Kumulative Habilitationsschrift

Coloring and Covering
—
Geometric Graphs and Hypergraphs

Dr. Torsten Ueckerdt

Karlsruhe Institute of Technology

December 1, 2016

Research Summary

Torsten Ueckerdt

This is a brief summary of the research I conducted in discrete mathematics and theoretical computer science, particularly in graph theory, combinatorics, discrete geometry, order theory and game theory. In many cases we are concerned with combinatorial problems in a geometric setting, being motivated and driven by the question of how discrete combinatorial properties can capture the continuous world of geometry. Within this summary I focus on coloring problems, intersection representations, and covering problems.

Geometrically defined graphs and hypergraphs are a classical topic in discrete mathematics. In fact, the Four-Color-Problem for planar graphs is generally recognized as the driving force that led to the development of modern graph theory. Nowadays, some of the most intriguing areas of combinatorics concern graphs, hypergraphs and partially ordered sets that arise from geometric settings, the majority of which seeks to color or cover the elements at hand. The interest in combinatorial geometry stems not only from its beauty and complexity, but also from the fact that geometric arrangements play a central role in many sciences, such as physics, biology and computer science, as well as in many applications, such as geographical maps, sensor networks, chip designs, or resource allocations.

1 Coloring Problems

Many combinatorial questions, and many important combinatorial questions, can be stated as a coloring problem. Colorings, being an illustrative model for labelings, assignments, partitions, and clusterings, are easily accessible and at the same time absolutely intriguing. A famous example is the Four-Color-Problem for planar graphs. The question of how the combinatorial property of admitting a proper 4-coloring is related to the geometric property of admitting a crossing-free embedding in the plane has attracted hundreds of researchers and resulted in rich theories with very deep, fundamental insights, before and even after the problem has been finally proven by Appel and Haken [3]. Besides plain inquisitiveness, coloring problems became a central topic in discrete mathematics because of their numerous and manifold applications in all areas of combinatorics and real-world problems. Even structural results for graphs and hypergraphs –especially those defined in a geometric setting– are often stated in terms of coloring the vertices, edges, relations, hyperedges, angles, or faces. Surely, fascinating coloring problems will continue to be the engine that drives the development of a deeper understanding of discrete geometry and combinatorics.

1.1 Range capturing hypergraphs

Let X be a locally finite point set in $\mathbb{R}^d$ and $\mathcal{R}$ be a class of subsets of $\mathbb{R}^d$, which we call *ranges*. Typical ranges are the class of all lines, all balls, or all axis-aligned octants. We then obtain the *range capturing hypergraph* $\mathcal{H} = \mathcal{H}(X, \mathcal{R})$ with vertex set X by defining the hyperedges to be exactly those $Y \subseteq X$ for which there exists a range $R \in \mathcal{R}$ satisfying $Y = X \cap R$. I.e., hyperedges are those subsets of vertices that can be captured by a range.

Range capturing hypergraphs appear naturally in applications. For example when X is the set of positions of radio masts and $\mathcal{R}$ is the class of all unit disks, then the corresponding range capturing hypergraph characterizes those subsets of radio masts that can communicate with each other. In proximity representations the range capturing hypergraph is pruned as to contain only the hyperedges of a given size k. The case when $k = 2$ and $\mathcal{R}$ is the class of all homothetic copies of a fixed convex set S, the resulting graphs are called convex distance function Delaunay triangulations [17]. In fact if $k = 2$ and S is a triangle, these planar graphs are closely related to Schnyder realizer and the dimension of the vertex-edge poset of the graph [49]. For $k > 2$ range capturing hypergraphs are the central objects in the study of weak ϵ-nets [41].

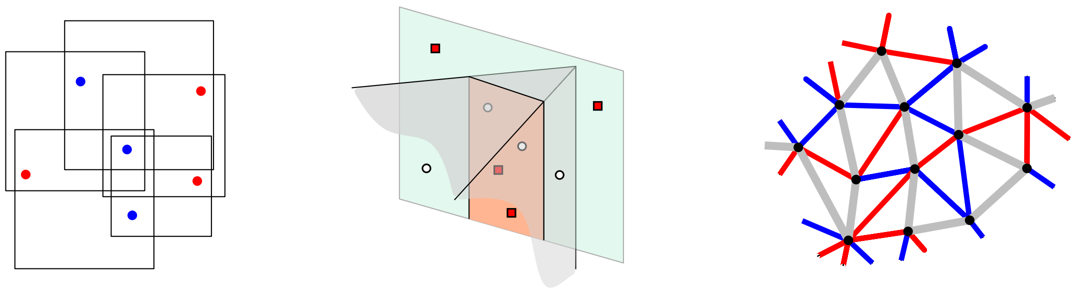

Figure 1: Left: Vertex coloring of proximity hypergraphs. Middle: 3-good 2-coloring of 7 points with respect to a family $\mathcal{R}$ of squares. Right: Improper edge coloring into two trees and one path.

A particularly important coloring problem for range capturing hypergraphs is the following: Given a finite point set $X \subset \mathbb{R}^2$, a class of ranges $\mathcal{R}$, and a natural number t, can we color the points in X with t colors such that every hyperedge in $\mathcal{H}(X, \mathcal{R})$ of size at least p, for some p, contains at least one point of each color, i.e.,

$$\forall R \in \mathcal{R} \text{ with } |R \cap X| \geq p \text{ we have that } R \cap X \text{ contains every color?}$$

Let us call such a coloring a *p-good t-coloring of X with respect to $\mathcal{R}$*, and define

$$p(t) = \min\{p \mid \forall \text{ finite } X \subset \mathbb{R}^2\ \exists\ p\text{-good } t\text{-coloring of } X \text{ with respect to } \mathcal{R}\}.$$

The left of Figure 1 shows a 3-good 2-coloring of a set X of 7 points with respect to a family $\mathcal{R}$ of squares. The dual version of this problem is known as cover-decomposability and can be stated as follows: Given a finite set $\mathcal{R}$ of ranges and a natural number t, can we color the ranges in $\mathcal{R}$ with t colors such that every point that is contained in at least $p(t)$ ranges, for some function $p(t)$, is contained in at least one range of each color? These problems, where we are interested for every t in the smallest $p(t)$ possible (if at all possible), have immediate implications for the existence of weak ϵ-nets. We proved some of the currently best known upper and lower bounds on $p(t)$ for ranges that are bottomless rectangles [4], triangles [12] and octants in 3D [13].

Theorem 1 (Asinowski *et al.* [4])**.**
For $\mathcal{R}$ being the class of all bottomless rectangles we have $\frac{5}{2}t \leq p(t) \leq 3t - 2$.

Theorem 2 (Cardinal, Knauer, Micek, Ueckerdt [13])**.**
For $\mathcal{R}$ being the class of negative octants in $\mathbb{R}^3$ we have $p(t) \leq p(2) \cdot t^{\log_2(2p(2)-1)}$.

Using the currently best upper bound $p(2) \leq 9$ due to Keszegh and Pálvölgyi [33], we conclude from Theorem 2 that $p(t) \leq 9t^{4{,}088}$. We remark that negative octants in $\mathbb{R}^3$ generalize both, bottomless rectangles and homothetic triangles in $\mathbb{R}^2$, c.f. the middle of Figure 1. In fact, Theorem 1 gives the best known lower bound for octants, while Theorem 2 gives the best known upper bound for homothetic triangles. The linear upper bound $p(t) = O(t)$ in Theorem 1 immediately implies the following.

Corollary 3. *For $\mathcal{R}$ being the class of all bottomless rectangles the following holds. For every finite point set $X \subset \mathbb{R}^2$ and every $\varepsilon > 0$ there exists $Y \subseteq X$ with $|Y| \leq 3/\varepsilon = O(1/\varepsilon)$ such that*

$$\forall R \in \mathcal{R} \textit{ with } |R \cap X| \geq \varepsilon|X| \textit{ we have } R \cap Y \neq \emptyset.$$

For given $\mathcal{R}$ and ε, a subset $Y \subseteq X$ as in Corollary 3 is called an *ε-net of X with respect to $\mathcal{R}$*. It is known that whenever the VC-dimension of $\mathcal{H}(X, \mathcal{R})$ is $O(1)$, there exists an ε-net $Y \subseteq X$ of size $|Y| = O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$. (See, for instance, Chapter 10 in Matoušek's lectures [42].) But for some family of ranges $\mathcal{R}$ there exist ε-nets of size $O(1/\varepsilon)$, which is asymptotically optimal. Corollary 3 indeed can be generalized to say that whenever $p(t) = O(t)$ for a given $\mathcal{R}$, then such linear size ε-nets Y with

$|Y| = O(1/\varepsilon)$ exist. As it is known that negative octants admit linear size ε-nets [19], it is interesting to see whether $p(t) = O(t)$ for homothetic triangles in $\mathbb{R}^2$ or even negative octants $\mathbb{R}^3$.

Finally, we mention that octants in $\mathbb{R}^3$ and homothetic triangles in $\mathbb{R}^2$ are self-dual and hence Theorem 2 also proves that $p(t) < \infty$ for the dual coloring problem with respect to these ranges. On the other hand, for the dual coloring problem with respect to bottomless rectangles, we only know that $p(2) = 3$ [32] and that we can not prove that $p(t) < \infty$ via any semi-online coloring [13] (a concept used to prove Theorem 1 for the primal coloring problem). understanding for which classes $\mathcal{R}$ of ranges it is true that $p(t) < \infty$ implies $p(t+1) < \infty$ is surely the most interesting open problem here.

In order to prove coloring results for a range capturing hypergraph $\mathcal{H} = \mathcal{H}(X, \mathcal{R})$, it is necessary to investigate the structure of this hypergraph, depending on the set $\mathcal{R}$ of ranges. Most of what is known here concerns only the convex distance function Delaunay triangulations, i.e., the graph (2-uniform subhypergraph) G arising from $\mathcal{H}$ by considering only hyperedges of size 2. Whenever $\mathcal{R}$ is the class of all homothetic copies of a fixed convex set S, then G is planar, and if X lies in general position with respect to $\mathcal{R}$ (no four points in X lie on the boundary of a range $R \in \mathcal{R}$), then every inner face of G is a triangle. Schnyder's Theorem [49] implies that every inner triangulated plane graph G can arise in this way for S being a triangle. On the other hand, not every triangulated planar graph G arises when S is not a triangle; for example Dillencourt proves that S being a disk gives rise to 1-tough planar graphs only [23], while it is easily seen that S being a square can never create a planar graph with a separating triangle.

But there is one property that all range capturing hypergraphs and all proximity hypergraphs share, as long as $\mathcal{R}$ is the set of all homothetic copies of a fixed convex shape S: the maximum number of (hyper)edges on a given number of vertices.

Theorem 4 (Axenovich, Ueckerdt [7])**.**
Let $S \subset \mathbb{R}^2$ be any convex compact set and $\mathcal{R}$ be the class of all homothetic copies of S. For any $k \geq 2$ and any finite point set $X \subset \mathbb{R}^2$, the number of hyperedges of size k in $\mathcal{H}(X, \mathcal{R})$ is at most

$$(2k-1)|X| - k^2 + 1 - \sum_{i=1}^{k-1} a_i,$$

where a_i is the number of i-element subsets of X that can be separated from the rest of X with a straight line.

Most interestingly, the inequality in Theorem 4 is tight whenever the boundary of S contains no corners and no straight segments, and X lies in general position with respect to $\mathcal{R}$. Summing over all $k \in \{1, \ldots, n\}$ one obtains that the total number of hyperedges in $\mathcal{H}(X, \mathcal{R})$ is at most $\binom{n}{3} + \binom{n}{2} + \binom{n}{1}$, which again is tight whenever the boundary of S has everywhere positive and finite curvature and no four points of X lies on the boundary of a homothetic copy of S.

1.2 Improper colorings

Surely, the most important colorings of graphs are the *proper vertex colorings*, i.e., colorings of the vertices such that any two adjacent vertices receive different colors. However, finding proper colorings is usually very hard, requiring high computational effort and possibly many colors. And coloring with fewer colors than needed for a proper coloring, necessarily results in at least one conflict. But can we color the graph improperly in such a way that the deficiency is not too bad? For example, can we guarantee that every vertex v has no more than four neighbors with the same color as v, or that each color induces only connected components of size at most ten? Let us refer to the right of Figure 1 and the middle of Figure 2 for an improper edge-coloring and improper vertex-coloring of a planar graph with three colors, respectively. It turns out that improper colorings with low deficiency have many connections to various areas of graph theory.

The majority of the enormous amount of literature on improper colorings concerns vertex colorings (and list colorings) of restricted planar graphs with two, three or four colors. Starting from proper

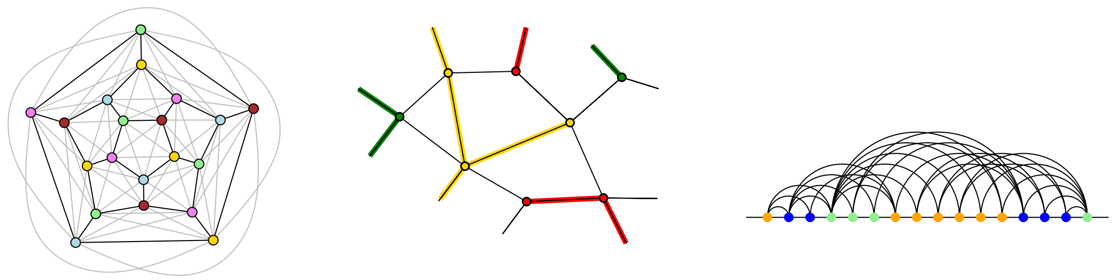

Figure 2: Left: Proper 5-coloring of a fan-planar graph. Middle: Improper vertex coloring with bounded monochromatic degree. Right: Proper 3-coloring of an ordered graph without nesting edges.

colorings where each color class induces a set of isolated vertices, one line of research is to restrict the subgraph induced by any color class to be a set of only short paths. It is known that every planar graph G of girth at least 7 admits a vertex 2-coloring such that any color class induces a set of paths of at most 2 vertices [10]. We have considered the case of girth 4, 5 and 6.

Theorem 5 (Axenovich, Ueckerdt, Weiner [8]).
Every planar graph G of girth at least 6 admits a vertex coloring in 2 colors such that every color class induces a forest in which each component is a path on at most 15 vertices.

Theorem 6 (Axenovich, Ueckerdt, Weiner [8]).
For every $k \in \mathbb{N}$ there exists a planar graph G_k of girth 4 such that for every vertex coloring in 2 colors one color class induces a path on at least k vertices.

Interestingly, in many cases of improper colorings of planar graphs girth 5 remains the only unknown case. For example, is it possible to 2-color the vertices of any planar graph of girth 5 so that each monochromatic component has at most 3 vertices?

Let us mention that improper vertex colorings of planar graphs are also closely related to the colorings of range capturing hypergraphs as defined above. For example, if the 2-regular proximity graph G is planar and every hyperedge on at least t vertices induces a triangle in G, this proves that $p(2) \leq t$ since planar graphs can be 2-colored with no monochromatic triangles. Surely, further developments in the field of improper colorings would have further applications for range capturing hypergraphs.

1.3 Coloring embedded graphs

The Four-Color-Problem is the classic example of a coloring problem in a geometric setting. In related questions we are given an embedded graph G, i.e., G is implicitly defined by a geometric arrangement of a certain kind, and the question is to determine or bound the chromatic number $\chi(G)$.

Most naturally, G is given with a classical node-link diagram in 2D and we have a forbidden pattern of how sets of edges are not allowed to cross. Then one way to upper bound the chromatic number is to show that the number of edges in G is only linear in its number of vertices. This has been done for 1-planar graphs [45], 4-quasiplanar graphs [1], and fan-crossing free graphs [16].

We have introduced a new class of almost planar graphs: the *fan-planar graphs*. A graph is fan-planar if it admits a simple topological drawing in which for each edge e the edges crossing e have a common endpoint on the same side of e. I.e., all crossings are of the form that a fan of incident edges at some vertex are crossed left-to-right by another edge e. See the left of Figure 2. This can be formulated by two forbidden patterns, one of which is the configuration where an edge e is crossed by two independent edges and the other where e is crossed by incident edges with the common endpoint on different sides of e. We remark that every 1-planar graph is also fan-planar, and every fan-planar graph is also 3-quasiplanar, where both inclusions are strict.

Theorem 7 (Kaufmann, Ueckerdt [31]).
Every n-vertex fan-planar graph has at most $5n-10$ edges and there exists an infinite family of fan-planar graphs with n vertices and $5n - 10$ edges.

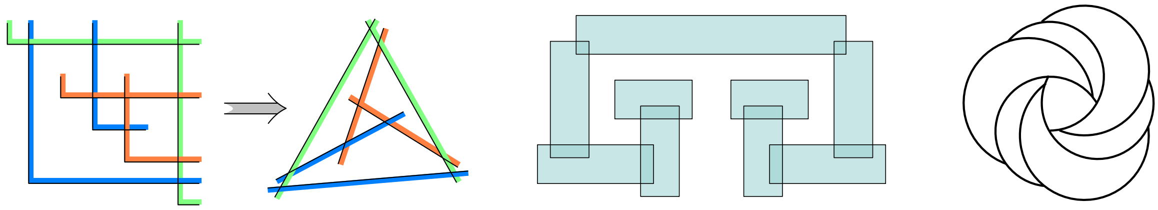

Figure 3: Left: Stretching an L-representation into a segment representation. Middle: Rectangle arrangement that is not stretchable into squares. Right: Contacts of circular arcs.

