencourt et al. [7] that gives a thickness-t representation of K_{4t}.

Dillencourt et al. [7] proved that the $4t$ vertices of a K_{4t} can be arranged in two rings of $2t$ vertices each, an outer ring and an inner ring, such that it can be embedded using exactly t planar layers. The vertices of the inner ring are arranged to form a regular $2t$-gon. For each pair of diametrically opposite vertices, a zigzag path is constructed as illustrated in Figure 8(a). This path has exactly one diagonal connecting diametrically opposite points (i.e., the diagonal connecting the two gray points in the figure.) The union of these zigzag paths, taken over all t pairs of diametrically opposite vertices, contains all the edges of K_{2t} in the inner ring, as shown in Figure 8(b). Consider now any zigzag path Z. For each pair of diametrically opposite vertices, we can draw rays in two opposite directions, so that none of the rays crosses any edge of Z. These rays, in each direction, meet at a common point (e.g., p or q) forming the outer ring, as shown in Figure 8(c).

Lemma 4. *$\overline{\theta}(H(G,t)) \leq t$, where $t > 1$ and $G \in \mathcal{G}_{t-1}$.*

Proof. We compute a geometric thickness-t representation of $H(G,t)$, as follows. Since $G \in \mathcal{G}_{t-1}$, $\overline{\theta}(G) \leq t-1$. Take a geometric thickness-$(t-1)$ representation of G and rotate it (if necessary) such that no two vertices lie on the same vertical line. Let Γ denote the resulting drawing. For the ith vertex in Γ, consider a thin vertical stripe S_i through it, as shown in Figure 8(d) in gray. Consider a horizontal stripe L below Γ that intersects all the vertical strips. For each i, we construct an inner ring (shown in black disk) that lies inside the intersection of S_i and L, where this ring corresponds to the drawing of H_i.

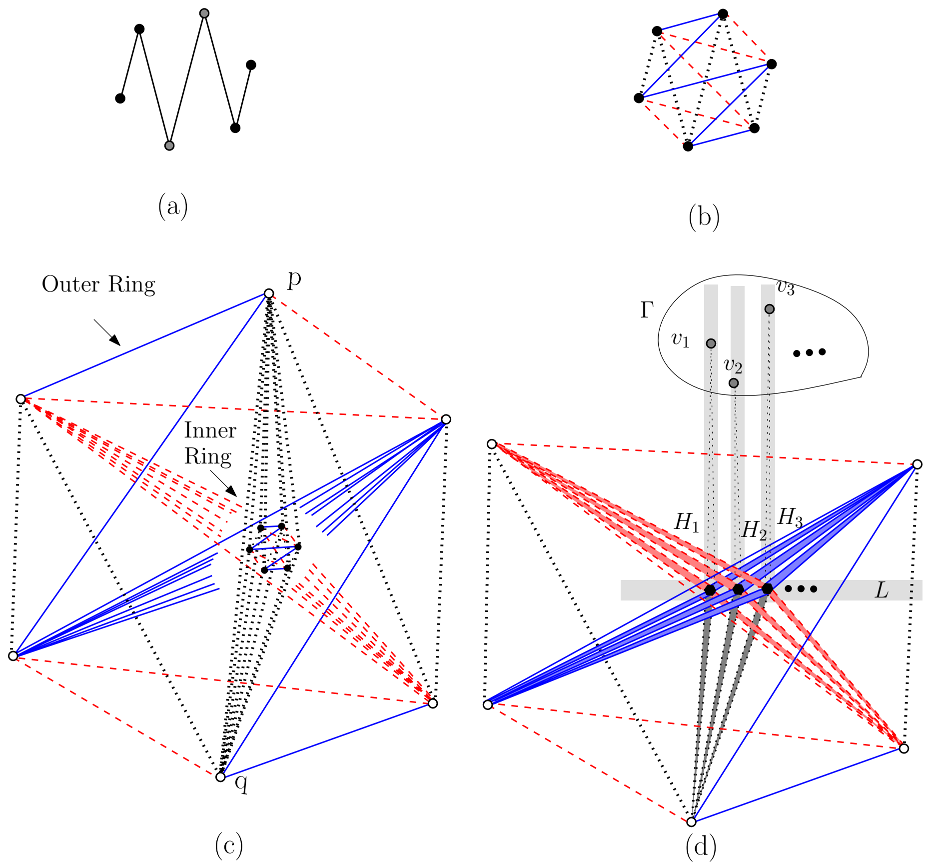

Figure 8: (a)–(c) Dillencourt et al.'s construction [7]. (a) A zigzag path in the inner ring. (b) K_{2t}, where t=3. (c) K_{4t}, where t=3. (d) The geometric thickness-three representation of $H(G, t)$, where t=3. Each subgraph H_i is determined by an inner ring, shown in black dot inside the horizontal strip L. The vertices of the outer ring are shown in unfilled circle.

The edges that connect the vertices of G with the vertices of H_i (i.e., the edges in the vertical stripes) lie in the t-th layer. Note that the inner rings must be scaled down small enough such that these edges do not create any edge crossing in any planar layer. Now construct an outer ring as in Dillencourt et al.'s construction [7], and delete a vertex from the ring to obtain a geometric thickness-t drawing of H', as shown in Figure 8(d). $\square$

5.2. Reduction

Given a geometric thickness-t graph G and a certificate coloring of G, one can check in polynomial time whether the number of colors used is at most $4t - 1$, and whether each edge of G receives two different colors at its end vertices. Therefore, the problem of coloring geometric thickness-t graphs with

$4t-1$ colors is in NP. The following theorem proves that the problem is NP-hard.

