On the Geometric Thickness of 2-Degenerate Graphs

Rahul Jain
FernUniversität in Hagen, Germany

Marco Ricci
FernUniversität in Hagen, Germany

Jonathan Rollin
FernUniversität in Hagen, Germany

André Schulz
FernUniversität in Hagen, Germany

Abstract

A graph is 2-degenerate if every subgraph contains a vertex of degree at most 2. We show that every 2-degenerate graph can be drawn with straight lines such that the drawing decomposes into 4 plane forests. Therefore, the geometric arboricity, and hence the geometric thickness, of 2-degenerate graphs is at most 4. On the other hand, we show that there are 2-degenerate graphs that do not admit any straight-line drawing with a decomposition of the edge set into 2 plane graphs. That is, there are 2-degenerate graphs with geometric thickness, and hence geometric arboricity, at least 3. This answers two questions posed by Eppstein [Separating thickness from geometric thickness. In *Towards a Theory of Geometric Graphs*, vol. 342 of *Contemp. Math.*, AMS, 2004].

2012 ACM Subject Classification Mathematics of computing → Graph theory; Theory of computation → Randomness, geometry and discrete structures

Keywords and phrases Degeneracy, geometric thickness, geometric arboricity

Digital Object Identifier 10.4230/LIPIcs.SoCG.2023.44

Related Version *Full Version*: https://arxiv.org/abs/2302.14721 [16]

1 Introduction

A graph is planar if it can be drawn without crossings on a plane. Planar graphs exhibit many nice properties, which can be exploited to solve problems for this class more efficiently compared to general graphs. However, in many situations, graphs cannot be assumed to be planar even if they are sparse. It is therefore desirable to define graph classes that extend planar graphs. Several approaches for extending planar graphs have been established over the last years [4, 14]. Often these classes are defined via drawings, for which the types of crossings and/or the number of crossings are restricted. A natural way to describe how close a graph is to being a planar graph is provided by the graph parameter *thickness*. The thickness of a graph G is the smallest number $\theta(G)$ such that the edges of G can be partitioned into $\theta(G)$ planar subgraphs of G. Related graph parameters are *geometric thickness* and *book thickness*. Geometric thickness was introduced by Kainen under the name *real linear thickness* [17]. The geometric thickness $\bar{\theta}(G)$ of a graph G is the smallest number of colors that is needed to find an edge-colored geometric drawing (i.e., one with edges drawn as straight-line segments) of G with no monochromatic crossings. For the book thickness $\mathsf{bt}(G)$, we additionally require that only geometric drawings with vertices in convex position are considered.

An immediate consequence from the definitions of thickness, geometric thickness and book thickness is that for every graph G we have $\theta(G) \leq \bar{\theta}(G) \leq \mathsf{bt}(G)$. Eppstein shows that the three thickness parameters can be arbitrarily "separated". Specifically, for any number k

there exists a graph with geometric thickness 2 and book thickness at least k [10] as well as a graph with thickness 3 and geometric thickness at least k [11]. The latter result is particularly notable since any graph of thickness k admits a k-edge-colored drawing of G with no monochromatic crossings if edges are not required to be straight lines. This follows from a result by Pach and Wenger [22], stating that any planar graph can be drawn without crossings on arbitrary vertex positions with polylines.

Related to the geometric thickness is the *geometric arboricity* $\bar{\mathsf{a}}(G)$ of a graph G, introduced by Dujmović and Wood [6]. It denotes the smallest number of colors that are needed to find an edge-colored geometric drawing of G without monochromatic crossings where every color class is acyclic. As every such plane forest is a plane graph, we have $\bar{\theta}(G) \leq \bar{\mathsf{a}}(G)$. Moreover, every plane graph can be decomposed into three forests [24], and therefore $3\bar{\theta}(G) \geq \bar{\mathsf{a}}(G)$.

Bounds on the geometric thickness are known for several graph classes. Due to Dillencourt et al. [5] we have $\frac{n}{5.646} + 0.342 \leq \bar{\theta}(K_n) \leq \frac{n}{4}$ for the complete graph K_n. Graphs with bounded degree can have arbitrarily high geometric thickness. In particular, as shown by Barárt et al. [2], there are d-regular graphs with n vertices and geometric thickness at least $c\sqrt{d}n^{1/2-4/d-\varepsilon}$ for every $\varepsilon > 0$ and some constant c. However, due to Duncan et al. [8], if the maximum degree of a graph is 4, its geometric thickness is at most 2. For graphs with treewidth t, Dujmović and Wood [6] showed that the maximum geometric thickness is $\lceil t/2 \rceil$. Hutchinson et al. [15] showed that graphs with n vertices and geometric thickness 2 can have at most $6n - 18$ edges. As shown by Durocher et al. [9], there are n-vertex graphs for any $n \geq 9$ with geometric thickness 2 and $2n - 19$ edges. In the same paper, it is proven that it is NP-hard to determine if the geometric thickness of a given graph is at most 2. Computing thickness [18] and book thickness [3] are also known to be NP-hard problems. For bounds on the thickness for several graph classes, we refer to the survey of Mutzel et al. [19]. A good overview on bounds for book thickness can be found on the webpage of Pupyrev [23].

A graph G is *d-degenerate* if every subgraph contains a vertex of degree at most d. So we can repeatedly find a vertex of degree at most d and remove it, until no vertices remain. The reversal of this vertex order (known as a *degeneracy order*) yields a construction sequence for G that adds vertex by vertex and each new vertex is connected to at most d previously added vertices (called its *predecessors*). Adding a vertex with exactly two predecessors is also known as a Henneberg 1 step [12]. In particular, any 2-degenerate graph is a subgraph of a Laman graph (i.e., a graph that is generically minimal rigid), however not every Laman graph is 2-degenerate. All d-degenerate graphs are (d, ℓ)-sparse, for any $\binom{d+1}{2} \geq \ell \geq 0$, that is, every subgraph on n vertices has at most $dn - \ell$ edges.

Our Results. In this paper, we study the geometric thickness of 2-degenerate graphs. Due to the Nash-Williams theorem [20, 21], every 2-degenerate graph can be decomposed into 2 forests and hence has arboricity at most 2 and therefore thickness at most 2. On the other hand, as observed by Eppstein [10], 2-degenerate graphs can have unbounded book thickness. Eppstein's examples of graphs with thickness 3 and arbitrarily high geometric thickness are 3-degenerate graphs [11]. Eppstein asks whether the geometric thickness of 2-degenerate graphs is bounded by a constant from above and whether there are 2-degenerate graphs with geometric thickness greater than 2. The currently best upper bound of $O(\log n)$ follows from a result by Duncan for graphs with arboricity 2 [7]. We improve this bound and answer both of Eppstein's questions with the following two theorems.

▶ **Theorem 1.** *For each 2-degenerate graph G we have $\bar{\theta}(G) \leq \bar{\mathsf{a}}(G) \leq 4$.*

▶ **Theorem 2.** *There is a 2-degenerate graph G with $\bar{\mathsf{a}}(G) \geq \bar{\theta}(G) \geq 3$.*

2 Proof of Theorem 1: The upper bound

In this section, we prove Theorem 1. To this end, we describe, for any 2-degenerate graph, a construction for a straight-line drawing such that the edges can be colored using four colors, avoiding monochromatic crossings and monochromatic cycles. This shows that 2-degenerate graphs have geometric arboricity, and hence geometric thickness, at most four.

Before we give a high-level description of the construction we introduce some definitions. For a graph G we denote its edge set with $E(G)$ and its vertex set with $V(G)$. Consider a 2-degenerate graph G with a given, fixed degeneracy order. We define the *height* of a vertex $\mathrm{height}(v)$ as the length t of a longest path $u_0 \cdots u_t$ with $u_t = v$ such that for each $1 \leq i \leq t$ vertex u_{i-1} is a predecessor of u_i. The set of vertices of the same height is called a *level* of G. By definition, each vertex has at most two neighbors of smaller height.