Theorem 7 immediately implies that every fan-planar graph G has a vertex of degree at most 9, proving that $\chi(G) \leq 10$. A similar statement for the class of k-quasiplanar graphs is a famous conjecture [11] in the field: Is it true that for every fixed k the number of edges in k-quasiplanar n-vertex graphs is linear in n? This would in particular imply that there is some $f(k)$ such that every for k-quasiplanar graph G we have $\chi(G) \leq f(k)$.

It is also sensible to consider graphs embedded in 1D where vertices are mapped to integers and edges to intervals. This setting, usually under the name *ordered graphs*, has become quite popular over the last few years, e.g., in terms of extremal functions [35] and Ramsey theory [44].

We recently investigated very general forbidden edge patterns, such as k pairwise crossing (overlapping) edges or k pairwise nesting edges, where we determine exactly the maximum chromatic number of such embedded ordered graphs [5]; see for example the right of Figure 2. We also considered the chromatic number of an ordered graph G with a forbidden ordered subgraph H [6]. In sharp contrast to the unordered setting, we prove that there are some ordered paths H such that ordered graphs avoiding H as an ordered subgraph can still have arbitrarily large chromatic number.

2 Intersection Representations

An *intersection representation* of a graph is a set of (geometric) objects, one for each vertex, such that two vertices are adjacent if and only if the corresponding objects have a non-empty intersection. Intersection representations arise naturally from applications, for example in constellations of objects (vertices), each with a geometric position and a sphere of influence, like radio towers with broadcast coverages or electric cables with fields of tension, but also when objects are moving entities and one is interested in the intersections of their trajectories.

When arbitrary objects are allowed to represent the vertices, every graph admits an intersection representation. But as soon as we restrict ourselves to objects only of a certain type or allow intersections to be only of a certain type, we naturally obtain the class of all those graphs admitting such restricted intersection representations. In the following I discuss several important examples of restricted intersection representations. Further examples will follow in the Section 3 on coverings problems.

2.1 Connected sets in the plane – String graphs

A very natural and well-studied type of intersection representations uses as objects path-connected sets in the Euclidean plane. The corresponding class of intersection graphs is the class of *string graphs*. String graphs contain numerous non-planar graphs, but are not closed under taking subgraphs, and indeed some graphs (like full subdivisions of non-planar graphs) are not string graphs.

Variants of such intersection representations with very restricted sets in the plane, such as disks [18], segments [40], or squares [39], are as numerous as relevant, and play a key role in applications such as chip designs and frequency assignment problems.

In 1985 Scheinerman conjectured that every planar graph is a *segment graph*, that is, it admits an intersection representation with segments in the plane [47]. After being a famous open problem for more than 20 years, it has been finally verified in 2009.

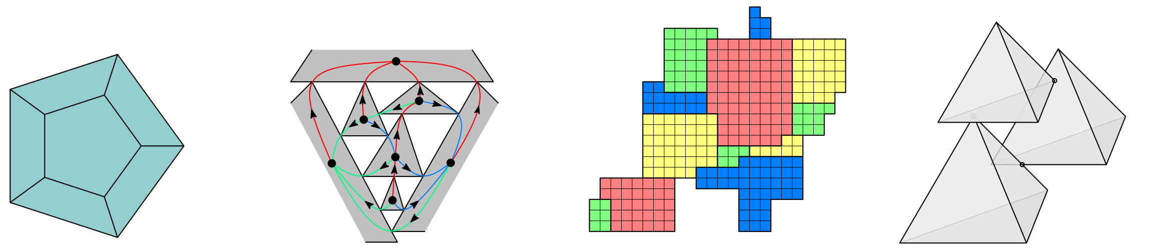

Figure 4: From left to right: Proper side-contacts of polygons; Triangle contacts induce a Schnyder realizer; Cartogram of central Europe with respect to CO_2-emissions in 2009; Contacts of 3D tetrahedra.

Theorem 8 (Chalopin, Gonçalves [14]).
Every planar graph is a segment graph.

A subclass of segment graphs are so-called *L-graphs*, that is, graphs that admit an intersection representation with axis-aligned paths with one bend and with the same orientation as the letter 'L' (in other words, the union of the lower and left side of an axis-aligned box). Indeed, Middendorf and Pfeiffer proved in 1992 that every L-representation can be "stretched" into a segment representation [43]; see the left of Figure 3. Generalizing L-graphs, one defines B_k-VPG graphs, which admit intersection representations with axis-aligned paths with k bends each.

In the light of Theorem 8 there are two main open problems in the field, both of which we could partially but not completely answer.

Problem 9. *Is every planar graph an L-graph? And is the complement of every planar graph a segment graph?*

Theorem 10 (Chaplick, Ueckerdt [15]).
Every planar graph is a B_2-VPG graph.

Theorem 11 (Felsner, Knauer, Mertzios, Ueckerdt [25]).
Every planar 3-tree is an L-graph. And every complement of a planar graph is a B_{19}-VPG graph.

Both questions in Problem 9 remain open in their full generality.

2.2 Contact representations

A *contact representation* is a special kind of intersection representation where we require the geometric objects representing the vertices to be interiorly disjoint. Then intersections can happen only along the boundaries of the objects, i.e., when the objects "touch" or "make contact". If we additionally require that contacts are not just isolated points and each object is a connected set, then in the plane only planar graphs admit such contact representations. We refer to Figure 4 and the right of Figure 3 for some illustrative examples.

Contact representations inherit a close relationship between the actual geometric realization of the arrangement and its combinatorial properties. Every contact representation naturally induces a plane embedding for the underlying planar graph G, but sometimes we can read off much more than that: When using polygonal objects such as polygons or polylines, every contact involves at least one corner or endpoint of at least one of the two objects involved. The information for each contact which corner/endpoint of which object is used, can be seen as a coloring and orientation of the edges of G that satisfies some local properties around each vertex depending on the type of objects. Examples are Schnyder realizer that arise from triangle contact representations [22], separating decompositions arising from 2-directional segment contact representations [21] and transversal structures arising from proper side-contact representations with rectangles [27].

Perhaps surprisingly, the local combinatorial information in all these combinatorial structures is enough to characterize the contact representation up to topological equivalence, showing that the geometry of touching polygons is combinatorially equivalent to graph theoretic criteria. Moreover, the existence of Schnyder realizer, separating decompositions and transversal structures characterizes the respective subclasses of planar graphs and so these became important tools in tasks such as enumeration, construction sequences, random generation, underlying poset structures and grid drawings.

We have extended the set of bijections between specific contact representations and coloring and orientations of the underlying contact graphs by two more entries. A *Laman graph* is a minimally rigid graph, or equivalently, a Laman graph on n vertices has exactly $2n-3$ edges, and any set of k vertices $(2 \leq k \leq n)$ induces at most $2k-3$ edges.

Theorem 12 (Kobourov, Ueckerdt, Verbeek [38]).
Every planar Laman graph admits a contact representation with axis-aligned one-bend paths and these representations are encoded by angular trees.

Theorem 13 (Klawitter, Nöllenburg, Ueckerdt [34]).
Every maximal triangle-free planar graph admits a contact representation with axis-aligned boxes and these representations are encoded by corner edge labelings.

In most recent investigations in this area, one is interested to see what happens beyond planarity. Here one tries to find higher-dimensional analogues of both, the geometric contact representations and the combinatorial graph structures. We succeeded in generalizing Schnyder realizer to any dimension [24] and most recently transversal structures to 3 dimensions [26].

Theorem 14 (Evans, Felsner, Kobourov, Ueckerdt [24]).
There is a d-dimensional analogues of Schnyder realizer, which encode analogous kinds of contact representations with d-dimensional boxes in $\mathbb{R}^d$.

3 Covering Problems

One might argue that scientific progress is just decomposing big problems into smaller pieces, until complex obstacles become series of doable steps. Decomposing graphs into smaller graphs, or equivalently, covering graphs by smaller graphs, is one of the most fundamental subjects of graph theory. Indeed, proper vertex-colorings and proper edge-colorings are just coverings by independent sets and matchings, respectively, and improper colorings can also be equivalently seen as coverings with simpler graphs. In applications, a covering by interval graphs is for example important for scheduling a set of interdependent jobs onto a number of processors.

Let G be a graph and let $\mathcal{C}$ denote the class of graphs with which we want to cover G. In the classical covering problem, one seeks to cover G with as few graphs from $\mathcal{C}$ as possible, that is, we ask for the smallest t such that $G = H_1 \cup H_2 \cup \cdots \cup H_t$ with $H_1, \ldots, H_t \in \mathcal{C}$. We proposed [37] a unifying approach to graph covering problems, capturing the classical model (which we call *t-global* coverings) as well as two relaxations of it: *t-local* coverings and *t-folded* coverings, which have been considered only in a few special cases before.

Definition 15 (Knauer, Ueckerdt [37]).
Let G be a graph and $\mathcal{C}$ be a class of graphs. A $\mathcal{C}$-cover *of G is an edge-surjective homomorphism $\varphi : C_1 \dot\cup \cdots \dot\cup C_k \to G$. A $\mathcal{C}$-cover $\varphi : C_1 \dot\cup \cdots \dot\cup C_k \to G$ is*

- t-global *if $\varphi|_{C_i}$ is injective for $i = 1, \ldots, k$, and $t = k$,*
- t-local *if $\varphi|_{C_i}$ is injective for $i = 1, \ldots, k$ and $|\varphi^{-1}(v)| \leq t$ for every $v \in V(G)$,*
- t-folded *if $|\varphi^{-1}(v)| \leq t$ for every $v \in V(G)$.*

Finally, and $c_g^{\mathcal{C}}(G)$ (respectively $c_\ell^{\mathcal{C}}(G)$ and $c_f^{\mathcal{C}}(G)$) is the smallest $t \in \mathbb{N}$ for which there exists a t-global (respectively t-local and t-folded) $\mathcal{C}$-cover of G.

In other words, the global covering number $c_g^{\mathcal{C}}(G)$ is the smallest t such that G can be covered with t graphs from $\mathcal{C}$, i.e., the global covering number corresponds to the classical covering problem. The local

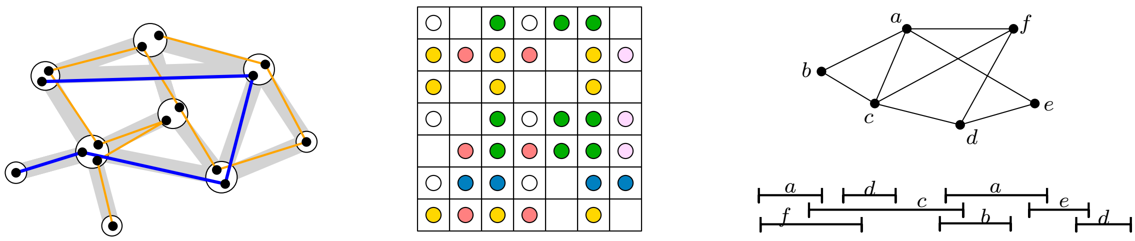

Figure 5: Left: 3-folded $\mathcal{C}$-cover of a graph with $\mathcal{C}$ being all paths. Middle: 6-global $\mathcal{C}$-cover of a bipartite graph with $\mathcal{C}$ being all complete bipartite graphs. Right: 2-folded $\mathcal{C}$-cover of a graph with $\mathcal{C}$ being all interval graphs.

covering number $c_\ell^{\mathcal{C}}(G)$ is the smallest t such that G can be covered with an arbitrary number of graphs from $\mathcal{C}$ but each vertex of G is contained in at most t such graphs. Finally, the folded covering number $c_f^{\mathcal{C}}(G)$ is the smallest t such that each vertex v of G can be split into at most t vertices, keeping each incident edge at v incident to at least one of the split vertices, such that the resulting graph is a disjoint union of graphs in $\mathcal{C}$. We refer to Figure 5 and the right of Figure 6 for some illustrative examples.

For any graph G and any class $\mathcal{C}$ we have $c_f^{\mathcal{C}}(G) \leq c_\ell^{\mathcal{C}}(G) \leq c_g^{\mathcal{C}}(G)$, because t-folded coverings are less restrictive than t-local coverings, which in turn are less restrictive than t-global coverings. In fact, in some cases where coverings are used as an ingredient in a bigger argumentation, the classical global model is used even though the local model would be sufficient. In other cases, two of the three cover variants have been investigated, but without realizing the close relation between them. With our framework we provide the concepts and the tools to discover general properties inherent to many covering problems, which we feel did not receive the appropriate attention before. Let us also mention that by considering local and folded covering numbers instead of global covering numbers, one can provide supporting evidence for some classic decomposition conjectures, such as the Double Cycle Cover Conjecture and the Linear Arboricity Conjecture.

3.1 General phenomena

During our work on the initiating paper [37], a follow-up paper [9], and third projected paper, we encountered several interesting phenomena yet to be explored and understood in full detail.

One phenomenon concerns the question by how much the folded, local and global covering number of the same graph G with respect to the same class $\mathcal{C}$ can differ.

Theorem 16 (Knauer, Ueckerdt [37])**.**
If $\mathcal{G}$ is the class of all line graphs and $\mathcal{C}$ is the class of all complete graphs, then

$$\max\{c_\ell^{\mathcal{C}}(G) \mid G \in \mathcal{G}\} = 2 \qquad \textit{and} \qquad \max\{c_g^{\mathcal{C}}(G) \mid G \in \mathcal{G}\} = \infty.$$

In other cases, for example when $\mathcal{G}$ is the class of all graphs and $\mathcal{C}$ is the class of all bipartite graphs, then $\max\{c_f^{\mathcal{C}}(G) \mid G \in \mathcal{G}\} = 2$ and $\max\{c_\ell^{\mathcal{C}}(G) \mid G \in \mathcal{G}\} = \infty$. On the other hand, we can prove that if $\mathcal{C}$ is closed under taking vertex-disjoint unions and topological minors, then there exists a function φ such that for every graph G we have $c_g^{\mathcal{C}}(G) \leq \varphi(c_f^{\mathcal{C}}(G))$, i.e., such a tremendous separation of local and global, or local and folded covering numbers is impossible. However, more precise conditions under which folded, local and global covering numbers have always comparable magnitude are yet to be discovered.

A second phenomenon concerns the computational complexity of computing the folded, local or global covering number. All graph classes $\mathcal{C}$ (that satisfy some reasonable assumptions) for which we know some computationally easy or hard cases show the following pattern: computing $c_f^{\mathcal{C}}(G)$ is "at most as difficult" as computing $c_\ell^{\mathcal{C}}(G)$, which in turn is "at most as difficult" as computing $c_g^{\mathcal{C}}(G)$. While we know cases for which all three covering numbers are efficiently computable, other cases for which this holds only for the folded and local variant, yet other cases for which folded covering numbers are

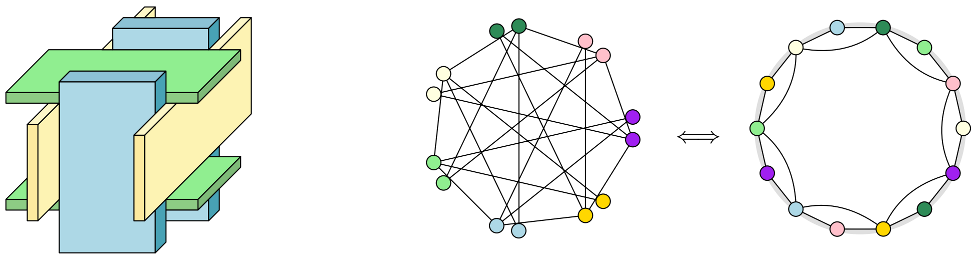

Figure 6: Left: 1-local boxes in $\mathbb{R}^3$. Right: 2-folded $\mathcal{C}$-cover of K_7 with $\mathcal{C}$ being all outerplanar graphs.

computationally easy while the local and global covering number are not, and also cases where the computation of all three covering numbers is NP-complete, we lack any explanation of this phenomenon.

Theorem 17 (Knauer, Ueckerdt [37]).
Let $\mathcal{C}$ be the class of all stars. Then for every graph G we have $c_f^{\mathcal{C}}(G) = c_\ell^{\mathcal{C}}(G)$. Moreover, computing $c_\ell^{\mathcal{C}}(G)$ can be done in polynomial time.

On the other hand, deciding whether $c_g^{\mathcal{C}}(G) \leq k$ is known to be NP-complete for $k = 2$ [30] and $k = 3$ [29]. Curiously, it remains open whether there is a union-closed graph class $\mathcal{C}$ such that computing the local or folded covering number is NP-complete, whereas the global covering number can be computed in polynomial time.

3.2 Coverings with interval graphs

For many classical covering problems we seek to cover a given graph G with graphs from $\mathcal{C}$ where $\mathcal{C}$ is the class of some particular geometric intersection graphs. A very important type of intersection graph, that was not mentioned in Section 2, is one in which vertices are represented by intervals on the real line and edges correspond to non-disjoint intervals. The graphs admitting such intersection representation are the famous and versatile *interval graphs*, which received a lot of attention and are today quite well understood. Less understood, although important in many real-world applications, are intersection representations in which each vertex is represented by a set of more than one interval.