Theorem 4. *It is NP-hard to color an arbitrary geometric thickness-t graph with* $4t-1$ *colors.*

PROOF. If $t = 1$, then coloring a planar graph (i.e., $t = 1$) with $4t-1 = 3$ colors is NP-hard [17]. Assume now that $t > 1$. Given an instance $I(G, t-1, 2(t-1)+1)$, where $G \in \mathcal{G}_{t-1}$, we construct the corresponding graph $H(G, t)$. We prove that G is $(2(t-1)+1)$-colorable if and only if $H(G, t)$ is $(4t-1)$-colorable.

If G is $(2(t-1)+1)$-colorable, then we can color the copy of G that is contained in $H(G, t)$ with $2(t-1)+1$ colors. Since there does not exist any edge in $H(G, t)$ that connects a vertex of H_i with a vertex of H_j, where $i \neq j$, we can color all H_is with $2t$ new colors. The remaining vertices (i.e., the vertices of H') induce a K_{2t-1} such that none of these vertices are adjacent to the vertices of G. Therefore, we can reuse the $2t-1$ colors that we used to color the vertices of G. Consequently, number of colors we used to color $H(G, t)$ is $2(t-1)+1+2t = 4t-1$.

On the other hand, assume that $H(G, t)$ is $(4t-1)$-colorable. Since the vertices of H' must have $2t-1$ different colors, and since each H_i is joined with H', the vertices of H_i must use the remaining $2t$ colors. Since every vertex v of G is joined with a copy of distinct H_i, v must be colored with a color from the $2t-1$ colors used to color H'. Therefore, G must be $(2t-1)$-colorable, i.e., $(2(t-1)+1)$-colorable. $\square$

6. Discussion and Possible Directions for Future Research

In Section 2 we constructed n-vertex geometric thickness-two graphs with $6n-19$ edges, where the best known upper bound on the number of edges is $6n-18$ [16]. It is still unknown whether there exists a geometric thickness-two graph with $6n-18$ edges.

Open Question 1. *Does there exist a geometric thickness-two graph with* n *vertices and* $6n-18$ *edges?*

In Section 3 we proved that any graph G with $\overline{\theta}(G) > \theta(G)$ must have at least 10 vertices. We constructed such a graph G with 10 vertices and 42 edges, where $\overline{\theta}(G) = 3 > \theta(G) = 2$. An interesting question is whether this inequality can be established for graphs with fewer edges.

Open Question 2. *Does there exist a graph* G *with fewer than* 42 *edges satisfying the inequality* $\overline{\theta}(G) > \theta(G)$*?*

In Section 4 we proved that recognizing geometric thickness-two graphs is NP-hard, which settles an open question posed in [6]. Therefore, it seems natural to examine whether geometric thickness can be approximated efficiently, a question also posed in [6, 13]. Note that it is possible to determine the thickness within a constant factor [13].

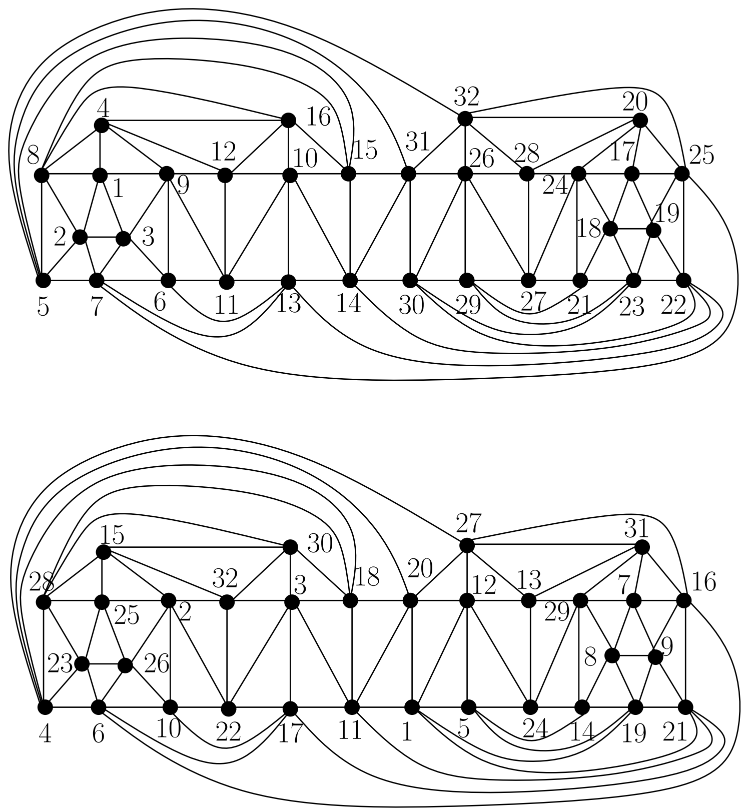

Figure 9: Two planar layers of a 11-regular 32-vertex graph with thickness two.

Open Question 3 ([6, 13]). *Does there exist a constant-factor approximation algorithm to determine geometric thickness?*

In Section 5 we proved the NP-hardness of coloring arbitrary geometric thickness-t graphs with $4t-1$ colors, which is particularly interesting since no graph with geometric thickness t is known that requires more than $4t$ colors for its proper coloring. Improving our complexity bound would be an interesting avenue to explore.

Open Question 4. *What is the complexity of coloring a geometric thickness-t graph with* $4t$ *colors? Does every geometric thickness-t graph admit a proper* $4t$*-coloring?*

Every planar graph contains a vertex of degree at most five, which is also the best possible since there exists 5-regular planar graphs. This upper bound on the minimum degree of planar graphs leads to a simple algorithm for constructing 6-colorings of planar graphs [17]. Ringel [22] observed that the average degree of the graphs with thickness-t is less than $6t$, which implies that every thickness-t graph contains a vertex of degree at most $6t - 1$. Consequently, the chromatic number of geometric thickness-t graphs is $6t$. For example, any graph G with $\theta(G) = 2$ must contain a vertex with degree at most 11, which leads to an algorithm for constructing 12-colorings of thickness-two graphs. In Figure 9 we show that this upper bound of 11 on the smallest degree is tight by constructing an 11-regular thickness-two graph. While this article was under review, Duncan [9] proved that there exist $(6t-1)$-regular graphs with thickness t. He also proved that there exist $5t$-regular graphs with geometric thickness at most t.