Our construction process embeds G level by level with increasing height. The levels are placed alternately either strictly below or strictly to the right of the already embedded part of the graph. If a level is placed below, then we use specific colors v and vs (short for “vertical” and “vertical slanted”, respectively) for all edges between this level and levels of smaller height. Similarly, we use specific colors h and hs (short for “horizontal” and “horizontal slanted”, respectively) if a level is placed to the right. See Figure 1 (right).

To make our construction work, we need several additional constraints to be satisfied in each step which we will describe next. For a point p in the plane, we use the notation $\mathsf{x}(p)$ and $\mathsf{y}(p)$ to refer to the x- and y-coordinates of p, respectively. Consider a drawing D of a 2-degenerate graph G of height k together with a coloring of the edges with colors $\{\mathsf{h}, \mathsf{hs}, \mathsf{v}, \mathsf{vs}\}$. For the remaining proof, we assume that each vertex of G has either 0 or exactly 2 predecessors. If not, we add a dummy vertex without predecessors to the graph and make it the second predecessor of all those vertices that originally only had 1 predecessor. We say that D is *feasible* if it satisfies the following constraints:

(C1) For each vertex in G the edges to its predecessors are colored differently. If $k > 0$, then each vertex of height k in G is incident to one edge of color h and one edge of color hs.

(C2) There exists some $x_D \in \mathbb{R}$ such that for each vertex $v \in V(G)$ we have $\mathsf{x}(v) > x_D$ if and only if $\mathrm{height}(v) = k$.

(C3) There is no monochromatic crossing.

(C4) No two vertices of G lie on the same horizontal or vertical line.

(C5) Each $v \in V(G)$ is h-*open to the right*, that is, the horizontal ray emanating at v directed to the right avoids all h-edges.

(C6) Each $v \in V(G)$ is v-*open to the bottom*, that is, the vertical ray emanating at v directed downwards avoids all v-edges.

These constraints are schematized in Figure 1.

We now show how to construct a feasible drawing for G. We prove this using induction on the height of the graph. The base case $k = 0$ is trivial, as there are no edges in the graph. Assume that $k \geq 1$ and the theorem is true for all 2-degenerate graphs with height $k - 1$. Let H denote the subgraph of G induced by vertices with height less than k. By induction, there is a feasible drawing D of H.

As a first step, we reflect the drawing D at the straight line $x = -y$. Thus, a point (x, y) before transformation becomes $(-y, -x)$. Additionally, we swap the colors hs and vs as well as the colors h and v. Let D' denote the resulting drawing. From now on, all appearing coordinates of vertices refer to coordinates in D'. By construction, D' satisfies (C3–C6). Applying (C1) to D shows that in D' each vertex of height $k - 1$ is incident to one edge of color v and one edge of color vs. Applying (C2) to D shows that there exists $y_{D'} \in \mathbb{R}$ such that for each vertex $v \in V(H)$ we have $\mathsf{y}(x) < y_{D'}$ in D' if and only if $\mathrm{height}(v) = k - 1$.

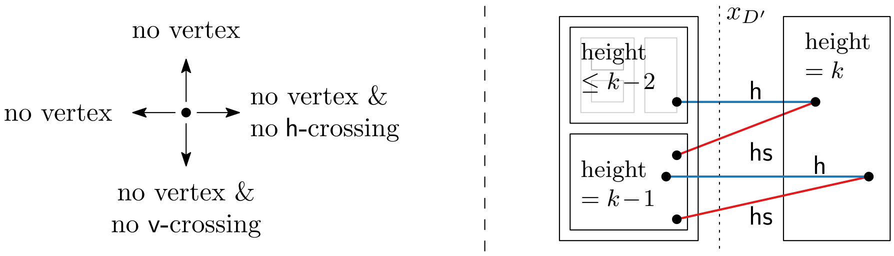

Figure 1 Left: For each vertex v in a feasible drawing, there are no other vertices on the vertical and the horizontal line through v. Moreover, v is h-open to the right and v-open to the bottom. Right: All vertices in the highest level (of height k) are placed to the right of all vertices of smaller height. Moreover, each vertex in that level is incident to one edge of color h and one edge of color hs.

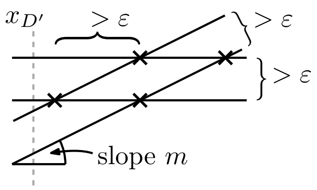

Figure 2 Horizontal lines intersecting straight lines of slope m. Conditions (ii–iv) are illustrated.

As the second (and last) step, we place the points of height k of G such that the resulting drawing is feasible. Let L_k denote the set of these vertices and let $x_{D'}$ denote the largest x-coordinate among all vertices in D'. Choose a sufficiently small slope m, with $m > 0$, and a sufficiently small ε, with $\varepsilon > 0$, such that the following holds.

(i) For any distinct u, $v \in V(H)$ with $\mathsf{y}(u) < \mathsf{y}(v)$, the horizontal line through v and the straight line through u with slope m intersect at a point p with $\mathsf{x}(p) > x_{D'}$.

(ii) For any distinct u, $v \in V(H)$ we have that $\varepsilon < |\mathsf{y}(u) - \mathsf{y}(v)|$.

(iii) For any distinct u, u', v, $v' \in V(H)$ let p be the intersection point of the straight line through u with slope m and the horizontal line through v and let p' be the intersection point of the straight line through u' with slope m and the horizontal line through v'. If $\mathsf{x}(p) \neq \mathsf{x}(p')$, then $\varepsilon < |\mathsf{x}(p) - \mathsf{x}(p')|$.

(iv) For any distinct u, $v \in V(H)$ we have that ε is smaller than the distance between the two straight lines of slope m through u and v, respectively.

The constraints are summarized in Figure 2. Such a choice of m and ε is possible, by choosing m according to Condition (i) first and then ε according to the Conditions (ii–iv).

For each vertex $w \in L_k$ let u and v be the two predecessors of w in H with $\mathsf{y}(u) < \mathsf{y}(v)$ and let p^w denote the intersection point of the straight line of slope m passing through u (called a slanted line) and the horizontal line passing through v. We will place w close to p^w and connect w to v using an edge of color h and we connect w to u using an edge of color hs. To determine the exact location of the vertices, we consider the horizontal lines through vertices $v \in V(H)$ from bottom to top (with increasing y-coordinate) and for each such line consider the intersections with slanted lines through vertices $u \in V(H)$ with $\mathsf{y}(u) < \mathsf{y}(v)$ from left to right (with increasing x-coordinate). Let $p_1, \ldots, p_t$ denote the intersection points in the order just described. For each intersection point p_i let ℓ_i denote the straight line through p_i with slope $-1/m$ (which is negative as $m > 0$), that is, ℓ_i is perpendicular to straight lines of slope m. Every vertex $w \in L_k$ with $p^w = p_i$ will be placed on ℓ_i at a certain distance from p^w (specified later). Note that there might be multiple points with the same

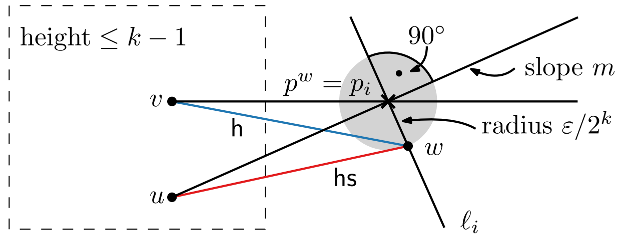

Figure 3 The placement of the k^{th} point w in order of vertices in L_k.