Now, an intersection representation of G in which every vertex is represented by a set of at most t intervals corresponds exactly to a t-folded cover of G where $\mathcal{C}$ is the class of interval graphs and every vertex is split into at most t vertices. Minimizing t leads to the so-called *interval number* of G, or equivalently the folded covering number $c_f^{\mathcal{C}}(G)$. On the other hand, the *track number* of G is the smallest t such that G is the union of t interval graphs, and hence the track number is the same as the global covering number $c_g^{\mathcal{C}}(G)$. Important results of Scheinerman and West, respectively Gonçalves, state that if G is planar then $c_f^{\mathcal{C}}(G) \leq 3$ [48] and $c_g^{\mathcal{C}}(G) \leq 4$ [28], where both results are best-possible. However, the corresponding local covering number $c_\ell^{\mathcal{C}}(G)$ has not been considered so far. In particular, it is open whether for any planar graph G we have $c_\ell^{\mathcal{C}}(G) \leq 3$ for $\mathcal{C}$ being the class of interval graphs, i.e., G can be split into interval graphs such that every vertex appears in only three interval graphs.

Theorem 18 (Knauer, Ueckerdt [37]).
Let $\mathcal{C}$ be the class of all interval graphs. Then for every planar graph G of tree-width 3 and every planar bipartite graph G we have $c_\ell^{\mathcal{C}}(G) \leq 3$.

We remark that the known planar graphs G with $c_g^{\mathcal{C}}(G) = 4$ are bipartite, for which by Theorem 18 we have $c_\ell^{\mathcal{C}}(G) \leq 3$. Trying to prove that $c_\ell^{\mathcal{C}}(G) \leq 3$ for every planar graph G, a gap in the proof of Scheinerman and West [48] that $c_f^{\mathcal{C}}(G) \leq 3$ for every planar G was found, which could not be fixed. Very recently, we could reprove their result with completely different arguments [36] and we hope to generalize our techniques to prove that even $c_\ell^{\mathcal{C}}(G) \leq 3$ holds for every planar G.

A higher-dimensional analog to intervals on the real line are axis-aligned boxes in $\mathbb{R}^t$, i.e., Cartesian products of t real intervals. The smallest integer t such that a given graph G admits an intersection representation with t-dimensional boxes is called the *boxicity* of G, denoted by $\mathrm{box}(G)$. The boxicity was introduced by Roberts [46] in 1969 and has many applications in as diverse areas as ecology and operations research [20].

The boxicity of any graph G can be equivalently seen as the smallest integer t such that $G = H_1 \cap \cdots \cap H_t$ for $H_1, \ldots, H_t$ being interval graphs. Note that this is equivalent to $G^c = H_1^c \cup \cdots \cup H_t^c$ where G^c denotes the complement of G and thus $H_1^c, \ldots, H_t^c$ are co-interval graphs (also known as comparability graphs of interval orders). Hence we can interpret the boxicity as a covering parameter by $\mathrm{box}(G) = c_g^{\mathcal{C}}(G^c)$ where $\mathcal{C}$ denotes here the class of all co-interval graphs. Using our framework of folded, local and global covering numbers [37], we recently introduced two boxicity-related concepts, which we call the *local boxicity* $\mathrm{box}_\ell(G)$ and the *union boxicity* $\overline{\mathrm{box}}(G)$. Indeed, the three parameters boxicity, local boxicity and union boxicity are non-trivial and reflect different aspects of the graph.

Theorem 19 (Bläsius, Stumpf, Ueckerdt [9])**.**
For every graph G we have $\mathrm{box}_\ell(G) \leq \overline{\mathrm{box}}(G) \leq \mathrm{box}(G)$. Moreover, for every positive integer k there exist graphs G_k, G'_k , G''_k with

- $\mathrm{box}_\ell(G_k) \geq k$,
- $\mathrm{box}_\ell(G'_k) = 2$ *and* $\overline{\mathrm{box}}(G'_k) \geq k$,
- $\overline{\mathrm{box}}(G''_k) = 1$ *and* $\mathrm{box}(G''_k) = k$.

We also give geometric interpretations of the local and union boxicity of a graph G in terms of intersecting high-dimensional boxes. For positive integers k, d with $k \leq d$ we call a d-dimensional box $B = I1 \times \cdots \times I_d$ *k-local* if for at most k indices $i \in \{1, \ldots, d\}$ we have $I_i \neq \mathbb{R}$. Thus a k-local d-dimensional box is the Cartesian product of d intervals, at least $d - k$ of which are equal to the entire real line $\mathbb{R}$. See the left of Figure 6 for an illustration.

Theorem 20 (Bläsius, Stumpf, Ueckerdt [9])**.**
Let G be a graph.

- *We have $\overline{\mathrm{box}}(G) \leq k$ if and only if there exist $d_1, \ldots, d_k$ such that G is the intersection graph of Cartesian products of k boxes, where the ith box is 1-local d_i-dimensional, $i = 1, \ldots, k$.*
- *We have $\mathrm{box}_\ell(G) \leq k$ if and only if there exists some d such that G is the intersection graph of k-local d-dimensional boxes.*

Let us remark that the boxicity of graphs is closely related to the dimension of posets [2]. It is a promising task to carry over the exciting connections between boxicity and poset dimension to the new concepts of local boxicity and union boxicity.

References

[1] E. Ackerman. On the maximum number of edges in topological graphs with no four pairwise crossing edges. In *Proceedings of the twenty-second annual symposium on Computational geometry*, pages 259–263. ACM, 2006.

[2] A. Adiga, D. Bhowmick, and L. S. Chandran. Boxicity and poset dimension. *SIAM Journal on Discrete Mathematics*, 25(4):1687–1698, 2011.

[3] K. I. Appel and W. Haken. *Every planar map is four colorable*, volume 98. AMS Providence, 1989.

[4] A. Asinowski, J. Cardinal, N. Cohen, S. Collette, T. Hackl, M. Hoffmann, K. B. Knauer, S. Langerman, M. Lason, P. Micek, G. Rote, and T. Ueckerdt. Coloring hypergraphs induced by dynamic point sets and bottomless rectangles. In F. Dehne, R. Solis-Oba, and J. Sack, editors, *Algorithms and Data Structures - 13th International Symposium, WADS 2013, London, ON, Canada, August 12-14, 2013. Proceedings*, volume 8037 of *Lecture Notes in Computer Science*, pages 73–84. Springer, 2013. `http://dx.doi.org/10.1007/978-3-642-40104-6_7`.

[5] M. Axenovich, J. Rollin, and T. Ueckerdt. The chromatic number of ordered graphs with constrained conflict graphs. *CoRR*, abs/1610.01111, 2016. http://arxiv.org/abs/1610.01111.

[6] M. Axenovich, J. Rollin, and T. Ueckerdt. Chromatic number of ordered graphs with forbidden ordered subgraphs. *CoRR*, abs/1603.00312, 2016. accepted at Combinatorica, http://arxiv.org/abs/1603.00312.

[7] M. Axenovich and T. Ueckerdt. Density of range capturing hypergraphs. *Journal of Computational Geometry*, 7(1):1–21, 2016. http://dx.doi.org/10.20382/jocg.v7i1a1.

[8] M. Axenovich, T. Ueckerdt, and P. Weiner. Splitting planar graphs of girth 6 into two linear forests with short paths. *Journal of Graph Theory*, pages n/a–n/a, 2016. http://dx.doi.org/10.1002/jgt.22093.

[9] T. Bläsius, P. Stumpf, and T. Ueckerdt. Local and union boxicity. *CoRR*, abs/1609.09447, 2016. http://arxiv.org/abs/1609.09447.

[10] O. V. Borodin, A. Kostochka, and M. Yancey. On 1-improper 2-coloring of sparse graphs. *Discrete Mathematics*, 313(22):2638–2649, 2013.

[11] P. Brass, W. Moser, and J. Pach. Research problems in discrete geometry. 2005.

[12] J. Cardinal, K. B. Knauer, P. Micek, and T. Ueckerdt. Making triangles colorful. *Journal of Computational Geometry*, 4(1):240–246, 2013. http://jocg.org/index.php/jocg/article/view/136.

[13] J. Cardinal, K. B. Knauer, P. Micek, and T. Ueckerdt. Making octants colorful and related covering decomposition problems. *SIAM Journal on Discrete Mathematics*, 28(4):1948–1959, 2014. http://dx.doi.org/10.1137/140955975.

[14] J. Chalopin and D. Gonçalves. Every planar graph is the intersection graph of segments in the plane: Extended abstract. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 631–638, New York, NY, USA, 2009. ACM.

[15] S. Chaplick and T. Ueckerdt. Planar graphs as VPG-graphs. *Journal of Graph Algorithms and Applications*, 17(4):475–494, 2013. http://dx.doi.org/10.7155/jgaa.00300.

[16] O. Cheong, S. Har-Peled, H. Kim, and H.-S. Kim. On the number of edges of fan-crossing free graphs. *Algorithmica*, 73(4):673–695, 2015.

[17] L. P. Chew. Constrained Delaunay triangulations. *Algorithmica*, 4(1-4):97–108, 1989.

[18] B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit disk graphs. *Discrete mathematics*, 86(1-3):165–177, 1990.

[19] K. L. Clarkson and K. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry*, 37(1):43–58, 2007.

[20] M. B. Cozzens and F. S. Roberts. Computing the boxicity of a graph by covering its complement by cointerval graphs. *Discrete Applied Mathematics*, 6(3):217–228, 1983.

[21] H. De Fraysseix and P. O. de Mendez. On topological aspects of orientations. *Discrete Mathematics*, 229(1):57–72, 2001.

[22] H. De Fraysseix, P. O. De Mendez, and P. Rosenstiehl. On triangle contact graphs. *Combinatorics, Probability and Computing*, 3(02):233–246, 1994.

[23] M. B. Dillencourt. Toughness and Delaunay triangulations. *Discrete & Computational Geometry*, 5(6):575–601, 1990.

[24] W. Evans, S. Felsner, S. G. Kobourov, and T. Ueckerdt. Graphs admitting d-realizers: Spanning-tree-decompositions and box-representations. In *EuroCG 2014: 30th European Workshop on Computational Geometry, Ein-Gedi, Israel, March 3 - 5, 2014. Proceedings*, 2014. http://www.cs.bgu.ac.il/~eurocg14/papers/paper_11.pdf.

[25] S. Felsner, K. B. Knauer, G. B. Mertzios, and T. Ueckerdt. Intersection graphs of L-shapes and segments in the plane. In E. Csuhaj-Varjú, M. Dietzfelbinger, and Z. Ésik, editors, *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part II*, volume 8635 of *Lecture Notes in Computer Science*, pages 299–310. Springer, 2014. http://dx.doi.org/10.1007/978-3-662-44465-8_26.

[26] S. Felsner, K. B. Knauer, and T. Ueckerdt. Contact graphs of axis-aligned rectangles in 3D. work in progress, 2016.

[27] É. Fusy. Transversal structures on triangulations: A combinatorial study and straight-line drawings. *Discrete Mathematics*, 309(7):1870–1894, 2009.

[28] D. Gonçalves. Caterpillar arboricity of planar graphs. *Discrete mathematics*, 307(16):2112–2121, 2007.

[29] D. Gonçalves and P. Ochem. On star and caterpillar arboricity. *Discrete Mathematics*, 309(11):3694–3702, 2009. 7th International Colloquium on Graph TheoryICGT 05.

[30] S. L. Hakimi, J. Mitchem, and E. Schmeichel. Star arboricity of graphs. *Discrete Mathematics*, 149(1):93–98, 1996.

[31] M. Kaufmann and T. Ueckerdt. The density of fan-planar graphs. *CoRR*, abs/1403.6184, 2014. http://arxiv.org/abs/1403.6184.

[32] B. Keszegh, N. Lemons, and D. Pálvölgyi. Online and quasi-online colorings of wedges and intervals. *Order*, 33(3):389–409, 2016.

[33] B. Keszegh and D. Pálvölgyi. More on decomposing coverings by octants. *Journal of Computational Geometry*, 6(1):300–315, 2015.

[34] J. Klawitter, M. Nöllenburg, and T. Ueckerdt. Combinatorial properties of triangle-free rectangle arrangements and the squarability problem. In E. Di Giacomo and A. Lubiw, editors, *Graph Drawing and Network Visualization*, volume 9411 of *Lecture Notes in Computer Science*, pages 231–244. Springer International Publishing, 2015. http://dx.doi.org/10.1007/978-3-319-27261-0_20.

[35] M. Klazar. Extremal problems for ordered (hyper) graphs: applications of Davenport–Schinzel sequences. *European Journal of Combinatorics*, 25(1):125–140, 2004.

[36] K. B. Knauer, J. Rollin, and T. Ueckerdt. Corrected proof of "The interval number of a planar graph: Three intervals suffice". work in progress, 2016.

[37] K. B. Knauer and T. Ueckerdt. Three ways to cover a graph. *Discrete Mathematics*, 339(2):745 – 758, 2016. http://dx.doi.org/10.1016/j.disc.2015.10.023.

[38] S. G. Kobourov, T. Ueckerdt, and K. Verbeek. Combinatorial and geometric properties of planar Laman graphs. In S. Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1668–1678. SIAM, 2013. http://dx.doi.org/10.1137/1.9781611973105.120.

[39] A. Kostochka. Coloring intersection graphs of geometric figures with a given clique number. *Contemporary mathematics*, 342:127–138, 2004.

[40] J. Kratochvíl and J. Matoušek. Intersection graphs of segments. *Journal of Combinatorial Theory, Series B*, 62(2):289–315, 1994.

[41] J. Matoušek. Epsilon-nets and computational geometry. In *New Trends in Discrete and Computational Geometry*, pages 69–89. Springer, 1993.

[42] J. Matoušek. *Lectures on discrete geometry*, volume 108. Springer New York, 2002.

[43] M. Middendorf and F. Pfeiffer. The max clique problem in classes of string-graphs. *Discrete mathematics*, 108(1):365–372, 1992.

[44] K. G. Milans, D. Stolee, and D. B. West. Ordered Ramsey theory and track representations of graphs. *Journal of Combinatorics*, 6(4), 2015.

[45] J. Pach and G. Tóth. Graphs drawn with few crossings per edge. *Combinatorica*, 17(3):427–439, 1997.

[46] F. S. Roberts. On the boxicity and cubicity of a graph. In *Recent Progress in Combinatorics (Proc. Third Waterloo Conf. on Combinatorics, 1968)*, pages 301–310. Academic Press, New York, 1969.

[47] E. R. Scheinerman. *Intersection classes and multiple intersection parameters of graphs*. PhD thesis, Princeton University, 1984.

[48] E. R. Scheinerman and D. B. West. The interval number of a planar graph: Three intervals suffice. *Journal of Combinatorial Theory, Series B*, 35(3):224–239, 1983.

[49] W. Schnyder. Planar graphs and poset dimension. *Order*, 5(4):323–343, 1989.

Articles in the Habilitationsschrift

Articles on Coloring Problems:

[1] A. Asinowski, J. Cardinal, N. Cohen, S. Collette, T. Hackl, M. Hoffmann, K. B. Knauer, S. Langerman, M. Lason, P. Micek, G. Rote, and T. Ueckerdt. Coloring hypergraphs induced by dynamic point sets and bottomless rectangles. In F. Dehne, R. Solis-Oba, and J. Sack, editors, *Algorithms and Data Structures - 13th International Symposium, WADS 2013, London, ON, Canada, August 12-14, 2013. Proceedings*, volume 8037 of *Lecture Notes in Computer Science*, pages 73–84. Springer, 2013. http://dx.doi.org/10.1007/978-3-642-40104-6_7.

[2] M. Axenovich and T. Ueckerdt. Density of range capturing hypergraphs. *Journal of Computational Geometry*, 7(1):1–21, 2016. http://dx.doi.org/10.20382/jocg.v7i1a1.

[3] M. Axenovich, T. Ueckerdt, and P. Weiner. Splitting planar graphs of girth 6 into two linear forests with short paths. *Journal of Graph Theory*, pages n/a–n/a, 2016. http://dx.doi.org/10.1002/jgt.22093.

[4] J. Cardinal, K. B. Knauer, P. Micek, and T. Ueckerdt. Making triangles colorful. *Journal of Computational Geometry*, 4(1):240–246, 2013. http://jocg.org/index.php/jocg/article/view/136.

[5] J. Cardinal, K. B. Knauer, P. Micek, and T. Ueckerdt. Making octants colorful and related covering decomposition problems. *SIAM Journal on Discrete Mathematics*, 28(4):1948–1959, 2014. http://dx.doi.org/10.1137/140955975.

[6] M. Kaufmann and T. Ueckerdt. The density of fan-planar graphs. *CoRR*, abs/1403.6184, 2014. http://arxiv.org/abs/1403.6184.

Articles on Intersection Representations:

[1] S. Chaplick and T. Ueckerdt. Planar graphs as VPG-graphs. *Journal of Graph Algorithms and Applications*, 17(4):475–494, 2013. http://dx.doi.org/10.7155/jgaa.00300.