We believe that for $t \geq 2$, the upper bound on the minimum degree of geometric thickness-t graphs is less than the bound for thickness-t graphs. While there exist 11-regular thickness-two graphs, no geometric thickness-two graph is known with minimum degree greater than 7.

Open Question 5. *What is the smallest integer k such that every geometric thickness-t graph contains a vertex of degree at most k?*

Acknowledgements. We thank the anonymous reviewers for their suggestions that helped improve the presentation of the paper.

References

[1] Aichholzer, O., Aurenhammer, F., Krasser, H.: Enumerating order types for small point sets with applications. Order 19(3), 265–281 (2002)

[2] Alekseev, V.B., Gonchakov, V.S.: Thickness of arbitrary complete graphs. Math USSR Sbornik 30(2), 187–202 (1976)

[3] Appel, K., Haken, W.: Every planar map is four colorable. Part I. Discharging. Illinois Journal of Mathematics 21, 429–490 (1977)

[4] Barát, J., Matoušek, J., Wood, D.R.: Bounded-degree graphs have arbitrarily large geometric thickness. Electronic Journal of Combinatorics 13 (2006)

[5] Beineke, L.W., Harary, F., Moon, J.W.: On the thickness of the complete bipartite graph. Math. Proc. of the Cambridge Philosophical Society 60, 1–6 (1964)

[6] Brandenburg, F.J., Eppstein, D., Goodrich, M.T., Kobourov, S.G., Liotta, G., Mutzel, P.: Selected open problems in graph drawing. In: Proceedings of the International Symposium on Graph Drawing. LNCS, vol. 2912, pp. 515–539. Springer (2003)

[7] Dillencourt, M.B., Eppstein, D., Hirschberg, D.S.: Geometric thickness of complete graphs. Journal of Graph Algorithms and Applications 4(3), 5–17 (2000)

[8] Dujmovic, V., Wood, D.R.: Graph treewidth and geometric thickness parameters. Discrete & Computational Geometry 37(4), 641–670 (2007)

[9] Duncan, C.: Maximizing the degree of (geometric) thickness-t regular graphs. In: Proceedings of the International Workshop on Graph Drawing & Network Visualization. Springer, To Appear (2015)

[10] Duncan, C.A.: On graph thickness, geometric thickness, and separator theorems. Computational Geometry: Theory and Applications 44(2), 95–99 (2011)

[11] Duncan, C.A., Eppstein, D., Kobourov, S.G.: The geometric thickness of low degree graphs. In: Proceedings of the International Symposium on Computational Geometry. pp. 340–346. ACM (2004)

[12] Durocher, S., Gethner, E., Mondal, D.: Thickness and colorability of geometric graphs. In: Proceedings of the International Workshop on Graph-Theoretic Concepts in Computer Science. LNCS, vol. 8165, pp. 237–248. Springer (2013)

[13] Eppstein, D.: Separating thickness from geometric thickness. Contemporary Mathematics 342, 75–86 (2004)

[14] Erten, C., Kobourov, S.G.: Simultaneous embedding of planar graphs with few bends. Journal Graph Algorithms and Applications 9(3), 347–364 (2005)

[15] Giacomo, E.D., Liotta, G.: Simultaneous embedding of outerplanar graphs, paths, and cycles. International Journal of Computational Geometry & Applications 17(2), 139–160 (2007)

[16] Hutchinson, J.P., Shermer, T.C., Vince, A.: On representations of some thickness-two graphs. Comp. Geom.: Theory and Applications 13(3), 161–171 (1999)

[17] Jensen, T.R., Toft, B.: Graph coloring problems. Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons Inc., New York (1995)

[18] Kainen, P.C.: Thickness and coarseness of graphs. Abhandlungen aus dem Mathematischen Seminar der Univ. Hamburg 39(88–95) (1973)

[19] Mansfield, A.: Determining the thickness of a graph is NP-hard. In: Mathematical Proceedings of the Cambridge Philosophical Society. vol. 93, pp. 9–23 (1983)

[20] McGrae, A.R.A., Zito, M.: The complexity of the empire colouring problem. Algorithmica 68(2), 483–503 (2014)

[21] Mutzel, P., Odenthal, T., Scharbrodt, M.: The thickness of graphs: A survey. Graphs Combin 14, 59–73 (1998)

[22] Ringel, G.: Fabungsprobleme auf flachen und graphen. VEB Deutscher Verlag der Wissenschaften (1950)

[23] Tutte, W.T.: The thickness of a graph. Indag. Math 25, 567–577 (1963)

Appendix A.

In the proof of Theorem 2 we claimed that at most six vertices of Γ' can be straight-line visible to a common point, where Γ' is a geometric thickness two drawing of K_9 minus an edge. In this section we give a detailed proof for this claim. Suppose for a contradiction that we can insert a vertex v such that it can be straight-line visible to seven distinct vertices of Γ'.

Delete all the black edges, i.e., the edges common to both triangulations, from Γ'. Let Γ'' denote the resulting drawing. Figures 5(a) and (b) show the candidate drawings that are obtained from Figures 3(a) and (c), respectively. We do not examine Figure 3(b) separately since its closed regions are similar to that of Figure 3(a). In each planar layer of Γ'', v must lie on some bounded region or on the unbounded region. Observe that the bounded regions in each planar layer are a collection of triangles and quadrangles. If v lies interior to an empty triangle in each planar layer, then it can have at most six adjacent edges. Therefore, v must lie on the unbounded region or on a quadrangle of some planar layer. Each such region Q is unfilled in Figure 5. We examine each candidate region Q, and show that Q cannot contain any point v which is straight-line visible to seven distinct vertices.