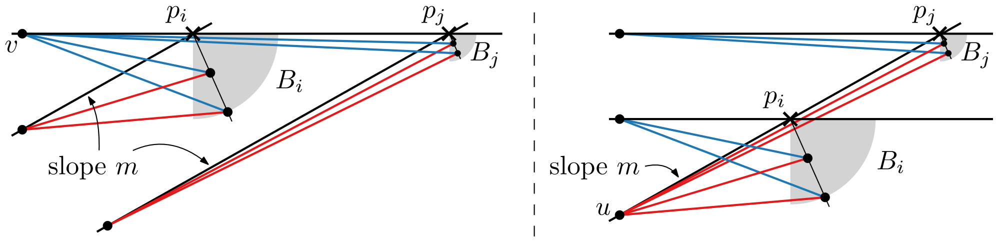

Figure 4 The placement of several points with a common “horizontal” predecessor v (left) or a common “slanted” predecessor u (right). Edges with color h are drawn blue, edges with color hs are drawn red.

predecessors and hence multiple vertices $w \in L_k$ with $p^w = p_i$. For each p_i we order all such vertices arbitrarily. This gives an ordering of all vertices in L_k based on the ordering $p_1, \ldots, p_t$. If w is the k^{th} vertex in this order, w is placed on ℓ_i to the bottom-right of p_i at distance $\varepsilon/2^k$ from p_i; see Figure 3. In this fashion, all vertices in L_k are placed with decreasing distance to their respective intersection point; see Figure 4.

We call the resulting drawing D_G. We claim that D_G demonstrates that the geometric arboricity of G is at most four.

Vertices on distinct points, edges intersect in at most one point in D_G. For each $i \leq t$ let δ_i denote the distance between p_i and the first vertex w placed close to p_i. Then $\delta_i \leq \delta_{i-1}/2$ for each i with $1 < i \leq t$. For each $i \leq t$ let B_i be the region formed by all points $q \in \mathbb{R}$ of distance at most δ_i to p_i with $\mathsf{x}(q) > \mathsf{x}(p_i)$ and $\mathsf{y}(q) < \mathsf{y}(p_i)$ (B_i is a quarter of a disk). Then all vertices $w \in L_k$ with $p^w = p_i$ are placed on distinct points along the intersection of the line ℓ_i with B_i; see Figure 4.

Due to Conditions (ii) and (iv), all the regions B_i are disjoint. By construction, no two vertices are placed on the same point within a region B_i. This shows that no two vertices in G are placed on the same point in D_G. Moreover, for the same reasons, for each vertex $v \in V(H)$ the edges between v and vertices in L_k do not contain vertices in their interior and intersect in v only. This shows no edge in G contains vertices in its interior and any two edges in G intersect in at most one point.

(C1) By construction, each vertex in L_k is incident to an edge of color h and an edge of color hs. Hence, D_G satisfies (C1).

(C2) By Condition (i), any horizontal line through some vertex of H and a slanted straight line through a vertex of height $k-1$ in H intersect in some point with x-coordinate larger than $x_{D'}$. Each vertex $w \in L_k$ is placed slightly to the right of such an intersection point. Hence, D_G satisfies (C2) with $x_{D_G} = x_{D'}$.

(C3) The edges in the drawing D' of H were not changed, so there are still no monochromatic crossings of those edges. Consider an edge vw with $v \in V(H)$ and $w \in L_k$.

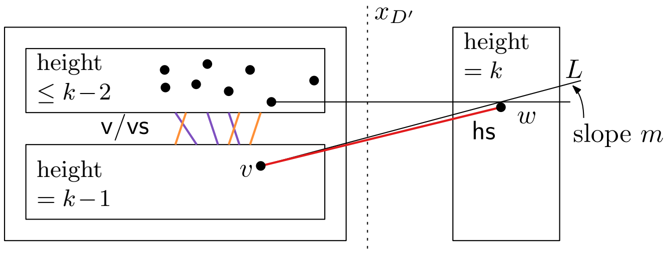

Figure 5 Checking Constraint (C3) for hs-colored edges.

First, assume that its color is h. Then $\mathsf{x}(w) > \mathsf{x}(v)$ and $\mathsf{y}(w) < \mathsf{y}(v)$ by construction. Consider an edge e of color h in H. We shall prove that e does not cross vw. If both endpoints of e lie above v, then e does not cross vw. If e crosses the horizontal line through v in some point p, then $\mathsf{x}(p) < \mathsf{x}(v)$ since v is h-open to the right in D'. Moreover, one endpoint of e lies above v while the other endpoint lies below w due to Condition (ii). So e does not cross vw. If both endpoints of e lie below v, then their y-coordinates are smaller than $\mathsf{y}(w)$ due to Condition (ii). Hence, e does not cross vw in either case.
Now consider an edge $v'w'$ of color h with $v' \in V(H)$, $\mathsf{y}(v') < \mathsf{y}(v)$ and $w' \in L_k$. As $\mathsf{y}(w) > \mathsf{y}(v')$ by Condition (ii) and $\mathsf{y}(w') < \mathsf{y}(v')$ by construction, these two edges do not cross. This shows that edges of color h do not cross in D_G.
Now assume that the color of vw is hs. By construction, v is the predecessor of w of the smallest y-coordinate. Since w has at least one predecessor of height $k-1$ and, by induction, all vertices of this height are placed below the vertices of smaller height in D', we have that height$(v) = k-1$. Consider the slanted straight line L (of slope m) through v. By Condition (i), L does not intersect the convex hull of vertices of height less than $k-1$ in D'; see Figure 5. By induction, all vertices of height $k-1$ in H are incident to edges of color v and vs only. Hence, L does not intersect any edge of color hs in D'. The edge vw has a positive slope slightly smaller than L and hence does not intersect any edge of color hs in D' either. It remains to show that vw does not intersect edges $v'w'$ of color hs with $v' \in V(H)$, $v' \neq v$, and $w' \in L_k$. Consider the slanted straight line L' (of slope m) through v'. Without loss of generality, assume that L is above L' (the case $L = L'$ produces no crossing since then $v = v'$). The edge $v'w'$ has a positive slope slightly smaller than L'. By Condition (iv), the distance between L and w is smaller than the distance between L and L'. Hence vw does not cross $v'w'$.
This shows that edges of color hs do not cross in D_G and hence D_G satisfies (C3).

(C4) No two vertices from H lie on a common vertical or horizontal line by induction. Consider $w \in L_k$ and the region B_i containing w. Due to Condition (ii) no horizontal line through B_i contains a vertex from H. Moreover, by (C2) no vertical line through B_i contains a vertex from H. Note that either two different regions B_i/B_j are separated by a horizontal line or $\mathsf{y}(p_i) = \mathsf{y}(p_j)$. In both cases, vertices placed in B_i/B_j cannot have the same y-coordinate. This is clear in the former case and in the latter it is true since we never select the same distance from p_i/p_j when placing the vertices. For the x-coordinates we can argue similarly. Hence, D_G satisfies (C4).

(C5) First, consider a vertex $v \in V(H)$ and the horizontal ray L emanating at v to the right. In the drawing D', each vertex in H is h-open to the right, so L does not intersect any h-colored edge from H. It remains to consider h-colored edges $v'w$ with $v' \in V(H)$ and $w \in L_k$. Then $\mathsf{x}(w) > \mathsf{x}(v')$ and $\mathsf{y}(v') > \mathsf{y}(w) > \mathsf{y}(v') - \varepsilon$ by construction. So if $\mathsf{y}(v') < \mathsf{y}(v)$, L does not intersect $v'w$. If $\mathsf{y}(v') > \mathsf{y}(v)$, then observe that $\mathsf{y}(w) > \mathsf{y}(v') - \varepsilon > \mathsf{y}(v)$ by Condition (ii). Hence L does not intersect $v'w$ in either case and v is h-open to the right in G_D.

Now consider a vertex $w \in L_k$ and the horizontal ray L emanating at w to the right. By (C2), L does not intersect any edge from H. It remains to consider h-colored edges $v'w'$ with $w' \in L_k$. Let v be the neighbor of w in H with vw colored h.