[2] W. Evans, S. Felsner, S. G. Kobourov, and T. Ueckerdt. Graphs admitting d-realizers: Spanning-tree-decompositions and box-representations. In *EuroCG 2014: 30th European Workshop on Computational Geometry, Ein-Gedi, Israel, March 3 - 5, 2014. Proceedings*, 2014. http://www.cs.bgu.ac.il/~eurocg14/papers/paper_11.pdf.

[3] S. Felsner, K. B. Knauer, G. B. Mertzios, and T. Ueckerdt. Intersection graphs of L-shapes and segments in the plane. In E. Csuhaj-Varjú, M. Dietzfelbinger, and Z. Ésik, editors, *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part II*, volume 8635 of *Lecture Notes in Computer Science*, pages 299–310. Springer, 2014. http://dx.doi.org/10.1007/978-3-662-44465-8_26.

[4] J. Klawitter, M. Nöllenburg, and T. Ueckerdt. Combinatorial properties of triangle-free rectangle arrangements and the squarability problem. In E. Di Giacomo and A. Lubiw, editors, *Graph Drawing and Network Visualization*, volume 9411 of *Lecture Notes in Computer Science*, pages 231–244. Springer International Publishing, 2015. http://dx.doi.org/10.1007/978-3-319-27261-0_20.

[5] S. G. Kobourov, T. Ueckerdt, and K. Verbeek. Combinatorial and geometric properties of planar Laman graphs. In S. Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1668–1678. SIAM, 2013. http://dx.doi.org/10.1137/1.9781611973105.120.

Articles on Covering Problems:

[1] T. Bläsius, P. Stumpf, and T. Ueckerdt. Local and union boxicity. *CoRR*, abs/1609.09447, 2016. http://arxiv.org/abs/1609.09447.

[2] K. B. Knauer and T. Ueckerdt. Three ways to cover a graph. *Discrete Mathematics*, 339(2):745 – 758, 2016. http://dx.doi.org/10.1016/j.disc.2015.10.023.

Articles on Coloring Problems

Coloring Hypergraphs Induced by Dynamic Point Sets and Bottomless Rectangles

Andrei Asinowski[1], Jean Cardinal[2], Nathann Cohen[3], Sébastien Collette[4], Thomas Hackl[5], Michael Hoffmann[6], Kolja Knauer[7], Stefan Langerman[2], Michał Lasoń[8], Piotr Micek[8], Günter Rote[9], and Torsten Ueckerdt[10]

[1] Freie Universität Berlin
asinowski@mi.fu-berlin.de
[2] Université Libre de Bruxelles
{jcardin,slanger}@ulb.ac.be
[3] Université Paris-Sud 11
nathann.cohen@gmail.com
[4] Université Libre de Bruxelles
me@scollette.com
[5] TU Graz
thackl@ist.tugraz.at
[6] ETH Zürich
hoffmann@inf.ethz.ch
[7] Université Montpellier 2
kolja.knauer@gmail.com
[8] Jagiellonian University in Krakow
{mlason,piotr.micek}@tcs.uj.edu.pl
[9] Freie Universität Berlin
rote@inf.fu-berlin.de
[10] Karlsruhe Institute of Technology
torsten.ueckerdt@kit.edu

Abstract. We consider a coloring problem on dynamic, one-dimensional point sets: points appearing and disappearing on a line at given times. We wish to color them with k colors so that at any time, any sequence of $p(k)$ consecutive points, for some function p, contains at least one point of each color.

We prove that no such function $p(k)$ exists in general. However, in the restricted case in which points appear gradually, but never disappear, we give a coloring algorithm guaranteeing the property at any time with $p(k) = 3k-2$. This can be interpreted as coloring point sets in $\mathbb{R}^2$ with k colors such that any bottomless rectangle containing at least $3k-2$ points contains at least one point of each color. Here a bottomless rectangle is an axis-aligned rectangle whose bottom edge is below the lowest point of the set. For this problem, we also prove a lower bound $p(k) > ck$, where $c > 1.67$. Hence, for every k there exists a point set, every k-coloring of which is such that there exists a bottomless rectangle containing ck points and missing at least one of the k colors.

Chen *et al.* (2009) proved that no such function $p(k)$ exists in the case of general axis-aligned rectangles. Our result also complements recent results from Keszegh and Pálvölgyi on cover-decomposability of octants (2011, 2012).

F. Dehne, R. Solis-Oba, and J.-R. Sack (Eds.): WADS 2013, LNCS 8037, pp. 73–84, 2013.

1 Introduction

It is straightforward to color n points lying on a line with k colors in such a way that any set of k consecutive points receive different colors; just color them cyclically with the colors $1, 2, \ldots, k, 1, \ldots$. What can we do if points can appear and disappear on the line, and we wish a similar property to hold at any time? More precisely, we fix the number k of colors, and wish to maintain the property that at any given time, any sequence of $p(k)$ consecutive points, for some function p, contains at least one point of each color.

We show that in general, such a function does not exist: there are dynamic point sets on a line that are impossible to color with two colors so that monochromatic subsequences have bounded length. This holds even if the whole schedule of appearances and disappearances is known in advance. This family of point sets is described in Section 2.

We prove, however, that there exists a linear function p in the case where points can appear on the line at any time, but *never disappear*. Furthermore, this is achieved in a constructive, *semi-online* fashion: the coloring decision for a point can be delayed, but at any time the currently colored points yield a suitable coloring of the set. The algorithm is described in Section 3.

In Section 4, we restate the result in terms of a coloring problem in $\mathbb{R}^2$: for any integer $k \geq 1$, every point set in $\mathbb{R}^2$ can be colored with k colors so that any *bottomless* rectangle containing at least $3k - 2$ points contains one point of each color. Here, an axis-aligned rectangle is said to be bottomless whenever the y-coordinate of its bottom edge is $-\infty$.

In Section 5, we give lower bounds on the problem of coloring points with respect to bottomless rectangles. We show that the number of points $p(k)$ contained in a bottomless rectangle must be at least $1.67k$ in order to guarantee the presence of at least one point of each color.

Finally, in Section 6, we consider an alternative problem in which we fix the size of the sequence to k, but we are allowed to increase the number of colors.

Motivation and previous works. The problem is motivated by previous intriguing results in the field of geometric hypergraph coloring. Here, a geometric hypergraph is a set system defined by a set of points and a set of geometric ranges, typically polygons, disks, or pseudodisks. Every hyperedge of the hypergraph is the intersection of the point set with a range.

It was shown recently [7] that for every convex polygon P, there exists a constant c, such that any point set in $\mathbb{R}^2$ can be colored with k colors in such a way that any translation of P containing at least $p(k) = ck$ points contains at least one point of each color. This improves on several previous intermediate results [15,17,2]. Similar positive results for other families of geometric hypergraphs are given by Aloupis et al. [3,1], and Smorodinsky and Yuditsky [18]. Discussions on the relation between this coloring problem and ε-nets can be found in Pach and Tardos [13].

The problem for translates of polygons can be cast in its dual form as a covering decomposition problem: given a set of translates of a polygon P, we wish to color them with k colors so that any point covered by at least $p(k)$ of them is covered by at least one of each color. The two problems can be seen to be equivalent by replacing the points by translates of a symmetric image of P centered on these points. The covering decomposition problem has a long history that dates back to conjectures by János Pach in the early 80s (see for instance [11,4], and references therein). The decomposability of coverings by unit disks was considered in a seemingly lost unpublished manuscript by Mani and Pach in 1986. Up to recently, however, surprisingly little was known about this problem.

For other classes of ranges, such as axis-aligned rectangles, disks, translates of some concave polygons, or arbitrarily oriented strips [5,12,14,16], such a coloring does not always exists, even when we restrict ourselves to two colors.

Keszegh [8] showed in 2007 that every point set could be 2-colored so that any bottomless rectangle containing at least 4 points contains both colors. Our positive result on bottomless rectangles (Corollary 2) is a generalization of Keszegh's results to k-colorings. Later, Keszegh and Pálvölgyi [9] proved the following cover-decomposability property of octants in $\mathbb{R}^3$: every collection of translates of the positive octant can be 2-colored so that any point of $\mathbb{R}^3$ that is covered by at least 12 octants is covered by at least one of each color. This result generalizes the previous one (with a looser constant), as incidence systems of bottomless rectangles in the plane can be produced by restricted systems of octants in $\mathbb{R}^3$. It also implies similar covering decomposition results for homothetic copies of a triangle. More recently, they generalized their result to k-colorings, and proved an upper bound of $p(k) < 12^{2^k}$ on the corresponding function $p(k)$ [10].

2 Coloring Dynamic Point Sets

A *dynamic point set* S in $\mathbb{R}$ is a collection of triples $(v_i, a_i, d_i) \in \mathbb{R}^3$, with $d_i \geq a_i$, that is interpreted as follows: the point $v_i \in \mathbb{R}$ appears on the real line at time a_i and disappears at time d_i. Hence, the set $S(t)$ of points that are present at time t are the points v_i with $t \in [a_i, d_i)$. A k-coloring of a dynamic point set assigns one of k colors to each such triple.

We now show that it is not possible to find a 2-coloring of such a point set while avoiding long monochromatic subsequences at any time.

Theorem 1. *For every $p \in \mathbb{N}$, there exists a dynamic point set S with the following property: for every 2-coloring of S, there exists a time t such that $S(t)$ contains p consecutive points of the same color.*

Proof. In order to prove this result, we work on an equivalent two-dimensional version of the problem. From a dynamic point set, we can build n horizontal segments in the plane, where the ith segment goes from (a_i, v_i) to (d_i, v_i). At any time t the visible points $S(t)$ correspond to the intervals that intersect the line $x = t$. It is therefore equivalent, in order to obtain our result, to build a collection of horizontal segments in the plane that cannot be 2-colored in such a way that any set of p segments intersecting some vertical segment contains one element of each color.

Our construction borrows a technique from Pach, Tardos, and Tóth [14]. In this paper, the authors provide an example of a set system whose base set cannot be 2-colored without leaving some set monochromatic. This set system $\mathcal{S}$ is built on top of the $1 + p + \cdots + p^{p-1} = \frac{1-p^p}{1-p}$ vertices of a p-regular tree $\mathcal{T}^p$ of depth p, and contains two kinds of sets :

- the $1 + p + \cdots + p^{p-2}$ sets of *siblings*: the sets of p vertices having the same father,
- the p^{p-1} sets of p vertices corresponding to a path from the root vertex to one of the leaves in $\mathcal{T}^p$.

It is not difficult to realize that this set system is not 2-colorable: by contradiction, if every set of siblings is non-monochromatic, we can greedily construct a monochromatic path from the root to a leaf.

We now build a collection $\mathcal{S}$ of horizontal segments corresponding to the vertices of $\mathcal{T}^p$, in such a way that for any set $E \in \mathcal{S}$ there exists a time t at which the elements of E are consecutive among those that intersect the line $x = t$. For any p (see Fig. 1), the construction starts with a building block B_p^1 of p horizontal segments, the ith segment going from $(-\frac{i}{p}, i)$ to $(0, i)$. Because these p segments represent *siblings* in $\mathcal{T}^p$, they are consecutive on the vertical line that goes through their rightmost endpoint, and hence cannot all receive the same color.

Block B_p^{j+1} is built from a copy of B_p^1 to which are added p resized and translated copies of B_p^j : the ith copy lies in the rectangle with top-right corner $(-\frac{i-1}{p}, i+1)$ and bottom-left corner $(-\frac{i}{p}, i)$. By adding to B_p^{p-1} a last horizontal segment below all others, corresponding to the root of $\mathcal{T}^p$, the ancestors of a segment are precisely those that are below it on the vertical line that goes through its leftmost point. When such sets of ancestors are of cardinality $p - 1$, which only happens when one considers the set of ancestors of a leaf, then the set formed by the leaf and its ancestors is required to be non-monochromatic.

With this construction we ensure that a feasible 2-coloring of the segments would yield a proper 2-coloring of $\mathcal{S}$, which we know does not exist. □

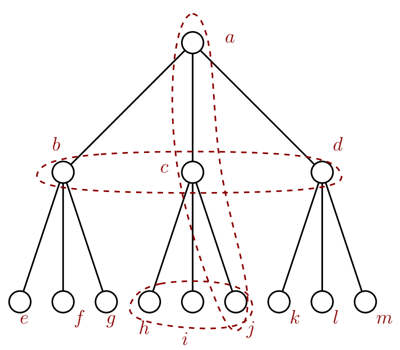

(a) The tree $\mathcal{T}^3$.

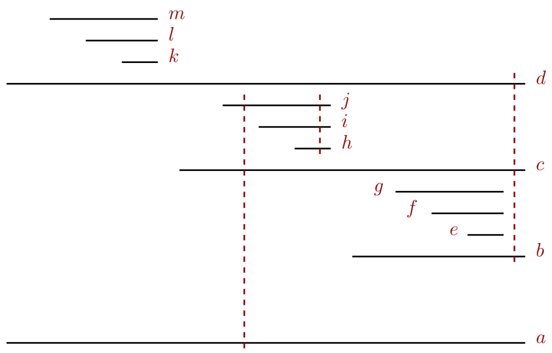

(b) The corresponding set of horizontal segments B_3^2, with a root segment a.

Fig. 1. The recursive construction of Theorem 1, for $p = 3$

The above result implies that no function $p(k)$ exists for any k that answers the original question. If it were the case, then we could simply merge color classes of a k-coloring into two groups and contradict the above statement. Theorem 1 can also be interpreted as the indecomposability of coverings by a specific class of unbounded polytopes in $\mathbb{R}^3$. We define a *corner* with coordinates (a, b, c) as the following subset of $\mathbb{R}^3$: $\{(x, y, z) \in \mathbb{R}^3 : a \leq x \leq b, y \leq c \leq z\}$. An example is given in Fig. 2. One can verify that a point (x, y, z) is contained in a corner a, b, c if and only if the vertical line segment with endpoints (x, y) and (x, z) intersects the horizontal line segment with endpoints (a, c) and (b, c). The corollary follows.

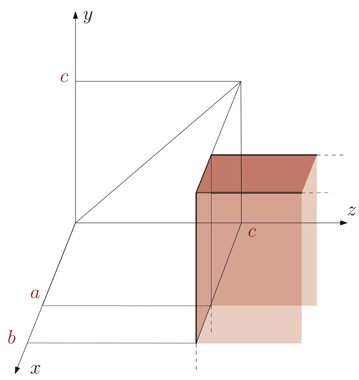

Fig. 2. A corner with coordinates (a, b, c)

Corollary 1. *For every $p \in \mathbb{N}$, there exists a collection S of corners with the following property: for every 2-coloring of S, there exists a point $x \in \mathbb{R}^3$ contained in exactly p corners of S, all of the same color. In other words, corners are not cover-decomposable.*

3 Coloring Point Sets under Insertion

Since we cannot bound the function $p(k)$ in the general case, we now consider a simple restriction on our dynamic point sets: we let the deletion times d_i be infinite for every i. Hence, points appear on the line, but never disappear.

A natural idea to tackle this problem is to consider an online coloring strategy, that would assign a color to each point in order of their arrival times a_i, without any knowledge of the points appearing later. However, we cannot guarantee any bound on $p(k)$ unless we delay some of the coloring decisions. To see this, consider the case $k = 2$, and call the two colors red and blue. An online algorithm must color each new point in red or blue as soon as it is presented. We can design an adversary such that the following invariant holds: at any time, the set of points is composed of a sequence of consecutive red points, followed by a sequence of consecutive blue points. The adversary simply chooses the new point to lie exactly between the two sequences at each step.

Our computation model will be *semi-online*: The algorithm considers the points in their order of the arrival time a_i. At any time, a point in the sequence either has one of the k colors, or is uncolored. Uncolored points can be colored later, but once a point is colored, it keeps its color for the rest of the procedure. At any time, the colors that are already assigned suffice to satisfy the property that any subsequence of $3k - 2$ points has one point of each color, i.e., $p(k) \leq 3k - 2$.

Theorem 2. *Every dynamic point set without disappearing points can be k-colored in the semi-online model such that at any time, every subsequence of at least $3k-2$ consecutive points contains at least one point of each color.*

Proof. We define a *gap for color i* as a maximal interval (set of consecutive points) containing no point of color i, that is, either between two successive occurrences of color i, or before the first occurrence (first gap), or after the last occurrence (last gap), or the whole line if no point has color i. A *gap* is simply a gap for color i, for some $1 \leq i \leq k$. We propose an algorithm for the semi-online model keeping the sizes of all gaps to be at most $3k-3$. This means every set of $3k-2$ consecutive points contains each color at least once and implies $p(k) \leq 3k-2$. The algorithm maintains two invariants:
(a) every gap contains at most $3k-3$ points; (b) if there is some point colored with i then every gap for color i, except the first and the last gap, contains at least $k-1$ points.