We distinguish two main cases depending on whether Q is a region in Figure 5(a) (Case A), or Figure 5(b) (Case B).

Case A (Q belongs to Figure 5(a)): We consider several sub-cases, as follows.

Case A1 ($\boldsymbol{Q = idhg}$) : In this case one can partition Q into 5 regions $Q_1, Q_2, \ldots, Q_5$, depending on how the neighboring segments of h intersects Q. Let q be a point in Q.

If $q \in Q_1 \cup Q_2$, then since $\{d, g\}$ and $\{i\}$ lie on the opposite sides of the line determined by segment hb, vertex g cannot be visible to q. Therefore, q can see only 3 points in the blue layer. Observe now that q is either inside a triangle in the red layer, or inside the quadrangle $fche$.

If q is inside a triangle in the red layer, then q can see at most six vertices in total, and hence assume that q lies inside $fche$. Since $\{f, h\}$ and $\{e, c\}$ lie on the opposite sides of the line determined by segment ig, vertex h cannot be visible to q. Therefore, also in this case q can see at most six vertices in total. The case when $q \in Q_4 \cup Q_5$ can be analyzed in a similar fashion, where d and g cannot be visible to q simultaneously, and similarly h and f cannot be visible to q simultaneously.

In the remaining case we have $q \in Q_3$, where q can see all the vertices $\{i, g, h, d\}$ in the blue layer. Since q is on the outer face in the red layer, only the vertices $\{b, c, i, h\}$ are visible to q. Therefore, q has straight-line visibility to at most six points in total.

Case A2 ($\boldsymbol{Q = fdha}$) : In this case v either lies in $\Delta dbf \cap Q$, or in $\Delta dch \cap Q$, or outside of Δbdc.

If $v \in \Delta dbf \cap Q$, then v can see either $\{d, b, g\}$ or $\{f, b, g\}$ in the red layer. Since both of these sets contain a vertex common to $\{f, d, h, a\}$, vertex v cannot have straight-line visibility to more than six distinct vertices.

If $v \in \Delta dch \cap Q$, then v can have straight-line visibility to one of the following sets in the red layer: $\{d, c, g\}$, $\{d, b, g\}$, $\{f, b, g\}$, $\{f, g, c\}$, $\{f, e, h, c\}$, and $\{f, e, g\}$. Each of these sets except $\{f, e, h, c\}$ contains a vertex common to $V(Q) = \{f, d, h, a\}$, whereas the set $\{f, e, h, c\}$ contains two vertices that are common to $V(Q) = \{f, d, h, a\}$. Therefore, vertex v cannot have straight-line visibility to more than six distinct vertices.

If v is outside of Δbdc, then it is in the unbounded face in red layer. This case is considered later in Case A5.

Case A3 ($Q = fehc$) : Consider first the case where v lies in $\Delta iae \cap Q$. Here v can have straight-line visibility to one of the following sets in the blue layer: $\{f, d, i\}$, $\{i, d, h, g\}$, $\{f, d, h, a\}$, $\{a, h, g\}$, $\{g, i, e\}$, and $\{a, g, e\}$. Each of these sets of cardinality three has a vertex common to $V(Q) = \{f, e, h, c\}$. The set $\{f, d, h, a\}$ has two vertices common to $V(Q) = \{f, e, h, c\}$. Therefore, v cannot see any of these sets and also have straight-line visibility to more than six distinct vertices. The remaining set is $\{i, d, h, g\}$, but by Case A1, v cannot belong to the quadrangle $idhg$.

If v is outside of Δiae, then v is in the unbounded face in the blue layer. This case is considered later in Case A4.

Case A4 (Q is the unbounded face in the blue layer): Let H_{ae} be the half-plane containing i, which is determined by the line through a and e. Similarly, let H_{be} be the half-plane containing i, which is determined by the line through b and e. We distinguish two cases depending on whether $v \in H_{ae} \cap H_{be}$ or not.

Consider first the case when $v \in H_{ae} \cap H_{be}$. If v is also in the half-plane H_{bf}, which contains $\{a, d, g\}$, then v cannot see h in any layer. Note that in this scenario v is also in the unbounded face of the red layer. Since there are only seven vertices $\{a, b, c, d, e, h, i\}$ in the unbounded faces of the red and blue layers, v cannot have straight-line visibility to more than six vertices.

We may thus assume that v lies inside the quadrangle $bfhe$. In this scenario v can see one of the following sets in the red layer: $\{b, f, h\}$, $\{b, h, c, i\}$, $\{h, e, f, c\}$. The sets $\{b, f, h\}$ and $\{b, h, c, i\}$ contains one and two vertices that are common to $V(Q) = \{a, e, b, i\}$, respectively. Therefore, v cannot see any of these sets and also have straight-line visibility to more than six distinct vertices. The remaining set is $\{h, e, f, c\}$, where $e \in V(Q)$. Therefore, it suffices to show that f cannot be straight-line visible to v. Since $v \in Q$, it lies in Δbie. Note that $\Delta bie \subset \Delta bhe \subset \Delta bhe$. Therefore, the edge vf must cross either edge eh or bh. Both eh or bh belong to the red layer. Since f is interior to the blue layer, vf must belong to the red layer, which implies that vf cannot exist.

Consider now the case when v is outside of $H_{ae} \cap H_{be}$. If v is also on the unbounded face of the red layer, then v have to see all the seven vertices $\{a, b, c, d, e, h, i\}$. Note that h and d can be visible to v only in the red layer. Therefore, in this scenario v must lie on the vertically opposite angles of $\angle dch$. Since i lies on the vertically opposite angles of $\angle dhc$, the visibility from v to h must be blocked by the edge ci of the red layer.