If $v' = v$, consider the region B_i containing w. If w' is in B_i, then w' and w lie on the diagonal ℓ_i in B_i. If w' is in B_j with $j < i$, then w' is placed to the left of w, and if w' is on B_j with $j > i$, then w' is placed above w. In either case, L does not intersect $v'w'$.

Now suppose that $v' \neq v$. Assume that $\mathsf{y}(v') < \mathsf{y}(v)$ then by Condition (ii) and by construction $\mathsf{y}(w) > \mathsf{y}(v') > \mathsf{y}(w')$. If on the other hand $\mathsf{y}(v') > \mathsf{y}(v)$ then $\mathsf{y}(v') > \mathsf{y}(w') > \mathsf{y}(v) > \mathsf{y}(w)$, again by Condition (ii) and by construction. In both cases, it follows that L does not intersect $v'w'$.

This shows that each vertex of G is h-open to the right in D_G.

(C6) In the drawing D', each vertex in H is v-open to the bottom. The vertices in L_k are not incident to any edges of color v. Hence, all vertices of G are v-open to the bottom in D_G. So (C6) is satisfied.

No monochromatic cycles. (C1–C6) are satisfied, thus D_G is feasible, and uses 4 colors. Consider any cycle in G and a vertex w of largest height in the cycle. Then its neighbors u and v in the cycle have to be its predecessors. Due to (C1), uw and vw do not have the same color. Hence there are no monochromatic cycles.

3 Proof of Theorem 2: The lower bound

In this section, we shall describe a 2-degenerate graph with geometric thickness at least 3. For a positive integer n let $G(n)$ denote the graph constructed as follows. Start with a vertex set Λ_0 of size n and for each pair of vertices from Λ_0 add one new vertex adjacent to both vertices from the pair. Let Λ_1 denote the set of vertices added in the last step. For each pair of vertices from Λ_1 add 89 new vertices, each adjacent to both vertices from the pair. Let Λ_2 denote the set of vertices added in the last step. For each pair of vertices from Λ_2 add one new vertex adjacent to both vertices from the pair. Let Λ_3 denote the set of vertices added in the last step. This concludes the construction. Observe that for each $i = 1, 2, 3$, each vertex in Λ_i has exactly two neighbors in Λ_{i-1}. Hence, $G(n)$ is 2-degenerate. We claim that for sufficiently large n the graph $G(n)$ has geometric thickness at least 3. To prove this result, we need several geometric and topological insights that are summarized in the following lemmas.

Let G_k denote the grid formed by k horizontal straight-line segments crossing k vertical straight-line segments. The grid G_k has four *sides*: the sets of left and right endpoints of the horizontal segments and the sets of lower and upper endpoints of the vertical segments form the four sides of G_k, respectively. The first and the last horizontal segment and the first and the last vertical segment form the *boundary* of G_k while all other segments are called the *inner edges* of G_k. We call an arrangement of straight-line segments combinatorially equivalent to G_k a *k-grid*. Here, we call two arrangements of straight lines or straight-line segments *combinatorially equivalent* if the embeddings given by the arrangement of their graphs (skeletons) are combinatorially equivalent. We point out that a k-grid sometimes refers to a set of disjoint red segments and a set of disjoint blue segments where every pair of red/blue segment intersects; e.g., [1]. Note that our definition is more restrictive. Among others, no two segments share an endpoint in our notion of a k-grid. The following lemma shows how both concepts are related. A proof is given in the full version [16, Section 3.1].

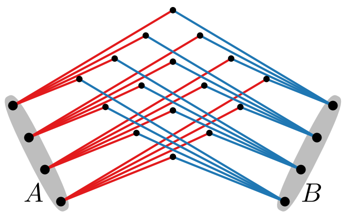

Figure 6 A tidy drawing of H_4, the full 1-subdivision of $K_{4,4}$. In particular, edges incident to A do not cross each other, edges incident to B do not cross each other, and, hence, there are no three pairwise crossing edges.

▶ **Lemma 3.** *Each arrangement of $k2^{k-1}$ disjoint red straight-line segments and k disjoint blue straight-line segments, where each red segment crosses each blue segment, contains a k-grid.*

In the following, we need a grid-structure with some additional properties summarized in the following definitions. For any point set Q in the plane, we call a straight-line segment in the plane a *Q-edge* if it has an endpoint in Q. We call two point sets A and B *separated* if $A \cup B$ is in convex position and the convex hull of A does not intersect the convex hull of B (that is, along the boundary of the convex hull of $A \cup B$ the sets do not interleave).

Consider a complete bipartite graph $K_{n,n}$ with bipartition classes A and B. Let H_n denote the graph obtained from $K_{n,n}$ by subdividing each edge exactly once. Let C denote the set of subdivision vertices of H_n. Observe that each edge of H_n has one endpoint in C and the other endpoint in $A \cup B$, and hence is either an A-edge or a B-edge. We call a geometric drawing of H_n *tidy*, if A and B are separated, there is no crossing between any two A-edges, and there is no crossing between any two B-edges. Figure 6 shows a tidy drawing of H_4. Note that we make no (convexity) assumptions on the positions of subdivision vertices. Since A and B are separated, a tidy drawing induces an ordering of A and B by traversing these points along the convex hull of $A \cup B$ in the counterclockwise direction starting with the vertices in A. An edge of H_n is called an *inner edge* if it is not incident to the first or last vertex of A and not incident to the first or last vertex of B in the order given above. Similarly, we call an edge of the underlying copy of $K_{n,n}$ an *inner edge* if it corresponds to two inner edges of H_n.

Consider a k-grid T in H_n with one side in A and one side in B (and the respective opposite sides in C). We call the sides of T that are contained in A or B the A-side and B-side, respectively. Let $a_1, \dots, a_k$ denote the vertices of the A-side of T in the order given by A and let $b_1, \dots, b_k$ denote the vertices of the B-side of T in the order given by B. For each i let x_i^A denote the crossing point between the A-edge of T with endpoint a_i and the B-edge in T farthest away from a_i. For i, $j \leq k$, with $i < j$, the *$A_{i,j}$-corridor* of T is the polygon enclosed by x_i^A, a_i, $a_{i+1}, \dots, a_j$, x_j^A. Crossing points $x_1^B, \dots, x_k^B$ and *$B_{i,j}$-corridors* are defined similarly. Figure 7 (right) shows examples of such corridors. A *tidy k-grid* is a topological subgraph T of a tidy drawing of H_n such that

- T is a k-grid with one side in A and one side in B (and the opposite sides in C),
- for each $i \leq k$, the segment $a_i x_i^A$ is contained in the $A_{1,k}$-corridor of T,
- for each $i \leq k$, the segment $b_i x_i^B$ is contained in the $B_{1,k}$-corridor of T.

Figure 7 shows a tidy 3-grid and a 4-grid that is not tidy.

Our arguments require a tidy grid such that every cell contains a (subdivision) vertex from C. Such a grid is called *dotted*. The following lemma shows that we can always find a suitable dotted grid. A proof is given in the full version [16, Section 3.2].

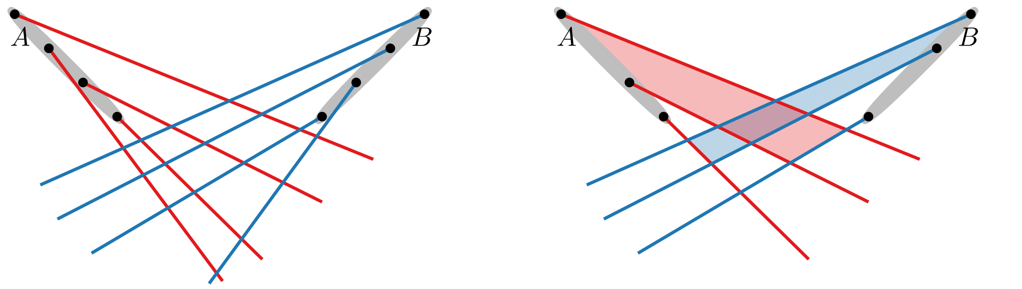

Figure 7 Left: A 4-grid with sides in A and B that is not tidy: there is a (red) A-edge not contained in the $A_{1,k}$-corridor as well as a (blue) B-edge not contained in the $B_{1,k}$-corridor. Right: A tidy (sub)grid. The $A_{1,2}$-corridor and the $B_{2,3}$-corridor are highlighted.