The two invariants are vacuous when the set of points is empty. Now, suppose that the invariants hold for an intermediate set of points and consider a new point on the line presented by an adversary. Clearly, invariant *(b)* cannot be violated in the extended set as no gaps decrease in size. However, there may arise some gaps of size $3k-2$ violating *(a)*. If not then the invariants hold for the extended set and the algorithm does not color any point in this step. Suppose there are some gaps of size $3k-2$. Consider one of them, say a gap of color i, and denote the points in the gap in their natural ordering on the line from left to right as $(\ell_1, \dots, \ell_{k-1}, m_1, \dots, m_k, r_1, \dots r_{k-1})$. Now, color i does not appear among these points. Invariant *(b)* yields that none of the $k-1$ remaining colors appears twice among $m_1, \dots, m_k$. Thus, there is some m_j, which is uncolored and the algorithm colors it with i. This splits the large gap into two smaller gaps. Moreover, since there are $k-1$ ℓ-points and $k-1$ r-points invariant *(b)* is maintained for both new i-gaps. The algorithm repeats that process until all gaps are of size at most $3k-3$.

This concludes the proof, as after the algorithm ends all remaining uncolored points can be arbitrary colored. □

4 Coloring Points with Respect to Bottomless Rectangles

A *bottomless rectangle* is a set of the form $\{(x, y) \in \mathbb{R}^2 : a \leq x \leq b, y \leq c\}$, for a triple of real numbers (a, b, c) with $a \leq b$. We consider the following geometric coloring problem: given a set of points in the plane, we wish to color them with k colors so that any bottomless rectangle containing at least $p(k)$ points contains at least one point of each color. It is not difficult to realize that the problem is equivalent to that of the previous section.

Corollary 2. *Every point set $S \subset \mathbb{R}^2$ can be colored with k colors so that any bottomless rectangle containing at least $3k-2$ points of S contains at least one point of each color.*

Proof. The algorithm proceeds by sweeping S vertically in increasing y-coordinate order. This defines a dynamic point set S' that contains at time t the x-coordinates of the points below the horizontal line of equation $y = t$. The set of points of S that are contained in a bottomless rectangle $\{(x, y) \in \mathbb{R}^2 : a \leq x \leq b, y \leq t\}$ correspond to the points in the interval $[a, b]$ in $S'(t)$. Hence, the two coloring problems are equivalent, and Theorem 2 applies. □

5 Lower Bound

We now give a lower bound on the smallest possible value of $p(k)$.

Theorem 3. *For any k sufficiently large, there exists a point set P such that for any k-coloring of P, there exists a color $i \in [k]$ and a bottomless rectangle containing at least $1.677k - 2.5$ points, none of which are colored with color i.*

Proof. Fix $k \geq 100$. For $n \in \mathbb{N}$ and $0 \leq a < k$ we define the point set $P = P(n, a)$ to be the union of point sets L, R and B (standing for left, right and bottom, respectively) as follows:

$$L := \{(i - n, 2i - 1) \in \mathbb{R}^2 \mid i \in [n]\}$$
$$B := \{(i, 0) \in \mathbb{R}^2 \mid i \in [a]\}$$
$$R := \{(a + i, 2n + 2 - 2i) \in \mathbb{R}^2 \mid i \in [n]\}$$

See Figure 3(a) for an illustration. Note that $|L| = |R| = n$ and $|B| = a$. Consider

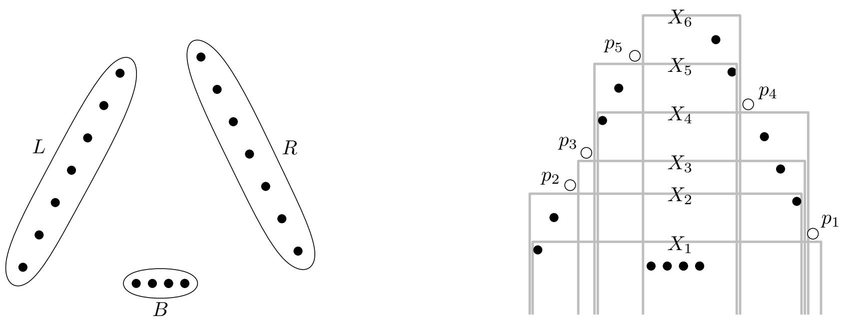

(a) (b)

Fig. 3. (a) The point set $P = P(n, a)$ with $n = 7$ and $a = 4$, and (b) the bottomless rectangles $X_1, \ldots, X_6$ corresponding to the color class $P(c^*) = \{p_1, \ldots, p_5\}$

any coloring of the points in P with colors from $[k]$. For a color $i \in [k]$ we define $P(i)$ to be the subset of points of P colored with i. We assume for the sake of contradiction that every bottomless rectangle that contains $b := \lfloor 1.677k - 2.5 \rfloor$ points, contains one point of each color. In the remainder of the proof we will identify a bottomless rectangle containing b' points but no point of one particular color. We give a lower bound for b' depending on n and a, but independent of the fixed coloring under consideration. Taking sufficiently large n and choosing $a = \lfloor 0.655k \rfloor$ we will prove $b' > b$, which contradicts our assumption and hence concludes the proof.

A color used at least once for the points in B is called a *low color* and a point colored with a low color is a *low point*. Note that there are low points outside of the set B. Let ℓ be the number of low colors. Clearly, $\ell \leq |B| = a$.

Claim 1.

(i) *For every non-low color c there are at least $\left\lfloor \frac{n}{b-a} \right\rfloor$ points of color c in L.*

(ii) *There are at least $\sum_{i=0}^{\ell-1} \left\lfloor \frac{n}{b-i} \right\rfloor$ low points in L.*

Proof. Fix a color $c \in [k]$ and assume that the j leftmost points in B are not colored with c. Order the points in L colored with c according to their x-coordinate: p_1, p_2,..., p_m. Now for each $1 < i \leq m$ there is a bottomless rectangle containing all points in L between p_{i-1} and p_i, and the leftmost j points in B, and nothing else. Additionally, there is a bottomless rectangle containing all points in L to the left of p_1 together with j leftmost points in B, and a bottomless rectangle containing all points in L to the right of p_m together with j leftmost points in B. Note that all these rectangles are disjoint within L and each point from L not colored with c lies in exactly one such rectangle. Since each such rectangle X avoids the color c we get that $|X \cap P| \leq b-1$ and $|X \cap L| \leq b-1-j$ and therefore

$$m + (m+1)(b-1-j) = m(b-j) + b - j - 1 \geq |L| = n,$$
$$m \geq \left\lfloor \frac{n}{b-j} \right\rfloor . \tag{1}$$

In order to prove (i) consider a non-low color c. As c is not used on points in B at all we can put $j = a$ in (1) and the statement of (i) follows. Now, if c is a low color, then j defined as the maximum number of leftmost points in B avoiding c is always less than a. However, for each low color c we obtain a different j. Thus the sum of inequality (1) over all low colors is minimized by $\sum_{i=0}^{\ell-1} \lfloor \frac{n}{b-i} \rfloor$, which gives (ii). □

By Claim 1 (i) and (ii) combined we get that there is a set S of $k-a$ non-low colors such that at most $n - \sum_{i=0}^{a-1} \lfloor \frac{n}{b-i} \rfloor$ points in L have a color from S. Analogously, at most $n - \sum_{i=0}^{a-1} \lfloor \frac{n}{b-i} \rfloor$ points in R have a color from S. Summing up we get:

$$\begin{aligned}\sum_{c\in S}|P(c)| &= \sum_{c\in S}\big(|P(c)\cap L| + |P(c)\cap R|\big)\\ &\le 2n - 2\sum_{i=0}^{a-1}\left\lfloor \frac{n}{b-i}\right\rfloor \le 2n - 2\sum_{i=0}^{a-1}\left(\frac{n}{b-i} - 1\right)\\ &= 2n\left(1 - \sum_{i=b-a+1}^{b}\frac{1}{i}\right) + 2a\\ &= 2n\left(1 - \sum_{i=1}^{b}\frac{1}{i} + \sum_{i=1}^{b-a}\frac{1}{i}\right) + 2a.\end{aligned}$$

Using that $\sum_{i=1}^{x}\frac{1}{i} = \ln(x+1) - \sum_{j=1}^{\infty}\frac{B_j}{j(x+1)^j} + \gamma$ for every $x \ge 1$, where B_j are the second Bernoulli numbers and γ is the Euler-Mascheroni constant, we obtain

$$\begin{aligned}\sum_{c\in S}|P(c)| &< 2n\,(1 - \ln(b+1) + \ln(b-a+1)) + 2a\\ &= 2n\left(1 - \ln\left(\frac{b+1}{b-a+1}\right)\right) + 2a.\end{aligned}$$

From the pigeonhole principle we know that there has to exist a color $c^* \in S$, such that

$$q := |P(c^*)| \le \left\lfloor \frac{2n(1-\ln(\frac{b+1}{b-a+1})) + 2a}{k-a}\right\rfloor. \tag{2}$$

Enumerate the points in $P(c^*)$ by $p_1, p_2, \ldots, p_q$ according to their increasing y-coordinates, i.e., we have $i < j$ iff p_i has smaller y-coordinate than p_j. Now we consider all maximal bottomless rectangles that completely contain B and contain no point of color c^*. There are exactly $q+1$ such rectangles: For every point $p_i \in P(c^*)$ there is a bottomless rectangle X_i whose top side lies immediately below p_i. And one further bottomless rectangle X_{q+1} containing the entire strip between L and R, and with sides bounded by the point in $P(c^*) \cap L$ and the point in $P(c^*) \cap R$ with the highest index. See Figure 3(b) for an illustration.

Claim 2. $\sum_{i=1}^{q}|X_i \cap (L \cup R)| \ge \frac{3}{2}\big(2n - q - b + a\big)$.

Proof. Let Y_1 and Y_{q+1} be the sets of points in $L \cup R$ with y-coordinate smaller than p_1 and larger than p_q, respectively. Let Y_i, $2 \le i \le q$, be the set of points with y-coordinate between p_{i-1} and p_i. Note that $Y_i \subset X_i \cap (L \cup R)$ for all $1 \le i \le q+1$, and that the $q+1$ sets $Y_1, \ldots, Y_{q+1}$ partition the points of $L \cup R$ that are not colored with c^*. Clearly, $|X_i \cap Y_i| = |Y_i|$. We claim that $|X_{i+1} \cap Y_i| \ge \frac{1}{2}|Y_i|$, for $i = 1, \ldots, q$.

Without loss of generality, let us assume that $p_i \in L$. Then either $Y_i = \emptyset$ or the point in Y_i with largest y-coordinate lies in R. Since points from L and R alternate in the ordering of $L \cup R$ with respect to increasing y-coordinate it follows that Y_i is almost equally partitioned into its left part $Y_i \cap L$ and its right part $Y_i \cap R$. Since the topmost point in Y_i lies in R we have $|Y_i \cap R| \geq \frac{1}{2}|Y_i|$. Now since $p_i \in L$ we have $X_{i+1} \supset Y_i \cap R$, and thus

$$|X_{i+1} \cap Y_i| \geq |Y_i \cap R| \geq \frac{1}{2}|Y_i|. \tag{3}$$

Note also that $|X_{q+1} \cap Y_q| + |Y_{q+1}| \leq |X_{q+1} \cap (L \cup R)| < b - a$ as X_{q+1} avoids color c^*, so $|X_{q+1}| < b$, and contains all a points in B.

Now we calculate

$$\begin{aligned}\sum_{i=1}^{q} |X_i \cap (L \cup R)| &\geq \Big(\sum_{i=1}^{q} |X_i \cap Y_i| + |X_{i+1} \cap Y_i|\Big) - |X_{q+1} \cap Y_q| \\ &\overset{(3)}{\geq} \sum_{i=1}^{q} \frac{3}{2}|Y_i| - |X_{q+1} \cap Y_q| \\ &= \frac{3}{2}\big(2n - |P(c^*)| - |Y_{q+1}|\big) - |X_{q+1} \cap Y_q| \\ &\geq \frac{3}{2}\big(2n - q - (|Y_{q+1}| + |X_{q+1} \cap Y_q|)\big) \geq \frac{3}{2}\big(2n - q - (b - a)\big).\end{aligned}$$

□

From Claim 2 we get from the pigeonhole principle that there is a bottomless rectangle $X^* \in \{X_1, \ldots, X_q\}$ with

$$\begin{aligned}|X^*| &\geq \frac{\frac{3}{2}(2n - q - b + a)}{q} + a = \frac{3n}{q} - \frac{3}{2} - \frac{3(b-a)}{2q} + a \\ &\overset{(2)}{\geq} \frac{3(k-a)}{2\big(1 - \ln\big(\frac{b+1}{b-a+1}\big) + \frac{2a}{n}\big)} + a - \frac{3}{2} - \frac{3(b-a)}{2q}\end{aligned}$$

Now, if we increase n, then $q = |P(c^*)|$ increases as well, and for sufficiently large n the terms $\frac{2a}{n}$ in the denominator and the additive term $\frac{3(b-a)}{2q}$ become negligible. In particular, with $a := \lfloor 0.655k \rfloor$ and $b = \lfloor 1.677k - 2.5 \rfloor$ and sufficiently large n we have

$$\begin{aligned}|X^*| &\geq \frac{3(k-a)}{2\big(1 - \ln\big(\frac{b+1}{b-a+1}\big)\big)} + a - \frac{3}{2} \\ &= \frac{3(k - \lfloor 0.655k \rfloor)}{2\big(1 - \ln\big(\frac{\lfloor 1.677k - 2.5 \rfloor + 1}{\lfloor 1.677k - 2.5 \rfloor - \lfloor 0.655k \rfloor + 1}\big)\big)} + \lfloor 0.655k \rfloor - \frac{3}{2} \\ &\sim \Big(\frac{1.035}{2\big(1 - \ln\big(\frac{1.677}{1.022}\big)\big)} + 0.655\Big)k > 1.68k.\end{aligned}$$

Hence if k is big enough ($k \geq 100$ is actually enough) the bottomless rectangle X^* contains strictly more than $1.677k - 2.5$ points but no point of color c^*, which is a contradiction and concludes the proof. □

6 Increasing the Number of Colors

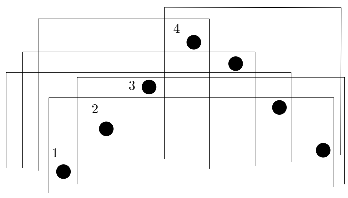

Fig. 4. A point set witnessing $c(k) \geq 2k-1$ for $k = 4$

There is another problem which can be tackled this time in an *online* model. The number $c(k)$ is the minimum number of colors needed to color the points on a line such that any set of at most k consecutive points is completely colored by distinct colors. The same problem has been considered for other types of geometric hypergraphs by Aloupis et al. [3]. Again, the algorithm considers the points in their order of the arrival time a_i but now colors them immediately.

Proposition 1. *Every dynamic point set without disappearing points can be* $(2k-1)$*-colored in the online model such that at any time, every subsequence of at least* k *consecutive points contains no color twice.*

Proof. At the arrival of a new point p denote by $(\ell_1, \ldots, \ell_{k-1})$ and $(r_1, \ldots, r_{k-1})$ the $k-1$ points to its left and to its right, respectively. Together they have at most $2k-2$ colors, Thus, there is at least one of the $2k-1$ colors unused among these points. The algorithm colors p with this color. □

Corollary 3. *Every point set* $S \subset \mathbb{R}^2$ *can be colored with* $2k-1$ *colors so that any bottomless rectangle containing at least* k *points of* S *contains no color twice.*

The number of colors used in Corollary 3 is smallest possible. This is witnessed by a point set S consisting of k points of the form $\{(i, 2i) \mid 0 \leq i \leq k-1\}$ and $k-1$ points of the form $\{(2k-i, 2i-1) \mid 1 \leq i \leq k-1\}$, see Fig. 4 for an example. It is easy to see that every pair of points in such a point set is in a common bottomless rectangle of size at most k. Finally, let us remark that an upper bound on $c(k)$ for dynamic point sets in which points can both appear and disappear, as in Section 2, can be obtained by bounding the chromatic number of the corresponding so-called *bar k-visibility graph*, as defined by Dean et al. [6]. In particular, they show that those graphs have $O(kn)$ edges, yielding $c(k) = O(k)$ for that case.

Acknowledgments. This research is supported by the ESF EUROCORES programme EuroGIGA, CRP ComPoSe, the Austrian Science Fund (FWF): P23629-N18 "Combinatorial Problems on Geometric Graphs" (Thomas Hackl), CRP GraDR and the Swiss National Science Foundation, SNF Project 20GG21-134306 (Michael Hoffmann). It was initiated at the ComPoSe kickoff meeting held at CIEM in Castro de Urdiales (Spain) on May 23–27, 2011, and pursued at the 2nd ComPoSe Workshop held at TU Graz (Austria) on April 16–20, 2012. The authors warmly thank the organizers of these two meetings as well as all the other participants. Part of this work was done during a stay of Kolja Knauer, Piotr Micek, and Torsten Ueckerdt at ULB (Brussels) and supported as EUROCORES short-term visit. A preliminary version was presented by a subset of the authors at EuroCG'12 in Assisi (Italy).