If v is outside of $H_{ae} \cap H_{be}$, but not on the unbounded face of the red layer, then v can see one of the following sets in the red layer: $\{d, g, c\}$, $\{f, g, c\}$, $\{f, e, h, c\}$. Since $v \in Q$, v can see only the vertices $\{a, e, b\}$ in the blue layer. Therefore, v can see at most six distinct vertices altogether.

Case A5 (Q is the unbounded face in the red layer): Let H_{bd} be the half-plane containing a, which is determined by the line through b and d. Similarly, let H_{dc} be the half-plane containing a, which is determined by the line through d and c. We distinguish two cases depending on whether $v \in H_{bd} \cup H_{dc}$ or not. We can concentrate only on the case when v is not in the unbounded face of the blue layer (otherwise, we refer to Case A4).

If $v \in H_{bd} \cup H_{dc}$, then v can see one of the following sets in the blue layer: $\{a, i, f\}$, $\{a, f, d, h\}$, $\{a, h, g\}$, $\{a, g, e\}$. Since $\Delta abi \subset \Delta abe \subset \Delta abc$, only the vertices $\{b, d, c\}$ are visible to v in the red layer. Therefore, at most six distinct vertices can be visible to v.

If v is outside of $H_{bd} \cup H_{dc}$, then v can see one of the following sets in the blue layer: $\{a, i, f\}$, $\{f, i, d\}$, $\{d, i, g, h\}$, $\{i, g, e\}$. Since $\Delta bci \subset \Delta bch \subset \Delta bcd$, only the vertices $\{b, h, c, i\}$ are visible to v in the red layer. Therefore, at most six distinct vertices can be visible to v.

Case B (Q belongs to Figure 5(b)): We consider several sub-cases, as follows.

Case B1 ($Q = igfe$) : Note that $\Delta bgc \subset \Delta bec \subset \Delta bfc$. Therefore, if $v \in \Delta egb$, then v can see one of the following sets in the red layer: $\{b, c, e, g\}$, $\{b, e, h\}$, $\{b, i, h\}$. The corresponding point sets that are visible to v from the blue layer are $\{e, f, g, i\}$, $\{e, g, i\}$, $\{e, g, i\}$, respectively. Therefore, v can see at most six distinct vertices in total.

Otherwise, if $v \in \Delta efg$, then v can see one of the following sets in the red layer: $\{b, c, e, g\}$, $\{c, e, h\}$, $\{b, e, h\}$, $\{b, i, h\}$, $\{i, h, c\}$, $\{i, c, d\}$, $\{i, f, d\}$. For the set $\{b, c, e, g\}$, the points that are visible in the blue layers are $\{i, e, g, f\}$, which implies visibility to only six distinct vertices. For the remaining sets, only three more points are visible in the blue layers, i.e., $\{e, g, f\}$.

Case B2 ($Q = aefh$) : If v is in the unbounded face in the red layer, then we refer to Case B5. Otherwise, v can lie in Δbef or in Δcfh.

If $v \in \Delta bef$, then v can see one of the following sets in the red layer: $\{b, i, f\}$, $\{d, i, f\}$, $\{i, c, d, f\}$, $\{i, c, h\}$, $\{b, h, i\}$, $\{b, e, h\}$. Since $\Delta ebc \subset \Delta fbc$ and e, h to the opposite sides of the edge fg, v cannot see both e and h simultaneously in the blue layer. Therefore, the set visible to v in the blue layer is $\{a, e, f\}$. Hence the number of distinct points visible to v is at most six.

If $v \in \Delta cfh$, then v can see one of the following sets in the red layer: $\{c, d, i, f\}$, $\{d, i, f\}$, $\{i, c, h\}$. Since $\Delta ebc \subset \Delta fbc$ and e, h to the opposite sides of the edge fg, v cannot see both e and h simultaneously in the blue layer. Therefore, the set visible to v in the blue layer is $\{a, h, f\}$. Hence the number of distinct points visible to v is at most six.

Case B3 ($Q = icfd$) : If v is in the unbounded face in the blue layer, then we refer to Case B4. Otherwise, v can lie in Δadf or in Δdgi.

If $v \in \Delta adf$, then v can see one of the following sets in the blue layer: $\{a, h, d\}$, $\{a, e, f, h\}$. Since $\Delta ifd \subset \Delta ifc$, the set visible to v in the red layer is $\{f, d, c\}$. Hence the number of distinct points visible to v is at most six.

If $v \in \Delta dgi$, then v can see one of the following sets in the blue layer: $\{d, g, h\}$, $\{a, d, h\}$, $\{a, e, f, h\}$, $\{f, g, h\}$, $\{i, g, e, f\}$, $\{a, g, i\}$. Since $\Delta ifd \subset \Delta ifc$, the set visible to v in the red layer is $\{i, d, c\}$. Hence the number of distinct points visible to v is at most six, except when we take the union of the sets $\{a, e, f, h\}$ and $\{i, d, c\}$. However, by Case B2, v cannot lie in the quadrangle $aefh$.

Case B4 (Q is the unbounded face in the blue layer): Since the vertices on the unbounded face of the blue layer are $\{a, g, d, b\}$, v cannot lie inside any red triangle that is adjacent to b or c. Therefore, v must lie in the quadrangle $icdf$ or in the unbounded face of both red and blue layer.

If v lies in the the quadrangle $icdf$, then since $\Delta abd \subset \Delta abc$, v can see only the vertices $\{a, b, d\}$ in the blue layer, which implies that v is visible to at most six distinct vertices.

Assume now that v is in the unbounded face of both red and blue layer. Then to make visible adjacent to seven vertices, v must see all the vertices $\{a, b, c, d, e, f, g\}$. Since e, f are internal vertices in the blue layer, they must be seen from the red layer. Therefore, v must lie in the angle vertically opposite to $\angle ecf$. Note that $\Delta bgf \subset \Delta bcf$, and e is lies interior to Δbgf inside the $\angle gcf$. Therefore, the visibility from v to e must be blocked by the edge cg.