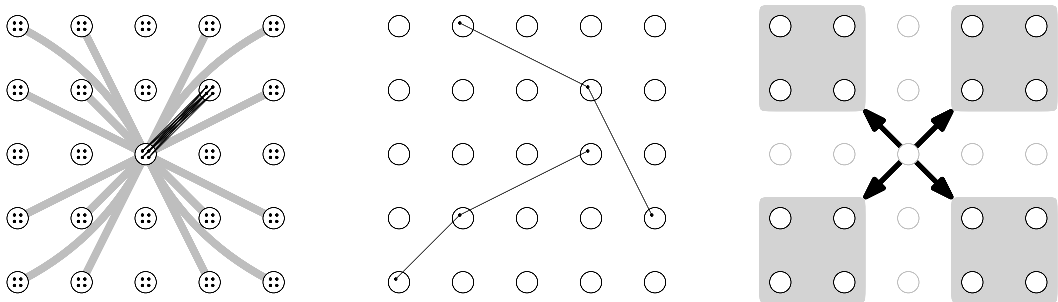

Figure 8 Left: An illustration of $\Gamma(5,4)$. Only edges incident to the central vertices are sketched. Middle: Two monotone paths in $\Gamma(5,4)$. Right: The four $(3,3)$-quadrants.

▶ **Lemma 4.** *There is a constant c_2 such that for any integers n and k, with $n \geq 2^{c_2 k^4 2^{8k}}$ and $k \geq 3$, each tidy drawing of H_n contains a dotted tidy k-grid.*

Next, we will consider connections between the vertices inside of a dotted grid. To find such connections running in certain directions within the grid, we shall use a Ramsey type argument, summarized in the following Lemma 5. We will apply this lemma in such a way that the mentioned color r corresponds to connections within the grid. For positive integers k and t let $\Gamma(k,t)$ denote the graph whose vertex set consists of disjoint sets V_i^j, $i,j \leq k$, on t vertices each, such that $u \in V_i^j$ and $v \in V_p^q$ are adjacent if and only if $i \neq p$ and $j \neq q$. See Figure 8 (left) for an illustration of $\Gamma(5,4)$. Let $r \geq 3$. We call an r-coloring of $E(\Gamma(k,t))$ *admissible* if each monochromatic copy of K_5 is of color r and any path uvw is not monochromatic in some color c with $3 \leq c < r$ in case $u \in V_i^j$, $v \in V_p^q$, and $w \in V_x^y$ with $1 \leq i < p < x \leq k$ and with $1 \leq j < q < y \leq k$ or $1 \leq y < q < j \leq k$. Loosely speaking, $\Gamma(k,t)$ is the t-blowup of the complement of a $k \times k$-grid graph, and an r-coloring is admissible if any monochromatic copy of K_5 has color r and each monotone monochromatic path on at least two edges is colored with some color in $\{1,2,r\}$. Given i and j, the *(i,j)-quadrants of* $\Gamma(k,t)$ are the four subgraphs induced by $\bigcup_{p<i,q<j} V_p^q$, $\bigcup_{p<i,q>j} V_p^q$, $\bigcup_{p>i,q<j} V_p^q$, and $\bigcup_{p>i,q>j} V_p^q$, respectively. See Figure 8 for an illustration. A proof of the following lemma is given in the full version [16, Section 3.3].

▶ **Lemma 5.** *Let r and t denote positive integers. There is a constant c_3 such that for each $k \geq c_3$ and each admissible r-coloring of $E(\Gamma(k,t))$ there are i, $j \leq k$ such that each vertex in V_i^j is incident to four edges of color r with endpoints in different (i,j)-quadrants.*

We also use the following bound on Erdős–Szekeres numbers.

▶ **Lemma 6** ([13]). *There is a constant c_4 such that for each positive integer k each set of $2^{k+c_4\sqrt{k\log k}}$ points in general position in the plane contains a subset of k points in convex position.*

Finally, we prove that the graph $G(n)$ described in the beginning of this section has geometric thickness at least 3.

▶ **Theorem 7.** *Let k, m, n be integers with $k \geq c_3$ (with c_3 from Lemma 5 for $r = 11$ and $t = 5$), $n \geq 2^{c_2 2^{10k^2}}$ (with c_2 from Lemma 4) and $m \geq 12^n$. For each $N \geq 2^{2m+c_4\sqrt{2m\log(2m)}}$ (with c_4 from Lemma 6) the graph $G(N)$ has geometric thickness at least 3.*

Proof. Consider any geometric drawing of $G = G(N)$. We assume that the vertices are in general position, otherwise we can apply a small perturbation at the vertices to achieve this without introducing any new crossings. For the sake of a contradiction, suppose that there is a partition of G into two plane subgraphs $\mathbb{A}$ and $\mathbb{B}$. We refer to the sets Λ_0, Λ_1, Λ_2, and Λ_3 as points sets like in the definition of G. Our proof proceeds as follows. We find a large tidy drawing of H_n with base points in Λ_0 and subdivision vertices in Λ_1. Lemma 4 guarantees a dotted grid in this drawing. Then we consider the connections of the vertices in the grid cells via Λ_2. We use Lemma 5 to show that many connections stay within the grid and hence many vertices of Λ_2 lie in the grid as well. Finally, we consider the connections of vertices from Λ_2 within the grid and use Lemma 5 again, to find a configuration of vertices from Λ_2 that leads to a contradiction.

Consider the point set Λ_0. Lemma 6 yields a set $\Lambda_0' \subseteq \Lambda_0$ of $2m$ points in convex position, since $N \geq 2^{2m+c_4\sqrt{2m\log(2m)}}$. We consider the points in Λ_0' in counterclockwise order with an arbitrary first vertex. Consider the copy of H_m in G between the set A of the first m vertices of Λ_0' and the set B of the last m vertices of Λ_0'. The edges of the underlying copy of $K_{m,m}$ are of four different types: in H_m they correspond to two edges from $\mathbb{A}$, or to two edges from $\mathbb{B}$, or one edge from $\mathbb{A}$ and one edge from $\mathbb{B}$ (where either the edge from $\mathbb{A}$ has an endpoint in A and the edge from $\mathbb{B}$ has an endpoint in B or vice versa). Since $m \geq 12^n$ there is, due to the bipartite Ramsey theorem (precise statement given in the full version [16, Lemma 6]), a copy of $K_{n,n}$ with all edges of the same type, leading to a corresponding copy H of H_n. Since $n \geq 3$, this type cannot be one of the types with edges only from $\mathbb{A}$ or only from $\mathbb{B}$ as both $\mathbb{A}$ and $\mathbb{B}$ are planar but $K_{3,3}$ is not. Without loss of generality, assume that all edges in H incident to A are in $\mathbb{A}$ and all edges incident to B are in $\mathbb{B}$. Observe that H is a tidy geometric drawing of H_n since $\mathbb{A}$ and $\mathbb{B}$ are crossing-free and the sets A and B are separated (their convex hulls do not intersect and $A \cup B = \Lambda_0'$ is in convex position). Further note that $2^{2k^2} \geq (k^2+1)^4 2^8$ for $k \geq 4$. Hence $n \geq 2^{c_2 2^{10k^2}} \geq 2^{c_2 (k^2+1)^4 2^{8(k^2+1)}}$, and there is, by Lemma 4, a dotted tidy (k^2+1)-grid T in H with vertices from Λ_1 in the cells.