References

1. Aloupis, G., Cardinal, J., Collette, S., Imahori, S., Korman, M., Langerman, S., Schwartz, O., Smorodinsky, S., Taslakian, P.: Colorful strips. Graphs and Combinatorics 27(3), 327–339 (2011)
2. Aloupis, G., Cardinal, J., Collette, S., Langerman, S., Orden, D., Ramos, P.: Decomposition of multiple coverings into more parts. Discrete & Computational Geometry 44(3), 706–723 (2010)
3. Aloupis, G., Cardinal, J., Collette, S., Langerman, S., Smorodinsky, S.: Coloring geometric range spaces. Discrete & Computational Geometry 41(2), 348–362 (2009)
4. Brass, P., Moser, W.O.J., Pach, J.: Research Problems in Discrete Geometry. Springer (2005)
5. Chen, X., Pach, J., Szegedy, M., Tardos, G.: Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles. Random Struct. Algorithms 34(1), 11–23 (2009)
6. Dean, A.M., Evans, W., Gethner, E., Laison, J.D., Safari, M.A., Trotter, W.T.: Bar k-visibility graphs. J. Graph Algorithms Appl. 11(1), 45–59 (2007)
7. Gibson, M., Varadarajan, K.R.: Optimally decomposing coverings with translates of a convex polygon. Discrete & Computational Geometry 46(2), 313–333 (2011)
8. Keszegh, B.: Weak conflict-free colorings of point sets and simple regions. In: CCCG, pp. 97–100 (2007)
9. Keszegh, B., Pálvölgyi, D.: Octants are cover-decomposable. Discrete & Computational Geometry 47(3), 598–609 (2012)
10. Keszegh, B., Pálvölgyi, D.: Octants are cover-decomposable into many coverings. CoRR, abs/1207.0672 (2012)
11. Pach, J.: Covering the plane with convex polygons. Discrete & Computational Geometry 1, 73–81 (1986)
12. Pach, J., Tardos, G.: Coloring axis-parallel rectangles. J. Comb. Theory, Ser. A 117(6), 776–782 (2010)
13. Pach, J., Tardos, G.: Tight lower bounds for the size of epsilon-nets. In: Proceedings of the 27th Annual ACM Symposium on Computational Geometry, SoCG 2011, pp. 458–463 (2011)
14. Pach, J., Tardos, G., Tóth, G.: Indecomposable coverings. In: CJCDGCGT, pp. 135–148 (2005)
15. Pach, J., Tóth, G.: Decomposition of multiple coverings into many parts. Comput. Geom. 42(2), 127–133 (2009)
16. Pálvölgyi, D.: Indecomposable coverings with concave polygons. Discrete & Computational Geometry 44(3), 577–588 (2010)
17. Pálvölgyi, D., Tóth, G.: Convex polygons are cover-decomposable. Discrete & Computational Geometry 43(3), 483–496 (2010)
18. Smorodinsky, S., Yuditsky, Y.: Polychromatic coloring for half-planes. J. Comb. Theory, Ser. A 119(1), 146–154 (2012)

DENSITY OF RANGE CAPTURING HYPERGRAPHS

Maria Axenovich,[*] *and Torsten Ueckerdt*[*]

Abstract. For a finite set X of points in the plane, a set S in the plane, and a positive integer k, we say that a k-element subset Y of X is captured by S if there is a homothetic copy S' of S such that $X \cap S' = Y$, i.e., S' contains exactly k elements from X. A k-uniform S-capturing hypergraph $\mathcal{H} = \mathcal{H}(X, S, k)$ has a vertex set X and a hyperedge set consisting of all k-element subsets of X captured by S. In case when $k = 2$ and S is convex these graphs are planar graphs, known as *convex distance function Delaunay graphs*.

In this paper we prove that for any $k \geq 2$, any X, and any convex compact set S, the number of hyperedges in $\mathcal{H}(X, S, k)$ is at most $(2k-1)|X| - k^2 + 1 - \sum_{i=1}^{k-1} a_i$, where a_i is the number of i-element subsets of X that can be separated from the rest of X with a straight line. In particular, this bound is independent of S and indeed the bound is tight for all "round" sets S and point sets X in general position with respect to S.

This refines a general result of Buzaglo, Pinchasi and Rote [2] stating that every pseudodisc topological hypergraph with vertex set X has $O(k^2|X|)$ hyperedges of size k or less.

Keywords: Hypergraph density, geometric hypergraph, range-capturing hypergraph, homothets, convex distance function, Delaunay graph.

1 Introduction

Let S and X be two subsets of the Euclidean plane $\mathbb{R}^2$ and k be a positive integer. In this paper, S is a convex compact set containing the origin and X is a finite set. An ***S*-*range*** is a homothetic copy of S, i.e., a set obtained from S by scaling with a positive factor with respect to the origin and an arbitrary translation. In other words, an S-range is obtained from S by first contraction or dilation and then translation, where two S-ranges are *contractions/dilations* of one-another if in their corresponding mappings the origin is mapped to the same point.

We say that an S-range S' *captures* a subset Y of X if $X \cap S' = Y$. An ***S*-capturing hypergraph** is a hypergraph $\mathcal{H} = (X, \mathcal{E})$ with vertex set X and edge set $\mathcal{E} \subseteq 2^X$ such that for every $Y \in \mathcal{E}$ there is an S-range S' that captures E.

In this paper we consider ***k*-uniform** S-capturing hypergraphs, that is, those hypergraphs $\mathcal{H} = \mathcal{H}(X, S, k)$ with vertex set X and hyperedge set consisting of *all* k-element subsets of X captured by S. I.e., the hyperedges correspond to S-ranges containing exactly k elements from X. These hypergraphs are often referred to as *range-capturing hypergraphs*

[*] *Karlsruher Institut für Technologie*, {maria.aksenovich,torsten.ueckerdt}@kit.edu

or *range spaces*. The importance of studying k-uniform S-capturing hypergraphs was emphasized by their connection to epsilon nets and covering problems of the plane [12, 13, 16]. See also some related literature for geometric hypergraphs, [1–4, 8–11, 14, 15, 17–19, 23, 24, 26].

The first non-trivial case $k = 2$, i.e., when $\mathcal{H}(X, S, k)$ is an ordinary graph, was first considered by Chew and Dyrsdale in 1985 [7]. They showed that if S is convex and compact, then $\mathcal{H}(X, S, 2)$ is a planar graph, called the *Delaunay graph* of X for the *convex distance function* defined by S. In particular, $\mathcal{H}(X, S, 2)$ has at most $3|X| - 6$ edges and this bound can be achieved. We remark that it follows from Schnyder's realizer [21] that every maximally planar graph can be written as $\mathcal{H}(X, S, 2)$ for some X and S being any triangle.

1.1 Related work.

Recently, Buzaglo, Pinchasi and Rote [2] considered the maximum number of hyperedges of size k or less in a pseudodisc topological hypergraph on n vertices. Here, a *family of pseudodiscs* is a set of closed Jordan curves such that any two of these curves either do not intersect or intersect in exactly two points. A hypergraph is called *pseudodisc topological hypergraph* if its vertex set X is a set of points in the plane and for every hyperedge Y there is a closed Jordan curve such that the bounded region of the plane obtained by deleting the curve contains Y and no point from $X \setminus Y$, and the set of all these Jordan curves is a family of pseudodiscs.

The authors of [2] observed that pseudodisc topological hypergraphs have VC-dimension [25] at most 3, and that using this fact the number of hyperedges can be bounded from above. For this, a version of the Perles-Sauer-Shelah theorem [20, 22] is applied. Let, for a set A and a positive integer d, $\binom{A}{\leq d}$ denote the set of all subsets of A of size at most d.

Theorem 1 (Perles-Sauer-Shelah Theorem). *Let $F = \{A_1, \ldots, A_m\}$ be a family of distinct subsets of $\{1, 2, \ldots, n\}$ and let F have VC-dimension at most d. Then $m \leq \left|\bigcup_{i=1}^{m} \binom{A_i}{\leq d}\right|$.*

Applying this theorem to the family of hyperedges in a pseudodisc topological hypergraph, one can see that the number of hyperedges in such a hypergraph is at most $O(n^3)$. In fact, if one considers only hyperedges of size k or less, a much stronger bound could be obtained.

Theorem 2 (Buzaglo, Pinchasi and Rote [2]). *Every pseudodisc topological hypergraph on n vertices has $O(k^2 n)$ hyperedges of size k or less.*

However, the methods used to prove Theorem 2 do not seem to give any non-trivial bound on the number of hyperedges of size *exactly* k. Tight bounds are only known is case $k = 2$. Indeed, every 2-uniform S-capturing hypergraph is a planar graph [7], called the *convex distance function Delaunay graph*, and thus has at most $3n - 6$ edges.

1.2 Our results.

In this paper, we consider the case when every hyperedge has **exactly k points**. In particular, we consider k-uniform S-capturing hypergraphs, for convex and compact sets S.

One can show that these hypergraphs are pseudodisc topological hypergraphs. Indeed, the family of all homothetic copies of a fixed convex set S is surely the most important example of a family of pseudodiscs. For a finite point set X of the plane, a subset Y of X can be *separated* with a straight line if there exists a line ℓ such that one halfplane defined by ℓ contains all points in Y and the other contains all points in $X - Y$. For a positive integer i, let a_i denote the number of i-subsets of X that can be separated with a straight line.

Theorem 3. *Let S be a convex compact set, X be a finite point set and k be a positive integer. Any k-uniform S-capturing hypergraph on vertex set X has at most $(2k-1)|X| - k^2 + 1 - \sum_{i=1}^{k-1} a_i$ hyperedges. Moreover, equality holds whenever S is nice and X is in general position with respect to S.*

Here a set S is called *nice* if S has "no corners" and "no straight segments on its boundary". We define such nice shapes formally later. Moreover, X is in *general position with respect to S* if no three points of X are collinear and no four points of X lie on the boundary of any S-range.

Note that for $k = 2$ the bound in Theorem 3 amounts for at most $3|X| - 3 - t$ edges, where $t = a_1$ is the number of corners of the convex hull of X. We obtain the following refinement of Theorem 2.

Corollary 4. *Let S be a convex compact set and k, n be positive integers. Any S-capturing hypergraph on n vertices has at most $k^2 n + O(k^3)$ hyperedges of size k or less.*

The paper is organized as follows. Section 2 provides general definitions. Here we also show how to reduce the general case of an arbitrary capturing hypergraph to one with a nice shape S and a point set X in general position with respect to S. Section 3 introduces different types of ranges. The number of ranges of Type I is determined exactly in Section 3.1. Section 3.2 gives an identity involving the number of ranges of both types in X. Finally, Theorem 3 is proven in Section 4.

2 Nice shapes, general position and next range

In this section we introduce nice shapes, the concepts of the next range and state their basic properties. For the ease of reading, the proofs of some results in this section are provided in the appendix because they are quite straightforward but also technical. We denote the boundary of a set S by ∂S. We denote the line through distinct points p and q by $\overline{pq}$. A halfplane defined by a line ℓ is a connected component of the plane after the removal of ℓ. In particular, such halfplanes are open sets. A closed halfplane is the closure of a halfplane, i.e., the union of the halfplane and its defining line. Typically we denote the two halfplanes defined by a line by L and R (standing for "left" and "right").

For a set X of n points in the plane and any $i \in [n]$ an *i-set* of X is a subset Y of X on i elements that can be separated with a straight line. In other words, Y is an i-set if it is captured by a closed halfplane. The number of i-sets of X is denoted by a_i. Note that some (but in general not all) i-sets can be captured by a halfplane that has two points of X on its boundary. Such halfplanes are called *representative halfplanes* and we denote the set of all representative halfplanes of i-sets of X by $\mathcal{A}_i$. Note that (even if $i \geq 2$) the number of representative halfplanes for a fixed i-set might be anything, including 0.

Lemma 5. *For any set X of n points in the plane, no three on a line, and any $i \in \{1, \dots, n-1\}$ we have $a_i = |\mathcal{A}_{i+1}|$.*

Proof. Let X be a finite point set with no three points in X on a line, $|X| = n$ and $i \in \{1, \dots, n-1\}$. We shall give a bijection between the set X_i of i-sets of X and $\mathcal{A}_{i+1}$.

Let $Y \in X_i$ be an i-set. Assume (by rotating the plane if needed) that Y is separated from $X - Y$ by a vertical line, such that Y is contained in the corresponding right halfplane. Consider all closed halfplanes that contain all points in Y and whose interior does not contain any point in $X - Y$. Among all lines defining such halfplane let ℓ be one with smallest slope. Then $\ell \cap X$ contains a point $p \in X - Y$ and a point $q \in Y$ and going from p to q along ℓ we have the corresponding halfplane that contains X on the right. In particular this right halfplane of ℓ is in $\mathcal{A}_{i+1}$.

On the other hand, for any closed halfplane $H \in \mathcal{A}_{i+1}$ consider the line ℓ defining H and the two points $p, q \in X \cap \ell$ so that going from p to q along ℓ we have H on the right. Then rotating ℓ slightly counterclockwise around any point on ℓ between p and q shows that $(H \cap X) - p$ is an i-set of X.

The above bijection shows that $a_i = |X_i| = |\mathcal{A}_{i+1}|$, as desired. □

2.1 Nice shapes and general position of a point set.

A convex compact set S is called a **nice shape** if

(i) for each point in ∂S there is exactly one line that intersects S only in this point and

(ii) the boundary of S contains no non-trivial straight line segment.

For example, a disc is a nice shape, but a rectangle is not. A nice shape has no "corners" and we depict nice shapes as discs in most of the illustrations.

Lemma 6. *If S is a nice shape, S_1 and S_2 are distinct S-ranges, then each of the following holds.*

(i) $\partial S_1 \cap \partial S_2$ is a set of at most two points.

(ii) If $\partial S_1 \cap \partial S_2 = \{p, q\}$ and L and R are the two open halfplanes defined by $\overline{pq}$, then

- *$S_1 \cap L \subset S_2 \cap L$ and $S_1 \cap R \supset S_2 \cap R$ or*
- *$S_1 \cap L \supset S_2 \cap L$ and $S_1 \cap R \subset S_2 \cap R$.*

(iii) Any three non-collinear points lie on the boundary of a unique S-range.

(iv) For a subset of points $X \subset \mathbb{R}^2$ and any $Y \subset X$, $|Y| \geq 2$, that is captured by some S-range there exists at least one S-range S' with $Y = X \cap S'$ and $|\partial S' \cap X| \geq 2$.

The proof of Lemma 6 is provided in the Appendix. We remark that only the last item of Lemma 6 remains true if S is convex compact but not nice. For example, if S is an axis-aligned square, then no three points with strictly monotone x- and y-coordinates lie on the boundary of any S-range, whereas three points, two of which have the same x- or y-coordinate lie on the boundary of infinitely many S-ranges.

For sets $X, S \subset \mathbb{R}^2$ we say that X **is in general position with respect to** S if

(i) no two points of X are on a vertical line,

(ii) no three points of X are collinear,

(iii) no four points of X lie on the boundary of any S-range.

Lemma 7. *For any point set X, positive integer k and a convex compact set S, there is a nice shape S' and a point set X' in general position with respect to S', such that $|X'| = |X|$ and the number of edges in $\mathcal{H}(X', S', k)$ is at least as large as the number of edges in $\mathcal{H}(X, S, k)$.*

To prove Lemma 7, we show that one can move the points of X slightly and modify S slightly to obtain the desired property. See the Appendix for a detailed account of the argument. From now on we will always assume that S is a nice shape and X is a finite point set in general position with respect to S.

2.2 Next Range.

For two distinct points p, q in the plane, we define $\mathcal{S}(p, q)$ to be the set of all S-ranges S_1 with $p, q \in \partial S_1$. The symmetric difference of two sets A and B is given by $A \triangle B = (A \setminus B) \cup (B \setminus A)$.

Lemma 8. *Let p, q be two points such that no four points in $X \cup \{p, q\}$ lie on the boundary of an S-range. Let L be a halfplane defined by $\overline{pq}$. Then the following holds.*

(i) The S-ranges in $\mathcal{S}(p, q)$ are linearly ordered, denoted by $\prec_{pq}$, by inclusion of their intersection with L:

$$S_1 \prec_{pq} S_2 \quad \Leftrightarrow \quad S_1 \cap L \subset S_2 \cap L \qquad \text{for all } S_1, S_2 \in \mathcal{S}(p, q)$$

(ii) For each $S_1 \in \mathcal{S}(p, q)$ there exists a $\prec_{pq}$-minimal $S_2 \in \mathcal{S}(p, q)$ with

$$(\partial S_2 \setminus \partial S_1) \cap X \neq \emptyset \qquad \text{and} \qquad S_1 \prec_{pq} S_2$$

if and only if $S_1 \triangle L$ contains a point from X in its interior.

The proof of Lemma 8 is provided in the Appendix.

Whenever no four points in $X \cup \{p, q\}$ lie on the boundary of an S-range, L is a halfplane defined by $\overline{pq}$ and $S_1 \in \mathcal{S}(p, q)$, we define $\text{next}_L(S_1)$, called the **next range** of S_1 in L, as follows.

- If the interior of $S_1 \triangle L$ contains a point from X, then $\text{next}_L(S_1) = S_2$ for the S-range $S_2 \in \mathcal{S}(p,q)$ in Lemma 8 (ii).
- If the interior of $S_1 \triangle L$ contains no point from X, then $\text{next}_L(S_1) = L$.

Informally, we can imagine continuously transforming S_1 into a new S-range containing p and q on its boundary and containing $S_1 \cap L$ and all points of $(S_1 \setminus \partial S_1) \cap X$. As soon as this new S-range contains a point from $X \setminus \{p,q\}$ on its boundary, we choose it as the next range of S_1 in L. Note that if $S_2 = \text{next}_L(S_1)$ and ∂S_1 contains a point from $X \setminus \{p,q\}$, then $\text{next}_R(S_2) = S_1$, where R denotes the other halfplane defined by $\overline{pq}$.