Case B5 (Q is the unbounded face in the red layer): Let H_{bf} be the half-plane containing a, which is determined by the line through b and f. Similarly, let H_{fc} be the half-plane containing a, which is determined by the line through f and c. We distinguish two cases depending on whether $v \in H_{bf} \cup H_{fc}$ or not. We can concentrate only on the case when v is not in the unbounded face of the blue layer (otherwise, we refer to Case B4).

If $v \in H_{bf} \cup H_{fc}$, then v can see one of the following sets in the blue layer: $\{a, i, g\}$, $\{a, i, e\}$, $\{a, e, f, h\}$, $\{a, h, d\}$. Since $\Delta bcg \subset \Delta bce \subset \Delta bcf$, only the vertices $\{b, f, c\}$ are visible to v in the red layer. Therefore, at most six distinct vertices can be visible to v.

If v is outside of $H_{bf} \cup H_{fc}$, then v can see one of the following sets in the blue layer: $\{a, i, g\}$, $\{i, g, e, f\}$, $\{g, f, h\}$, $\{g, h, d\}$, $\{g, h, d\}$. Since $\Delta bcg \subset \Delta bce \subset \Delta bcf$, only the vertices $\{b, c, e, g\}$ are visible to v in the red layer. Therefore, at most six distinct vertices can be visible to v.

Appendix B.

Let Γ be a thickness-two drawing of $K_9 - e$. Assume that the edges of one layer of Γ are assigned red color, and the remaining edges are assigned blue color. Observe that each of the unsaturated vertices $\{d, e\}$ of Γ is enclosed by a cycle of distinct color, as follows. If Γ corresponds to Figure 3(a) or (b), then d is enclosed inside a blue cycle $C(d) = (a, f, i, g, a)$, while e is enclosed inside a red cycle $C(e) = (c, f, h, c)$. If Γ corresponds to Figure 3(c), then d is enclosed inside a red cycle $C(d) = (c, i, f, c)$, while e is inside a blue cycle $C(e) = (a, i, g, h, a)$. Note that d lies outside of $C(e)$, while e lies outside of $C(d)$. Furthermore, d and e are adjacent to some vertex of $C(d)$ and $C(e)$, respectively. Therefore, we can represent these configurations as shown in Figures B.1(a)–(b). Let $[C(d)]$ and $[C(e)]$ be the closed interior of the cycle $C(d)$ and $C(e)$, respectively. As illustrated in Figures B.1(a)–(b), we have the following observation.

Fact 2. *Let H be a graph obtained by adding a vertex v to the unsaturated vertices d, e of the graph $K_9 - e$. Then any thickness-two drawing of graph H satisfies the following properties:*

- *The vertex v lies either inside $[C(d)] \cap [C(e)]$, or outside of $[C(d)] \cup [C(e)]$.*
- *The edges (v, e) and (v, d) must lie in different layers.*

Let $G_1, G_2, \ldots, G_k$ be $k \geq 9$ copies of $K_9 - e$. Let d_1, e_1 be the unsaturated vertices of G_1. For each G_j, $j > 1$, make d_1 adjacent to some unsaturated vertex d_j of G_j. Refer to the remaining unsaturated vertex of G_j as e_j. Add a vertex v and make v adjacent to all the unsaturated vertices of $G_1, G_2, \ldots, G_k$. Let $\mathcal{H}_k$ denote the resulting graph, which we refer to as a *rigid graph*.

Lemma 1. *Let $\mathcal{H}_k$ be a rigid graph. Then in any thickness-two drawing Γ of $\mathcal{H}_k$, the subgraph G' induced by the edges (v, d_1), (v, d_j) and (d_1, d_j), where $1 < j \leq k$, lies in the same layer.*

Proof. To prove that the edges of G' lies in the same layer, we show that for every $j > 1$, the edges of the triangle v, d_1, d_j lie in the same layer.

Let D_i be the drawing of G_i in Γ. Consider the graphs G_1, G_j. Since v can see both d_1 and e_1, by Fact 2, v lies either inside $[C(d_1)] \cap [C(e_1)]$ or outside of $[C(d_1)] \cup [C(e_1)]$. Similarly, since v can see both d_j and e_j, by Fact 2, v lies either inside $[C(d_j)] \cap [C(e_j)]$ or outside of $[C(d_j)] \cup [C(e_j)]$. Without loss of generality assume that $C(d_1)$ is a blue cycle in Γ.

First assume that v lies inside $[C(d_1)] \cap [C(e_1)]$. By Facts 1 and 2, the vertices v, d_j and e_j are connected both in red and blue layers. Therefore, d_j and e_j would also lie inside $[C(d_1)] \cap [C(e_1)]$, i.e., $\{v, d_j, e_j\} \in [C(d_1)] \cap [C(e_1)]$. We now consider Cases 1–2 depending on whether $C(d_j)$ is blue or red.

Case 1 ($C(d_j)$ is blue): Since $\{v, d_j, e_j\} \in [C(d_1)] \cap [C(e_1)]$ and d_j is adjacent to some vertex of $C(d_j)$ through some blue edge, the cycle $C(d_j)$ cannot enclose $C(d_1)$. Furthermore, since the cycle $C(d_1)$ is blue and d_1