Let $\Lambda_1' \subseteq \Lambda_1$ denote a set of vertices consisting of one vertex from each cell of T. Consider the graph Γ_1 with vertex set Λ_1' where two vertices are adjacent if and only if they are in distinct rows and distinct columns of T. Then Γ_1 forms a copy of $\Gamma(k^2, 1)$. We will define an edge coloring Φ of Γ_1 based on the drawing of the edges between Λ_1' and Λ_2. Consider two vertices x, $x' \in \Lambda_1'$. There are 89 vertices in Λ_2 adjacent to both x and x'. We will distinguish 11 different cases how the edges between such $y \in \Lambda_2$ and x, x' are drawn. Then, by the pigeonhole principle, there will be nine vertices from Λ_2 with the same type of drawing of xy and $x'y$. The cases are not disjoint from each other and we break ties arbitrarily. If there are nine vertices $y \in \Lambda_2$ with xy, $x'y \in E(\mathbb{A})$, then $\Phi(xx') = 1$. If there are nine such vertices with xy, $x'y \in E(\mathbb{B})$, then $\Phi(xx') = 2$. Now assume that there are no such nine vertices, so there are 73 such vertices where one edge is from $\mathbb{A}$ and the other edge is from

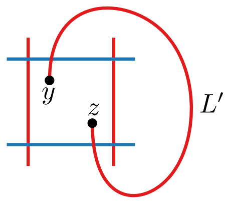

Figure 9 This arrangement is not realizable by straight-line segments, since the straight line through L' does not intersect any of the other lines twice and does not intersect itself.

$\mathbb{B}$. These edges either leave T or stay within T. If we have at least nine vertices that stay within T, we pick $\Phi(xx') = 11$. Otherwise, we can assume that there are at least 65 vertices y, for which the bicolored path xyx' leaves T. The cell containing x is the intersection of an A-corridor and a B-corridor of T. So an edge xy intersects the boundary of T either at one of the two "ends" of the A-corridor (if $xy \in E(\mathbb{A})$) or at one of the two "ends" of the B-corridor (if $xy \in E(\mathbb{B})$). Similarly, an edge $x'y$ has four options to leave T. Also observe that each of xy and $x'y$ can intersect the boundary of T only once, see Figure 9. The figure shows the boundary edges of T and a supposedly straight-line segment L' intersecting the boundary twice. This arrangement can't be realized by straight lines as the straight line through L' intersects itself once or some other line twice otherwise. This gives 8 possibilities how the intersections can be located (under the assumption that xy and $x'y$ are not both in $\mathbb{A}$ and not both in $\mathbb{B}$). We use colors $3, \ldots, 10$ to encode these possibilities. Whenever there is a set $\hat{Y}$ of nine vertices from Λ_2 such that the paths xyx' have the same locations of intersections for all $y \in \hat{Y}$, the edge xx' receives the corresponding color. If xx' is neither colored with 1, 2, or 11, we have at least 65 vertices connected via leaving T, and therefore at least one of the eight possibilities how to leave T occurs nine times. So Φ is well defined (up to breaking ties arbitrarily).

We claim that Φ is admissible. We first prove that colors $3, \ldots, 10$ do not induce a monotone monochromatic path on two edges. For the sake of a contradiction, suppose that there is such a path $xx'x''$. By symmetry, we assume that there are vertices y, y' and edges xy', $x'y \in E(\mathbb{A})$, and $x'y'$, $x''y \in E(\mathbb{B})$ such that xy' and $x'y$ leave T at the same sides of their respective A-corridors and $x'y'$ and $x''y$ leave T at the same sides of their respective B-corridors. The situation is depicted in Figure 10. We claim that this arrangement is not stretchable. To see this consider the 4-cycle between the intersections of $xy', x''y$ and the grid boundary as depicted in Figure 10 (right). This cycle needs to be embedded as a quadrilateral. For two opposing corners (the depicted crossings L_1/L_2 and $L_x/L_{x''}$) we have to embed the edges such that the "stubs" lie in the inside of the quadrilateral. To achieve this for one corner we need an incident concave angle in the quadrilateral and hence the realization of the quadrilateral would require at least two concave angles, which is not possible. Hence, such an arrangement is not stretchable. As a consequence, the colors $3, \ldots, 10$ do not induce a monotone monochromatic path on two edges. This immediately shows that these colors also do not induce a monochromatic copy of K_5. The color classes 1 and 2 correspond to subgraphs of the plane graphs $\mathbb{A}$ and $\mathbb{B}$, respectively. Hence, they do not induce monochromatic copies of K_5 as well. This shows that all monochromatic copies of K_5 are of color $r = 11$. Therefore, Φ is admissible.

Now divide the (k^2+1)-grid T into k^2 many $(k+1)$-grids T_i^j, with i, $j, \leq k$, where T_i^j consists of the A-edges on position $(i-1)k+1, \ldots, ik+1$ (in the ordering of A) and the B-edges with positions $(j-1)k+1, \ldots, jk+1$ (in the ordering of B). See Figure 11. Let Γ_i^j denote the subgraph of Γ_1 corresponding to T_i^j. Then Γ_i^j is a copy of $\Gamma(k,1)$ and Φ is an

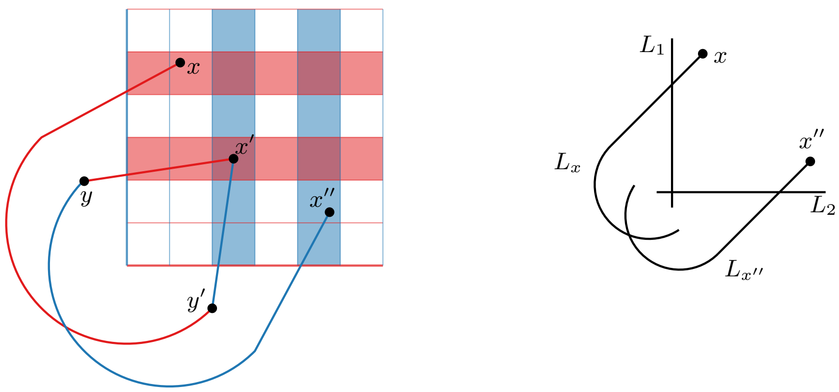

Figure 10 Left: A monotone path that is monochromatic under Φ in some color in $\{3, \ldots, 10\}$. Note that it is not possible that $x'y$ and $x'y'$ intersect. Right: The edges from the left part forming an arrangement that can't be realized by straight-line segments.

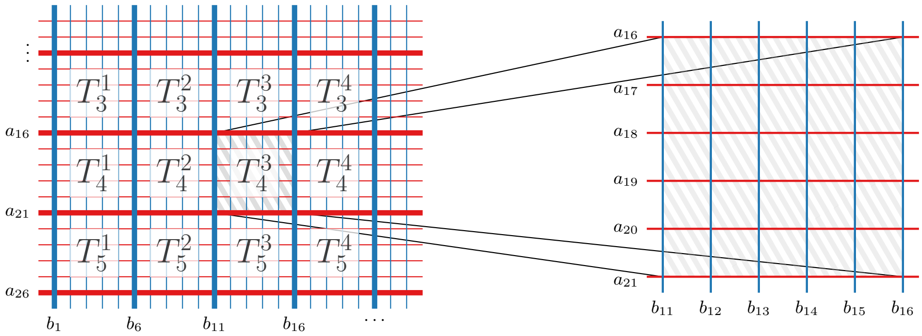

Figure 11 A (k^2+1)-grid T contains k^2 many $(k+1)$-subgrids. (Here $k = 5$.)

admissible 11-coloring of Γ_i^j. Consider some fixed $i, j \leq k$. Due to the choice of k there is, by Lemma 5, an edge xx' in Γ_i^j of color $r = 11$ (we do not need the stronger statement of the lemma here). Hence, there is a set $Y_i^j \in \Lambda_2$ of nine vertices such that for each $y \in Y_i^j$ the edges xy and $x'y$ stay within T. Let $\mathcal{A}_x$ and $\mathcal{B}_x$ denote the A-corridor and B-corridor whose intersection forms the cell containing x. Similarly, let $\mathcal{A}_{x'}$ and $\mathcal{B}_{x'}$ denote the respective corridors for the cell containing x'. As argued above, edges within T cannot leave their respective corridors. So each $y \in Y_i^j$ lies either in the cell $\mathcal{A}_x \cap \mathcal{B}_{x'}$ or in the cell $\mathcal{B}_x \cap \mathcal{A}_{x'}$. By the pigeonhole principle, there is a set $\tilde{Y}_i^j \subseteq Y_i^j$ of five vertices that lie in the same cell of T. Note that this cell is contained in T_i^j.