As no four points in $X \cup \{p,q\}$ lie on the boundary of an S-range, we have $|\partial S_1 \cap (X \setminus \{p,q\})| \leq 1$ for each $S_1 \in \mathcal{S}(p,q)$. This implies that if S_1 captures k elements of X then $\text{next}_L(S_1)$ captures $k-1$, k or $k+1$ elements of X. Indeed, if $\text{next}_L(S_1) = L$, then $|X \cap L| = k$ and if $\text{next}_L(S_1) \neq L$, then the following holds.

$$\text{If } R \cap \partial S_1 \cap X \neq \emptyset, \text{ then } \text{next}_L(S_1) \text{ captures } k \text{ or } k-1 \text{ points.} \tag{1}$$

$$\text{If } R \cap \partial S_1 \cap X = \emptyset, \text{ then } \text{next}_L(S_1) \text{ captures } k \text{ or } k+1 \text{ points.} \tag{2}$$

See Figure 1 for the three possible case scenarios.

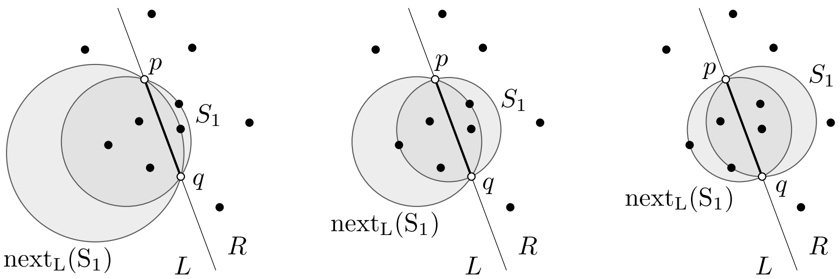

Figure 1: Three cases of an S-range S_1 with two boundary points p and q, and the next S-range of S_1 in a halfplane L defined by $\overline{pq}$. Note that $|X \cap \text{next}_L(S_1)| - |X \cap S_1|$ is -1 on the left, 0 in the middle, and 1 on the right.

3 Representative S-ranges and Types I and II

Let X be a set in a general position and S be a nice shape. Let Y be any hyperedge in $\mathcal{H}(X,S,k)$. An S-range S' is a *representative S-range for Y* if $Y = X \cap S'$ and among all such S-ranges S' has the maximum number of points from Y on its boundary. From Lemma 6 (iv) it follows that each hyperedge has a representative range and if S' is a representative S-range for Y, then S' has two or three points of X on its boundary.

We say that S' is of **Type I** if $|\partial S' \cap X| = 3$ and of **Type II** if $|\partial S' \cap X| = 2$.

We say that Y is of *Type I* if it has a representative range of Type I, otherwise it is of *Type II*. Note that in total we have at most $\binom{k}{3}$ many Type I ranges representing a Type I hyperedge since by Lemma 6 (iii) any three points of X are on the boundary of only one S-range. On the other hand, every Type II hyperedge has infinitely many representative ranges, see Figure 2. The **representative set** of Y contains all representative ranges for a Type I set Y and it contains one arbitrarily chosen representative range for a Type II set Y. We denote a representative set of Y by $\mathcal{R}(Y)$.

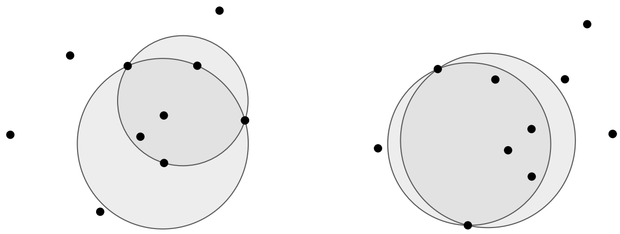

(a) (b)

Figure 2: (a) A Type I hyperedge and with only two representative ranges. (b) A Type II hyperedge with two of its infinitely many representative ranges.

We define

$$\begin{aligned}\mathcal{E}_1^k &= \{Y \subseteq X \mid Y \text{ is a Type I hyperedge}\},\\ \mathcal{R}_1^k &= \{S \in \mathcal{R}(Y) \mid Y \in \mathcal{E}_1^k\},\\ \mathcal{E}_2^k &= \{Y \subseteq X \mid Y \text{ is a Type II hyperedge}\} \text{ and}\\ \mathcal{R}_2^k &= \{S \in \mathcal{R}(Y) \mid Y \in \mathcal{E}_2^k\}.\end{aligned}$$

Lemma 9. *All Type II representative S-ranges for the same hyperedge have the same pair of X in their boundary.*

Proof. Consider a representative range S_1 for a Type II hyperedge Y with $\{p_1, q_1\} = Y \cap \partial S_1$. Assume for the sake of contraction that there is another representative range, S_2 for Y with $\{p_2, q_2\} = Y \cap S_2$ and $\{p_2, q_2\} \neq \{p_1, q_1\}$. We have that $p_2, q_2 \in S_1$, $p_1, q_1 \in S_2$. Assume, without loss of generality that $q_2 \notin \{p_1, q_1\}$ and $q_1 \notin \{p_2, q_2\}$. Then $q_2 \in S_1 - \partial S_1$ and $q_1 \in S_2 - \partial S_2$.

We need to distinguish the following cases: segments p_2q_2 and p_1q_1 cross properly, p_2q_2 and p_1q_1 share a vertex, i.e., $p_2 = p_1$, and finally p_1q_1 is to the left of p_2q_2. The first case does not occur by the argument presented in the previous item. In the other two cases, let L and R be the halfplanes defined by $\overline{p_1q_1}$ *not* containing q_2 and containing q_2, respectively. Consider the S-range $S_3 = \text{next}_L(S_1)$, which exists by Lemma 8 as $S_1 \Delta L$ contains q_2. We see that R contains at least one point in $\partial S_3 \cap \partial S_2$, see Figure 3. Moreover, S_3 must contain a point $z \in X \setminus Y$ because it is a next range and Y is not of Type I. It follows that $z \in L \setminus S_2$, which implies that the closure of L contains two points in $\partial S_3 \cap \partial S_2$, a contradiction to Lemma 6 (i). □

From Lemma 9 and the definitions above we immediately get the following.

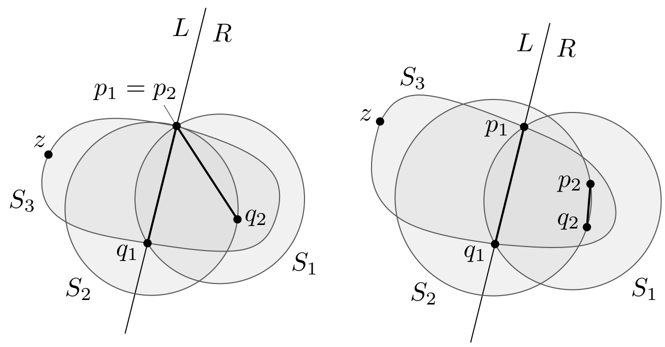

Figure 3: The second and third case in the proof of Lemma 9.

$$|\mathcal{E}_1^k| = |\mathcal{R}_1^k| - \sum_{Y \in \mathcal{E}_1^k} (|\mathcal{R}(Y)| - 1) \tag{3}$$

$$|\mathcal{E}_2^k| = |\mathcal{R}_2^k| \tag{4}$$

For a Type I hyperedge Y, let $Y' \subseteq Y$ be the subset of vertices that are on the boundary of at least one representative range for Y. We define the graph $G(Y) = (Y', E_Y)$ with vertex set Y' and edge set $E_Y = \{\{p, q\} \mid p, q \in \partial S',\ S' \in \mathcal{R}(Y)\}$. Then $G(Y)$ is the union of triangles, one for each representative range of Y. We call an edge of $G(Y)$ *inner edge* if it is contained in at least two triangles.

Lemma 10. *For every Type I hyperedge Y the graph $G(Y)$ is a maximally outerplanar graph. In particular, $G(Y)$ has exactly $|\mathcal{R}(Y)| - 1$ inner edges.*

Proof. We shall show that $G = G(Y)$ is maximally outerplanar by finding a planar embedding of G in which every vertex lies on the outer face, every inner edge lies in two triangles and every outer edge lies in one triangle. To this end draw every vertex of G at the position of its corresponding point in Y' and every edge as a straight-line segment.

First observe that every point in Y' lies on the convex hull of Y'. Hence, if G is drawn crossing-free, then every vertex of G lies on the outer face of G.

Second, assume for the sake of contraction that two edges $x_1y_1, x_2y_2 \in E(G)$ cross. Without loss of generality these four vertices appear on the convex hull of Y' in the clockwise order x_1, x_2, y_1, y_2. But then the S-ranges S_1 and S_2 with $Y \subset S_i$ and $\{x_i, y_i\} = \partial S_i \cap Y$ for $i = 1, 2$ have at least four intersections on their boundaries. See Figure 4(a) for an illustration. This is a contradiction to Lemma 6 (i), i.e., that $|\partial S_1 \cap \partial S_2| \leq 2$.

Finally, for any edge xy in G we shall show that x and y are consecutive points on the convex hull of Y', or xy is contained in two triangles. Note that this proves that G is maximally outerplanar.

As xy is an edge in G, there is a representative range $S_1 \in \mathcal{R}(Y)$ with $x, y \in \partial S_1$. Let L, R be the two open halfplanes defined by $\overline{xy}$. Assume without loss of generality that the third point in $\partial S_1 \cap X$ lies in L. We distinguish two cases.

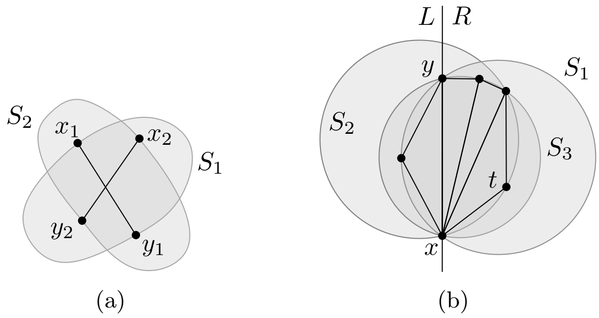

Figure 4: (a) If two edges x_1y_1 and x_2y_2 in $G(Y)$ cross, then the corresponding S-ranges have at least four intersections on their boundaries. (b) Illustration of the proof of Lemma 10.

If $R \cap Y' = \emptyset$, then x and y are consecutive on the convex hull of Y' and the edge xy is an outer edge of G. Otherwise, there exists some point $t \in R \cap Y'$ and we shall show that the edge xy lies in two triangles. Let $S_2 \in \mathcal{R}(Y)$ be a representative range with $t \in \partial S_2$, which exists as $t \in Y'$. Also consider $S_3 = \text{next}_L(S_1)$. By (2) we have $Y \subset S_3$ and $|X \cap S_3| \in \{k, k+1\}$. If $S_2 = S_3$, then the edge xy lies in two triangles, one for S_1 and one for $S_2 = S_3$. Otherwise the situation is illustrated in Figure 4(b). We have $x, y \in S_2$ and $t \in S_3 - \partial S_3$. It follows that ∂S_2 and ∂S_3 intersect (at least) twice in the closure of R. By Lemma 6 (i),(ii), we have $L \cap S_2 \subset L \cap S_3$. Thus, $S_3 \subset S_1 \cup S_2$, which implies that $S_3 \cap X = Y$, i.e., $S_3 \in \mathcal{R}(Y)$ and the edge xy lies in two triangles of G, as desired.

To summarize, G is drawn crossing-free with all vertices on the outer face, and every edge of G lies on the convex hull of Y' or in two triangles. This implies that G is maximally outerplanar. □

3.1 The number of Type I ranges.

Recall that for $i = 2, \ldots, |X|$ we denote by $\mathcal{A}_i$ the set of representative halfplanes for i-sets of X. In the next two proofs we treat representative halfplanes similarly to S-ranges. Indeed, one can think of a halfplane as a homothet of S with infinitely large dilation and at the same time infinitely large translation (however, formally this is incorrect!). In the light of Lemma 8, a halfplane defined by $\overline{pq}$ arises as a kind of limit object of a sequences of S-ranges in $\mathcal{S}(p, q)$. Accordingly, we defined $\text{next}_L(S_1) = L$ if $(S_1 \triangle L) \cap X = \emptyset$.

Proposition 11. *For $k \geq 3$ we have*

$$|\mathcal{R}_1^k| = 2(k-2)|X| - \sum_{i=2}^{k-1} |\mathcal{A}_i| - (k-1)(k-2).$$

Proof. For a point $p \in X$ and a set S' that is either a Type I S-range or a representative halfplane, we say that p is the *second point* of S' if

- $\partial S' \cap X = \{p, q, r\}$ and the x-coordinate of p lies strictly between the x-coordinates of q and r, or

- $\partial S' \cap X = \{p, q\}$ and S' is on the right when going from p to q along $\overline{pq}$.

Clearly, every representative halfplane has a unique second point. Note that also every Type I S-range has a unique second point, because no two points in X have the same x-coordinate. Moreover, if p is the second point of S' and ℓ denotes the vertical line through p, then $S' \cap \ell$ is a vertical segment (if S' is an S-range) or ray (if S' is a halfplane) with an endpoint p. We say that p is a *lower*, respectively *upper*, second point of S' if $S' \cap \ell$ has p as its lower, respectively upper endpoint.

Now, fix $p \in X$ and the vertical line ℓ through p. Let L and R denote the left and right open halfplanes defined by ℓ, respectively. We want to show that roughly $k-2$ S-ranges in $\mathcal{R}_1^k$ have the lower second point p. We say $S' \in \mathcal{R}_1^k \cup \bigcup_{i \geq 2} \mathcal{A}_i$ has *property* (a, b) if

- S' has the lower second point p,
- $L \cap S'$ contains exactly a points from X, one on $\partial S'$ if $a \geq 1$, and
- $R \cap S'$ contains exactly b points from X.

Claim. Let $m \geq 0$ be the number of points in X whose x-coordinate is smaller than the x-coordinate of p. Then for each $a = 1, \dots, \min(m, k-2)$ there exists $S_{a,p} \in \mathcal{R}_1^k \cup \bigcup_{i \geq 2} \mathcal{A}_i$ with property (a, b), so that

- either $a + b \leq k - 2$ and $S_{a,p}$ is a halfplane,
- or $a + b = k - 1$ and $S_{a,p}$ is a Type I range.

Proof of Claim. Let $a \in \{1, \dots, \min(m, k-2)\}$ be fixed. We shall first construct S-ranges with properties $(0,0), (1,0), \dots, (a,0)$, respectively, and then S-ranges with properties $(a,1), (a,2), \dots, (a,b)$, respectively. We start with any S-range S_0 with property $(0,0)$. Let q denote the upper endpoint of $\ell \cap S_0$. We choose S_0 so that no four points of $X \cup q$ lie on the boundary of any S-range. Then we define for $i = 1, \dots, a$ S_i to be the next S-range of S_{i-1} in L, i.e., $S_i = \mathrm{next}_L(S_{i-1})$. By (1),(2), S_i has property $(i, 0)$. In particular, S_a has property $(a, 0)$. See Figure 5(a) for an illustration.

Next, we shall construct a sequence $T_0, T_1, \dots T_t$ of S-ranges and possibly one halfplane, and a sequence $r_0, r_1, \dots r_t$ of elements of X with the following properties.

A) $r_i \in \partial T_i \cap L$ $(i = 1, \dots, t)$.

B) $\ell \cap T_i$ is a segment strictly shorter than $\ell \cap T_{i+1}$ $(i = 1, \dots, t-1)$.

C) $|T_i \cap R \cap X| = x_i$ $(i = 1, \dots, t)$ with $x_i \leq x_{i+1}$ $(i = 1, \dots, t-1)$.

D) $|(T_i \setminus \gamma_i) \cap L \cap X| = a$, where γ_i denotes the component of $\partial T_i \setminus \{p, r_i\}$ that is completely contained in L $(i = 1, \dots, t)$.

E) When $x_i < x_{i+1}$, then T_{i+1} has property (a, x_{i+1}), i.e., T_{i+1} is a Type I range with the second point p $(i = 1, \dots, t-1)$.

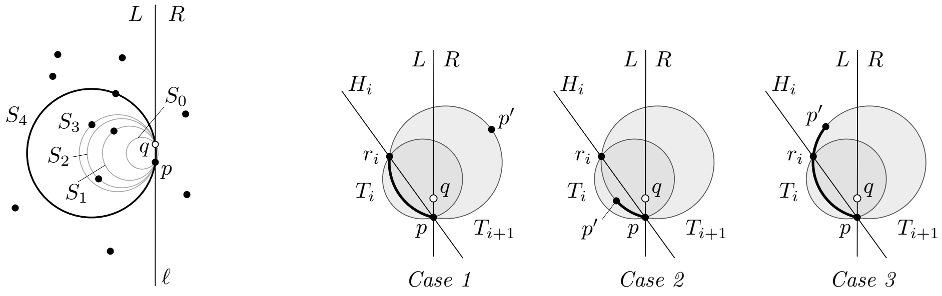

(a) (b)

Figure 5: (a) The S-ranges $S_0, \dots, S_4$ have property $(0,0), \dots, (4,0)$, respectively. (b) The three cases in the construction of T_{i+1} based on T_i and r_i. The component γ_{i+1} of $\partial T_{i+1} \setminus \{p, r_{i+1}\}$ that is completely contained in L is drawn bold.