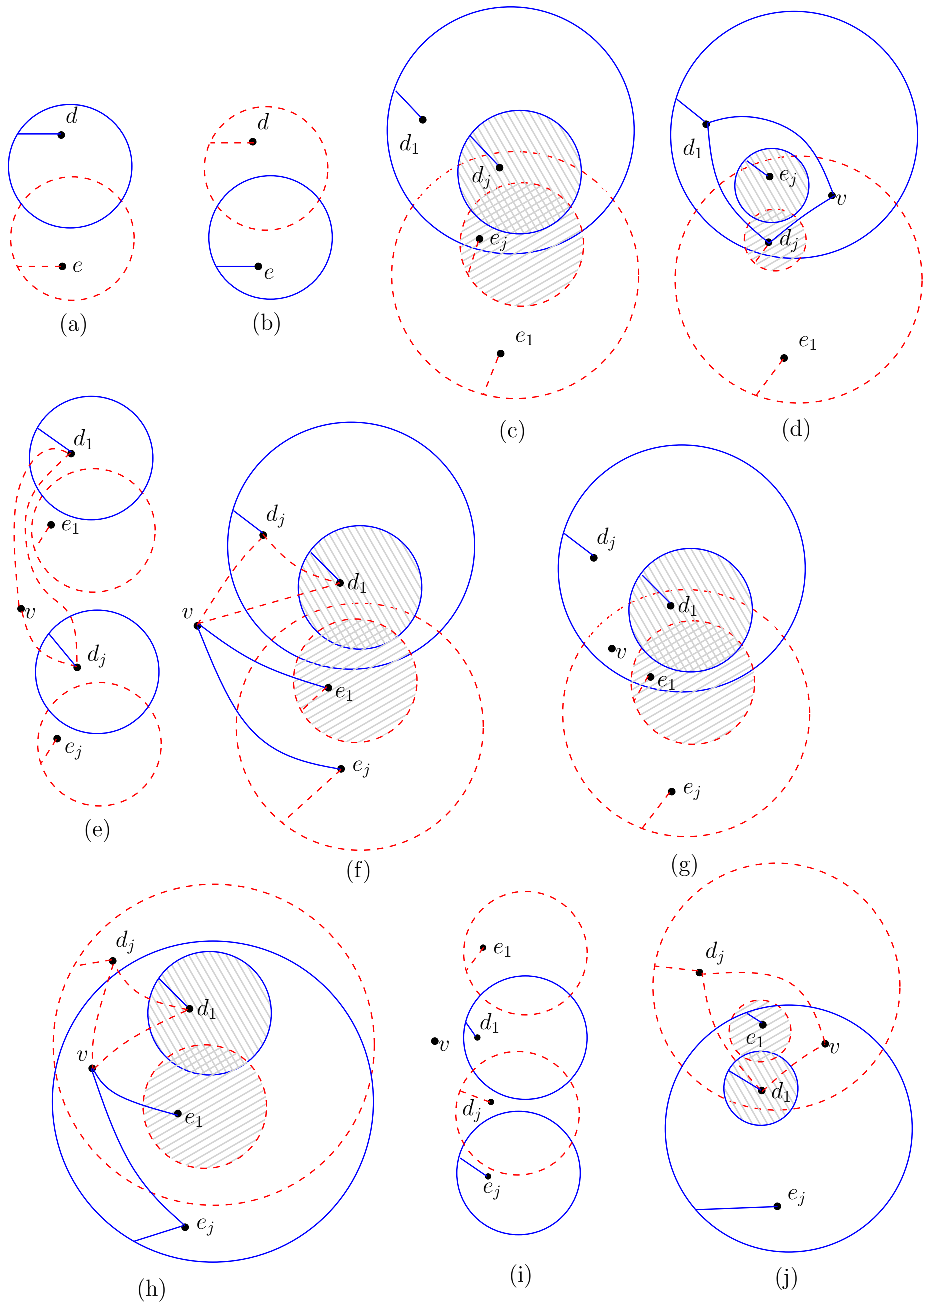

Figure B.1: (a)–(b) Schematic geometric thickness-two representations of K_9'. (c)–(j) Illustration of the proof of Lemma 1.

is adjacent to some vertex of $C(d_1)$ through some blue edge, $C(d_j)$ cannot enclose d_1. Therefore, $C(d_j)$ must lie interior to $C(d_1)$. In this scenario, the edge (d_1, d_j) crosses both a blue cycle and a red cycle, i.e., to reach d_j from d_1, we need to enter both the blue cycle $C(d_j)$ and the red cycle $C(e_1)$, which contradicts that Γ is a thickness-two drawing. Figure B.1(c) illustrates such a scenario.

Case 2 ($C(d_j)$ is red): Since $\{v, d_j, e_j\} \in [C(d_1)] \cap [C(e_1)]$, the cycle $C(d_j)$ must lie inside the red circle $C(e_1)$. Since d_1 lies outside of $C(e_1)$, the cycle $C(d_j)$ does not enclose d_1. Hence the edge (d_1, d_j) must cross $C(d_j)$, and therefore, (d_1, d_j) would lie in the blue layer. Similarly, the edge (v, d_1) crosses $C(e_1)$ and hence lies in the blue layer. Finally, observe that v cannot lie inside $[C(d_j)] \cup [C(e_j)]$. Therefore, the edge (v, d_j) crosses $C(d_j)$, and hence lies in the blue layer. Figure B.1(d) illustrates such a scenario.

Assume now that v lies outside of $[C(d_1)] \cup [C(e_1)]$. By Facts 1 and 2, the vertices v, d_j and e_j are connected both in red and blue layers. Therefore, d_j, e_j must lie outside of $[C(d_1)] \cup [C(e_1)]$. We now consider Cases 3–4 depending on whether $C(d_j)$ is blue or red.

Case 3 ($C(d_j)$ is blue): Since the cycle $C(d_1)$ is blue, $C(d_j)$ must lie outside of $C(d_1)$, i.e., either $C(d_1)$ and $C(d_j)$ are interior disjoint, or $[C(d_1)] \subset [C(d_j)]$. Note from the analysis of Case 1 that the case when $[C(d_1)] \subset [C(d_j)]$ and $v \in [C(d_j)] \cap [C(e_j)]$ cannot appear. On the other hand, if v is outside of $[C(d_j)] \cap [C(e_j)]$, then (d_1, d_j) crosses $C(d_1)$, which implies that (d_1, d_j) must be red. Since (v, d_1) and (v, d_j) both cross the cycle $C(d_1)$, they must lie on the red layer. See Figure B.1(f).