Consider the copy of $\Gamma(k, 5)$ whose vertex set consists of the union of all sets $\tilde{Y}_i^j$, with $i, j \leq k$, where two vertices $y \in \tilde{Y}_i^j$ and $y' \in \tilde{Y}_{i'}^{j'}$ are connected if and only if $i \neq i'$ and $j \neq j'$. For any two vertices $y, y' \in V(\Gamma(k, 5))$ there is a (unique) vertex in Λ_3 adjacent to both vertices. We define a coloring Ψ of the edges of $\Gamma(k, 5)$ similar to the coloring Φ above, except that the color of an edge yy' in $\Gamma(k, 5)$ is determined by the drawing of the unique edges yz and $y'z$, $z \in \Lambda_3$ (instead of a set of nine edge pairs behaving identically). Then Ψ is admissible by arguments similar to those applied for Φ. Due to the choice of k there are, by Lemma 5, indices $i, j \leq k$ such that each vertex in $\tilde{Y}_i^j$ is incident to four edges of color 11 under Ψ with endpoints in different (i, j)-quadrants of $\Gamma(k, 5)$. See Figure 12 for illustrations.

Let $Y = \tilde{Y}_i^j$ for the specific indices i and j from above. Consider the A-corridor $\mathcal{A}$ and B-corridor $\mathcal{B}$ of T whose intersection forms the cell containing the set Y. For a vertex $y \in Y$ consider four vertices $y_1, \ldots, y_4$ from different quadrants with $\Psi(yy_\ell) = 11$, $\ell = 1, \ldots, 4$.

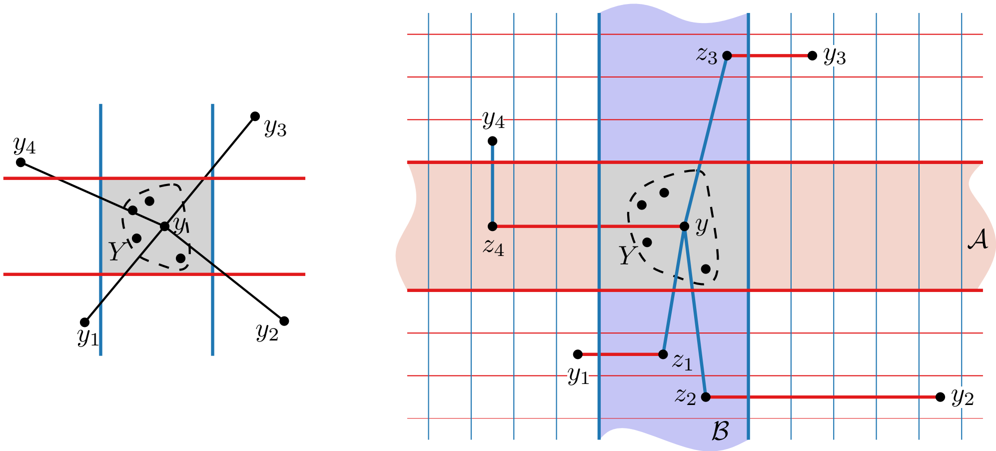

Figure 12 Left: Every vertex $y \in Y$ is incident to four edges in $\Gamma(k, 5)$ of color 11 with endpoints in different quadrants. Right: In G, each $y \in Y$ has two edges of the same type ($\mathbb{A}/\mathbb{B}$) that leave in the same direction relative to y; here yz_1 and yz_2.

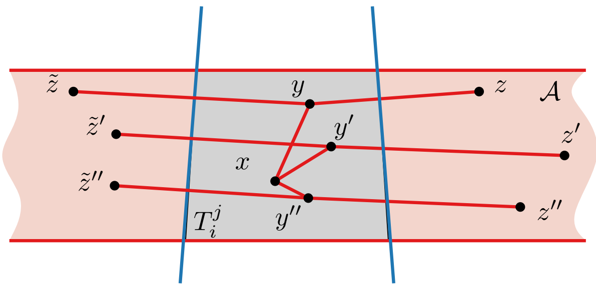

Figure 13 Construction in the proof of Theorem 7. Obtaining a monochromatic crossing at xy, xy' or xy'' is unavoidable.

Each edge $yy_\ell \in \Gamma(k, 5)$ corresponds to two edges (of G) yz_ℓ and $y_\ell z_\ell$ for some $z_\ell \in \Lambda_3$ such that z_ℓ lies within T. In particular, z_ℓ lies either in $\mathcal{A}$ or in $\mathcal{B}$ but not in the cell containing y. As $y_1, \ldots, y_4$ are from four different quadrants, two of the vertices $z_1, \ldots, z_4$ lie in $\mathcal{A}$ or two lie in $\mathcal{B}$. Moreover, for either $\mathcal{A}$ or $\mathcal{B}$ two vertices lie on different "sides" of y within the corridor. If for y we have $|\mathcal{A} \cap \{z_1, \ldots, z_4\}| \geq 2$ and at least two of these vertices lie on different sides in $\mathcal{A}$ relative to y, we call y an $\mathcal{A}$-vertex, otherwise we call y a $\mathcal{B}$-vertex.

To get a contradiction we now show that Y contains at most two $\mathcal{A}$-vertices and at most two $\mathcal{B}$-vertices, which violates $|Y| = 5$. Due to the choice of $Y \subseteq Y_i^j$, there are vertices x, $x' \in V(T_i^j) = V(\Gamma_i^j)$ such that there are edges $xy \in E(\mathbb{A})$ and $x'y \in E(\mathbb{B})$ with $\Phi(xy) = \Phi(x'y) = 11$. That is, $xy \in \mathcal{A}$ and $x'y \in \mathcal{B}$. For the sake of a contradiction, suppose that there are three $\mathcal{A}$-vertices y, y', y'' in Y. Then there are three vertices $\tilde{y}$, $\tilde{y}'$, $\tilde{y}'' \in V(\Gamma(k, 5)) \subseteq \Lambda_2$ and three vertices z, z', $z'' \in \Lambda_3$ such that yz, $y'z'$, $y''z'' \in E(\mathbb{A})$, $\tilde{y}z$, $\tilde{y}'z'$, $\tilde{y}''z'' \in E(\mathbb{B})$, and z, z', z'' lie in $\mathcal{A}$ on the same side relative to y, but not in T_i^j. By the same reasoning we can find three vertices $\tilde{z}, \tilde{z}', \tilde{z}''$ such that $y\tilde{z}, y'\tilde{z}', y''\tilde{z}'' \in E(\mathbb{A})$, but now these vertices lie on the other side in $\mathcal{A}$ relative to x (but also outside T_i^j). The edges $L = \{y\tilde{z}, yz, y'\tilde{z}', y'z', y''\tilde{z}'', y''z''\}$ split T_i^j in four zones. In one of these zones, x has to be located. No matter which zone we pick, there will always be a crossing of an edge from $\{xy, xy', xy''\} \subseteq E(\mathbb{A})$ with an edge in $L \subseteq E(\mathbb{A})$ (see Figure 13), a contradiction. Consequently, there are no three $\mathcal{A}$-vertices in Y.

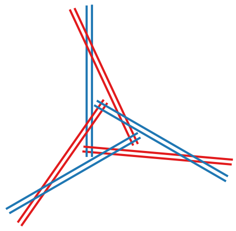

Figure 14 This arrangement of $3k$ red segments and $3k$ blue segments contains no copy of G_{k+1}. For each color and each slope there are k parallel segments (here $k = 2$ is depicted).

Similarly, there are no three $\mathcal{B}$-vertices in Y. This contradicts $|Y| \geq 5$. Hence, the geometric thickness of G is at least 3. ◀

Theorem 2 is a direct consequence of Theorem 7.

4 Conclusions

We proved that the largest geometric thickness among 2-degenerate graphs is either 3 or 4, answering two questions posed by Eppstein [11]. It remains open to decide whether there is a 2-degenerate graph of geometric thickness or geometric arboricity 4.

Our proof of the lower bound shows a geometric thickness of at least 3 for a tremendously large 2-degenerate graph. This is mainly due to using several rounds of Ramsey type arguments. We make little attempts to reduce this size and there are several places in the proof where a smaller size could be attained easily, for instance by using better or more specific Ramsey numbers (Lemma 5). In one step in the proof (Lemma 3) we are given a collection of red and blue straight-line segments in the plane and we need to find k red segments and k blue segments forming a grid combinatorially equivalent to G_k (which is formed by k horizontal segments crossing k vertical lines). We need exponentially many segments to be given, however it seems that a linear number suffices. An arrangement of $3k$ red segments and $3k$ blue segments without copy of G_{k+1} is given in Figure 14.

▶ **Question 1.** *Given an arrangement of $3k$ disjoint red straight-line segments and $3k$ disjoint blue straight-line segments, where each red segment crosses each blue segment, are there always k red segments and k blue segments forming a grid combinatorially equivalent to G_k?*

The 2-degenerate graphs form a subclass of Laman graphs, which in turn form a subclass of all graphs of arboricity 2. Our lower bound gives a graph of geometric thickness 3 in either of these classes. However, for both larger classes it is unknown whether the geometric thickness is bounded by a constant from above.

References

1 Eyal Ackerman, Jacob Fox, János Pach, and Andrew Suk. On grids in topological graphs. *Computational Geometry*, 47(7):710–723, 2014. doi:10.1016/j.comgeo.2014.02.003.

2 János Barát, Jiří Matoušek, and David R. Wood. Bounded-degree graphs have arbitrarily large geometric thickness. *The Electronic Journal of Combinatorics*, 13:R3, 2006. doi:10.37236/1029.

3 Fan R. K. Chung, Frank T. Leighton, and Arnold L. Rosenberg. Embedding graphs in books: A layout problem with applications to VLSI design. *SIAM Journal on Algebraic Discrete Methods*, 8(1):33–58, 1987. doi:10.1137/0608002.

4 Walter Didimo, Giuseppe Liotta, and Fabrizio Montecchiani. A survey on graph drawing beyond planarity. *ACM Computing Surveys*, 52:1–37, January 2020. doi:10.1145/3301281.

5 Michael B. Dillencourt, David Eppstein, and Daniel S. Hirschberg. Geometric thickness of complete graphs. *Journal of Graph Algorithms & Applications*, 4(3):5–17, 2000. arXiv:math/9910185, doi:10.7155/jgaa.00023.

6 Vida Dujmović and David R. Wood. Graph treewidth and geometric thickness parameters. *Discrete & Computational Geometry*, 37:641–670, 2007. arXiv:math/0503553, doi:10.1007/s00454-007-1318-7.

7 Christian A. Duncan. On graph thickness, geometric thickness, and separator theorems. *Computational Geometry*, 44:95–99, February 2011. doi:10.1016/j.comgeo.2010.09.005.

8 Christian A. Duncan, David Eppstein, and Stephen G. Kobourov. The geometric thickness of low degree graphs. In *SCG '04: Proceedings of the Twentieth Annual Symposium on Computational Geometry*, pages 340–346, New York, NY, USA, June 2004. Association for Computing Machinery. doi:10.1145/997817.997868.

9 Stephane Durocher, Ellen Gethner, and Debajyoti Mondal. Thickness and colorability of geometric graphs. *Computational Geometry*, 56:1–18, 2016. doi:10.1016/j.comgeo.2016.03.003.

10 David Eppstein. Separating geometric thickness from book thickness. Preprint, 2001. arXiv:math/0109195.

11 David Eppstein. Separating thickness from geometric thickness. In János Pach, editor, *Towards a Theory of Geometric Graphs*, volume 342 of *Contemporary Mathematics*. American Mathematical Society, 2004. arXiv:math/0204252.

12 Lebrecht Henneberg. Die Graphische Statik der Starren Körper. In Felix Klein and Conrad Müller, editors, *Encyklopädie der Mathematischen Wissenschaften mit Einschluss ihrer Anwendungen: Vierter Band: Mechanik*, pages 345–434, Wiesbaden, Germany, 1908. B. G. Teubner Verlag. doi:10.1007/978-3-663-16021-2_5.

13 Andreas F. Holmsen, Hossein Nassajian Mojarrad, János Pach, and Gábor Tardos. Two extensions of the Erdős–Szekeres problem. *Journal of the European Mathematical Society*, 22(12):3981–3995, 2020. doi:10.4171/JEMS/1000.

14 Seok-Hee Hong and Takeshi Tokuyama, editors. *Beyond Planar Graphs. Communications of NII Shonan Meetings*. Springer, 2020. doi:10.1007/978-981-15-6533-5.

15 Joan P. Hutchinson, Thomas C. Shermer, and Andrew Vince. On representations of some thickness-two graphs. *Computational Geometry*, 13(3):161–171, 1999. doi:10.1016/S0925-7721(99)00018-8.

16 Rahul Jain, Marco Ricci, Jonathan Rollin, and André Schulz. On the geometric thickness of 2-degenerate graphs. Preprint, 2023. arXiv:2302.14721.

17 Paul C. Kainen. Thickness and coarseness of graphs. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 39:88–95, 1973. doi:10.1007/BF02992822.

18 Anthony Mansfield. Determining the thickness of graphs is NP-hard. *Mathematical Proceedings of the Cambridge Philosophical Society*, 93(1):9–23, 1983. doi:10.1017/S030500410006028X.

19 Petra Mutzel, Thomas Odenthal, and Mark Scharbodt. The thickness of graphs: A survey. *Graphs and Combinatorics*, 14:59–73, 1998. doi:10.1007/PL00007219.

20 Crispin St.J. A. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *Journal of the London Mathematical Society*, 36:445–450, 1961. doi:10.1112/jlms/s1-36.1.445.

21 Crispin St.J. A. Nash-Williams. Decomposition of finite graphs into forests. *Journal of the London Mathematical Society*, 39:12, 1964. doi:10.1112/jlms/s1-39.1.12.

22 János Pach and Rephael Wenger. Embedding planar graphs at fixed vertex locations. *Graphs and Combinatorics*, 17(4):717–728, 2001. doi:10.1007/PL00007258.

23 Sergey Pupyrev. Linear Graph Layouts. Accessed: 2022-11-30. URL: https://spupyrev.github.io/linearlayouts.html.

24 Walter Schnyder. Embedding planar graphs on the grid. In *SODA '90: Proceedings of the first annual ACM-SIAM Symposium on Discrete Algorithms*, pages 138–148. Society for Industrial and Applied Mathematics, January 1990. doi:10.5555/320176.320191.