The case when $C(d_1)$ and $C(d_j)$ are interior disjoint is illustrated Figure B.1(e). Here the edge (d_1, d_j) crosses $C(d_1)$, which implies that (d_1, d_j) must be red. Since v lies outside of $[C(d_1)] \cup [C(e_1)]$, the edge (v, d_1) is also red. We now examine the edge (v, d_j). By Fact 2, v lies either inside $[C(d_j)] \cap [C(e_j)]$ or outside of $[C(d_j)] \cup [C(e_j)]$. If v is outside of $[C(d_j)] \cup [C(e_j)]$, then (v, d_j) must be red. If $v \in [C(d_j)] \cap [C(e_j)]$, then an analysis similar to Case 1 shows that (d_1, d_j) must create an edge crossing in either red or blue layer. See Figure B.1(g).

Case 4 ($C(d_j)$ is red): First assume that $d_1 \in [C(d_j)]$. In this scenario v cannot lie outside of $[C(d_j)] \cup [C(e_j)]$, because the edge (v, d_1) then crosses both a red cycle $C(d_j)$ and a blue cycle $C(d_1)$. Therefore, by Facts 1 and 2, v, d_1, e_1 must lie inside $[C(d_j)] \cap [C(e_j)]$. An analysis similar to Case 2 shows that the triangle v, d_1, d_j lies in the same (red) layer. See Figure B.1(h).

If $d_1 \notin [C(d_j)]$, then let R be the region $[C(d_1)] \cap [C(d_j)]$. If v lies outside of $[C(d_j)] \cup [C(e_j)]$, then d_1 or d_j cannot lie in R. In this scenario the edge (d_1, d_j) cannot exist, i.e., to reach d_j from d_1 we have to cross both a red cycle $C(d_j)$ and a blue cycle $C(d_1)$. See Figure B.1(i). Now consider the case when v lies inside $[C(d_j)] \cap [C(e_j)]$. An analysis similar to Case 2 shows that the triangle v, d_1, d_j must be of red color. See Figure B.1(j).

□

Observe that G' is a bipartite graph $K_{2,k-1}$ with vertex-partition $\{v, d_1\}$, $\{d_2, \ldots, d_k\}$, plus the edge (v, d_1). We call the graph G' the *core graph* and the vertex v the *pole vertex*. By Facts 1 and 2 one can derive the following.

Fact 3. *Let Γ be a thickness-two drawing of some rigid graph $\mathcal{H}_k$. Then the pole and the unsaturated vertices of $\mathcal{H}_k$ are connected in each layer of Γ.*

Let H be a chain of three rigid graphs H_a, H_b, H_c, each a copy of $\mathcal{H}_{17}$. For any H_q, where $q \in \{a, b, c\}$, let $\{d_i^q, e_i^q\}$, where $1 \leq i \leq 17$, be the unsaturated vertices of H_q, and let v^q be the pole of H_q. We now add the edges $(e_2^a, e_2^b), (e_3^a, e_3^b), \ldots, (e_9^a, e_9^b)$ and $(e_2^c, e_{10}^b), (e_3^c, e_{11}^b), \ldots, (e_9^c, e_{17}^b)$. Let H' denote the resulting graph. Figures 6(a)–(b) illustrate a schematic representation of H'.

Lemma 2. *In any thickness-two drawing Γ of H', there exists a path from $v^a, \ldots, v^b, \ldots, v^c$ that lies on the same layer.*

Proof. Let G_a, G_b and G_c be the core graphs of H_a, H_b, H_c, respectively. By Lemma 1, all the edges of a core graph are of same color. Assume without loss of generality that G_a is red in Γ. We first prove that there exists a blue edge between the unsaturated vertices of H_a and H_b. We distinguish two cases depending on the color of G_b.

Case 1(G_b is red): Since G_a and G_b are red, they are either interior disjoint or one is interior to some inner face of the other.

Consider first the case when G_a and G_b are interior disjoint. In this case choose some j, where $2 \leq j \leq 9$, such that d_j^a and d_j^b are inner vertices in G_a and G_b, respectively. Since d_j^a is inside some red cycle C_a in G_a, by Fact 2, e_j^a must lie inside C_a. On the other hand, since d_j^b is inside some red cycle C_b in G_b, by Fact 2, e_j^b must lie inside C_b. Consequently, (e_j^a, e_j^b) must cross C_a, and this will be the required blue edge.

Consider now the case when one of G_a and G_b is in some inner face of the other. Without loss of generality assume that G_b lies inside some inner face f_a of G_a. Let C_a be a red cycle of G_a that does not contain f_a, but encloses an unsaturated vertex d_j^a in its interior, where $2 \leq j \leq 9$. Then the vertex e_j^a must lie interior to C_a. Since G_b lies inside f_a, the vertex e_j^b must lie interior to f_a. Hence (e_j^a, e_j^b) must cross C_a, and this will be the required blue edge.

Case 2(G_b is blue): In the following we show that this case cannot occur. Let C_a be some red cycle of G_a such that there exists a vertex d_p^a interior to

C_a and a vertex d^a_q outside of C_a, where $2 \leq p, q \leq 9$. If all the edges between H_a and H_b are red, then the vertex e^b_p must lie interior to C_a. Since e^b_q and e^b_p are connected in the red layer, e^b_q must also lie interior to C_a. Consequently, the edge (e^a_q, e^b_q) must cross C_a, which contradicts that all the edges between H_a and H_b are red.

Similarly, we can prove that there exists a blue edge between the unsaturated vertices of H_b and H_c. By Fact 3, we can find a path $v_a, \ldots, e^a_p$, $e^b_p \ldots, v_b, \ldots, e^b_{q+8}, e^c_q, v_c$ that lies on the same layer, where $2 \leq p, q \leq 9$. $\square$