F) T_t is a halfplane.

We construct the sequence $T_0, T_1, \dots, T_t$ as follows. Let $T_0 = S_a$, $r_0 \in \partial T_0 \cap (L \cap X)$. Assume that $T_0, \dots, T_i$ and $r_0, \dots, r_i$ have been constructed and T_i is not a halfplane. Let H_i denote the right halfplane defined by $\overline{pr_i}$, i.e., $q \in H_i$. We define $T_{i+1} = \text{next}_{H_i}(T_i)$. Then by Lemma 8 we have $T_{i+1} \cap H_i \supset T_i \cap H_i$ and hence the segment $\ell \cap T_i$ is strictly shorter than $\ell \cap T_{i+1}$. Moreover, Lemma 8 implies that $|(T_{i+1} \setminus \gamma_{i+1}) \cap L \cap X| = a$ and $x_{i+1} = |T_{i+1} \cap R \cap X| \in \{x_i, x_{i+1}\}$.

Finally, we shall define the point r_{i+1}. If $T_{i+1} = H_i$ is a halfplane, we set $r_{i+1} = r_i$, $t = i + 1$ and the sequence is complete. Otherwise, T_{i+1} is a bounded S-range, and we consider the unique point p' in $(\partial T_{i+1} \setminus \partial T_i) \cap X$. We distinguish three cases.

Case 1: $p' \in R$. We have that $|T_{i+1} \cap L \cap X| = a$, $|T_{i+1} \cap R \cap X| = x_i + 1$ and $\partial T_{i+1} \cap X = \{r_i, p, p'\}$. So T_{i+1} has property (a, x_{i+1}) with $x_{i+1} = x_i + 1$. Set $r_{i+1} = r_i$, which implies $\gamma_{i+1} \cap X = \emptyset$.

Case 2: $p' \in L \setminus H_i$. Then $|T_{i+1} \cap L \cap X| = a$ and $|T_{i+1} \cap R \cap X| = x_i$, just like T_i. In this case we set $r_{i+1} = p'$, which gives again $\gamma_{i+1} \cap X = \emptyset$.

Case 3: $p' \in L \cap H_i$. Then $|T_{i+1} \cap L \cap X| = a + 1$ and $|T_{i+1} \cap R \cap X| = x_i$. We set $r_{i+1} = p'$, which implies $r_i \in \gamma_{i+1}$ and hence $|(T_{i+1} \setminus \gamma_{i+1}) \cap L \cap X| = a$.

We refer to Figure 5(b) for an illustration. We see that if T_{i+1} is not a halfplane, we either have $r_{i+1} \neq r_i$ or $x_{i+1} > x_i$. Since there are finitely many possibilities for r_i and x_i and no pair $\{r_i, x_i\}$ occurs twice, at some point T_{i+1} is a halfplane.

Note that T_t is a halfplane with property (a, x_t). Hence, if $x_t < k - 1 - a$, then T_t is the desired S-range $S_{a,p} \in \bigcup_{i \geq 2} \mathcal{A}_i$. Otherwise, a subsequence of $T_0, T_1, \dots T_t$ consists of Type I ranges with properties $(a, 0), (a, 1), \dots, (a, b)$ with $a + b = k - 1$ and the last element of this subsequence is the desired S-range $S_{a,p} \in \mathcal{R}_1^k$. This completes the proof of the claim. $\triangle$

Note that for every $S' \in \mathcal{R}_1^k \cup \bigcup_{i\geq 2} \mathcal{A}_i$ we have $S' = S_{a,p}$ for at most one pair (a,p) of a number $a \in \{1,\ldots,k-2\}$ and a point $p \in X$. The above claim states that for given (a,p) we find $S_{a,p}$, unless fewer than a points in X have smaller x-coordinate than p. This rules out $\binom{k-1}{2}$ of the $(k-2)|X|$ pairs (a,p) and we conclude that

$$(k-2)|X| - \binom{k-1}{2} \leq |\mathcal{R}_1^{k,\downarrow}| + \sum_{i=2}^{k-1} |\mathcal{A}_i^{\downarrow}|, \tag{5}$$

where $\mathcal{R}_1^{k,\downarrow} \subseteq \mathcal{R}_1^k$, respectively $\mathcal{A}_i^{\downarrow} \subseteq \mathcal{A}_i$ $(i = 2,\ldots,k-2)$, are the subsets of Type I ranges, respectively representative halfplanes, with a *lower* second point.

By symmetry, we obtain an analogous inequality for the Type I ranges $\mathcal{R}_1^{k,\uparrow}$ and representative halfplanes $\mathcal{A}_i^{\uparrow}$ with an *upper* second point. Together with $\mathcal{R}_1^{k,\downarrow} \cap \mathcal{R}_1^{k,\uparrow} = \emptyset$ and $\mathcal{A}_i^{\downarrow} \cap \mathcal{A}_i^{\uparrow} = \emptyset$ for $i \geq 2$ this shows that

$$2(k-2)|X| - (k-1)(k-2) \leq |\mathcal{R}_1^{k,\downarrow}| + |\mathcal{R}_1^{k,\uparrow}| + \sum_{i=2}^{k-1}(|\mathcal{A}_i^{\downarrow}| + |\mathcal{A}_i^{\uparrow}|) \leq |\mathcal{R}_1^k| + \sum_{i=2}^{k-1} |\mathcal{A}_i|. \tag{6}$$

In order to finish the proof of Proposition 11 it remains to prove that (5) (and hence also (6)) holds with equality. For this, we show that for every point $p \in X$ and every $a \in \{1,\ldots,k-2\}$ there is at most one S-range in $\mathcal{R}_1^k$ with property (a,b) for $a+b = k-1$ and at most one halfplane in $\bigcup_{i\geq 2} \mathcal{A}_i$ with property (a,b) for $a+b \leq k-2$. If $S_1, S_2 \in \mathcal{R}_1^k$ are two distinct S-ranges with the same lower second point p, then by Lemma 6 (ii) $S_1 \cap L \subset S_2 \cap L$ or $S_2 \cap L \subset S_1 \cap L$. Observe that, if $S_1, S_2 \in \bigcup_{i\geq 2} \mathcal{A}_i$ are distinct closed halfplanes with the same lower second point p, then we also have $S_1 \cap L \subset S_2 \cap L$ or $S_2 \cap L \subset S_1 \cap L$. So in either case we may assume that $S_1 \cap L \subset S_2 \cap L$. Now, if S_2 has property (a,b), then $\partial S_2 \cap L \cap X \neq \emptyset$ and hence $|S_2 \cap L \cap X| > |S_1 \cap L \cap X|$. Thus S_1 and S_2 can not both have property (a,b) for the same a.

We conclude that inequality (5) holds with equality. Thus (6) also holds with equality, which is the statement of Proposition 11. □

3.2 Relation between the number of Type I and Type II ranges.

Recall that for a fixed Type I hyperedge Y in $\mathcal{H}(X,S,k)$ we denote by $\mathcal{R}(Y)$ the set of representative ranges for Y.

Proposition 12. *For $k \geq 3$ we have*

$$3|\mathcal{R}_1^k| + 2|\mathcal{R}_2^k| = 3|\mathcal{R}_1^{k+1}| + |\mathcal{A}_k| + 2\sum_{Y\in\mathcal{E}_1^k}(|\mathcal{R}(Y)| - 1).$$

Proof. Consider the set P of all ordered pairs (S_1,S_2) of an S-range S_1 and an S-range or representative halfplane S_2, such that

(A) $S_1 \in \mathcal{R}_1^k \cup \mathcal{R}_2^k$.

(B) $\partial S_1 \cap \partial S_2$ is a pair p, q of points in X.

(C) $S_2 = \text{next}_L(S_1)$, where L is one of the halfplanes defined by $\overline{pq}$.

(D) $X \cap S_1 \subseteq X \cap S_2$.

For a pair $(S_1, S_2) \in P$ we say that S_2 *is an image of* S_1 and S_1 *is a preimage of* S_2. Note that, if S_2 is an image of S_1, then S_1 contains k points from X and thus by (1),(2) S_2 contains either k or $k+1$ points from X. In the former case, S_1 and S_2 are either distinct representative S-ranges for the same Type I hyperedge in $\mathcal{H}(X, S, k)$ or S_2 is a halfplane, see Figure 6(a), while in the latter case S_2 is a Type I range in $\mathcal{H}(X, S, k+1)$, see Figure 6(a) and (b).

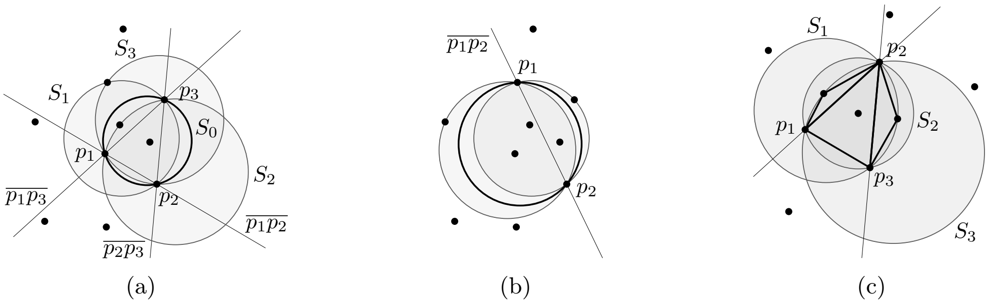

Figure 6: (a) The three images S_1, S_2, S_3 of a Type I range S_0 (in bold). Note that S_1 and S_3 correspond to the same hyperedge in $\mathcal{H}(X, S, k+1)$, and that S_0 and S_2 correspond to the same hyperedge in $\mathcal{H}(X, S, k)$. (b) The two images of a Type II range (in bold); one being also an image of the Type I range in (a). (c) Three representative ranges for the same hyperedge Y and the outerplanar graph $G(Y)$ (in bold). Here $(S_1, S_2) \in P_4$ and $(S_2, S_1) \in P_4$, both with respect to $\{p_2, p_3\}$, as well as $(S_2, S_3) \in P_4$ and $(S_3, S_2) \in P_4$, both with respect to $\{p_1, p_2\}$.

We partition P in two different ways; once with respect to the possibilities for preimages and once with respect to the possibilities for preimages. Firstly, $P = P_1 \dot\cup P_2$, where P_1 and P_2 contain all pairs (S_1, S_2) with $S_1 \in \mathcal{R}_1^k$ and $S_1 \in \mathcal{R}_2^k$, respectively. Secondly, $P = P_3 \dot\cup P_4 \dot\cup P_5$, where P_3, P_4 and P_5 contain all pairs (S_1, S_2) with $S_2 \in \mathcal{R}_1^{k+1}$, $S_2 \in \mathcal{A}_k$ and $S_2 \in \mathcal{R}_1^k$, respectively. We summarize:

$$P_i = \{(S_1, S_2) \in P \mid S_1 \in \mathcal{R}_i^k\}, \quad i = 1, 2,$$
$$P_3 = \{(S_1, S_2) \in P \mid S_2 \in \mathcal{R}_1^{k+1}\},$$
$$P_4 = \{(S_1, S_2) \in P \mid S_2 \in \mathcal{A}_k\},$$
$$P_5 = \bigcup_{Y \in \mathcal{E}_1^k} \{(S_1, S_2) \in P \mid S_1, S_2 \in \mathcal{R}(Y)\},$$
$$P = P_1 \dot\cup P_2 \quad \text{and} \quad P = P_3 \dot\cup P_4 \dot\cup P_5.$$

We shall show that, on one hand, $|P_1| = 3|\mathcal{R}_1^k|$ and $|P_2| = 2|\mathcal{R}_2^k|$, while on the other hand, $|P_3| = 3|\mathcal{R}_1^{k+1}|$, $|P_4| = |\mathcal{A}_k|$ and $|P_5| = 2\sum_{Y \in \mathcal{E}_1^k}(|\mathcal{R}(Y)| - 1)$. Together with $|P_1| + |P_2| = |P| = |P_3| + |P_4| + |P_5|$ this will conclude the proof.

To prove that $|P_1| = 3|\mathcal{R}_1^k|$, consider $S_1 \in \mathcal{R}_1^k$ and let $\partial S_1 \cap X = \{p_0, p_1, p_2\}$. For $i = 0, 1, 2$ let H_i be the halfplane defined by $\overline{p_{i-1}p_{i+1}}$ containing p_i, where indices are taken modulo 3. By (B),(C) and (D), every image S_2 of S_1 is the next S-range of S_1 in H_i for some $i \in \{0, 1, 2\}$, whose existence and uniqueness is given by Lemma 8 (ii). As every such next S-range contains $X \cap S_1$, S_1 has exactly three images.

To prove that $|P_2| = 2|\mathcal{R}_2^k|$, consider $S_1 \in \mathcal{R}_2^k$ and let $X \cap \partial S_1 = \{p_1, p_2\}$. If S_2 is an image of S_1, then by (B),(C) $S_2 = \text{next}_H(S_1)$, where H is a halfplane defined by $\overline{p_1p_2}$. As $\partial S_1 \cap X = \{p, q\}$, the next S-range in either halfplane contains $S_1 \cap X$, i.e., (D) is satisfied. Hence S_1 has two images, one for each halfplane.

To prove that $|P_3| = 3|\mathcal{R}_1^{k+1}|$, consider any $S_2 \in \mathcal{R}_1^{k+1}$ and let $X \cap \partial S_2 = \{p_0, p_1, p_2\}$. Let H_i and $\bar{H}_i$ be the halfplanes defined by $\overline{p_{i-1}p_{i+1}}$ containing p_i and *not* containing p_i, respectively, where indices are taken modulo 3 again. By (B), (C) and (D) every preimage S_1 of S_2 corresponds to a point $p_i \in \{p_0, p_1, p_2\}$ with $p_i \in S_2 \setminus S_1$, such that $S_2 = \text{next}_{H_i}(S_1)$. Indeed, if $S' = \text{next}_{\bar{H}_i}(S_2)$ captures k points from X, then $(S', S_2) \in P_1$. Whereas, if S' captures $k + 1$ points from X, then $Y = X \cap S' \cap S_2$ is a Type II hyperedge and for its representative range S'' we have $(S'', S_2) \in P_2$. Finally, S' can not capture $k+2$ points since $p_i \in S_2 \setminus S'$. Hence S_2 has exactly three preimages.

To prove that $|P_4| = |\mathcal{A}_k|$, we need to show that every halfplane $H \in \mathcal{A}_k$ is the image of exactly one S-range in $\mathcal{R}_1^k \cup \mathcal{R}_2^k$. In fact, if $\{p, q\} = X \cap \partial H$ and $Y = H \cap X$, then consider S-ranges defined by p, q and a third point from Y. By Lemma 6 these S-ranges are well-defined and by Lemma 8 they are linearly ordered by inclusion in H. For the S-range S' with $S' \cap H$ being inclusion-maximal we have $(S', H) \in P_4$, as desired.

Finally, we prove that $|P_5| = 2\sum_{Y \in \mathcal{E}_1^k}(|\mathcal{R}(Y)| - 1)$. Consider any hyperedge $Y \in \mathcal{E}_1^k$ and the graph $G(Y)$ defined above, whose edges are all pairs $\{p, q\} \subseteq Y$ such that $p, q \in \partial S'$ for a representative S-range $S' \in \mathcal{R}(Y) \subset \mathcal{R}_1^k$. By Lemma 10, connecting any two points in Y that are adjacent in $G(Y)$ with a straight line segment gives a maximally outerplanar drawing of $G(Y)$. Now if $(S_1, S_2) \in P_5$, then $\partial S_1 \cap \partial S_2 = \{p, q\}$ is an inner edge of $G(Y)$, see Figure 6(c). Moreover, exactly two S-ranges in $\mathcal{R}(Y)$ have p and q on their boundary, because S-ranges in $\mathcal{R}(Y)$ correspond to triangles in $G(Y)$ and $G(Y)$ is maximally outerplanar.

Thus, every pair $(S_1, S_2) \in P_5$ with $S_2 \in \mathcal{R}_1^k$ gives rise to an inner edge of $G(Y)$ and every inner edge $\{p, q\}$ of $G(Y)$ gives exactly two such ordered pairs in P_5. Because a maximally outerplanar graph with $|\mathcal{R}(Y)|$ triangles has $|\mathcal{R}(Y)| - 1$ inner edges, we have the desired equality.

Now we conclude that $|P| = |P_1| + |P_2| = 3|\mathcal{R}_1^k| + 2|\mathcal{R}_2^k|$, whereas $|P| = |P_3| + |P_4| + |P_5| = 3|R_1^{k+1}| + |\mathcal{A}_k| + 2\sum_{Y \in \mathcal{E}_1^k}(|\mathcal{R}(Y)| - 1)$. Together this gives the claimed equality. □

4 Proof of Theorem 3.

Proof of Theorem 3. For $k = 1$ and any X, the hypergraph $\mathcal{H}(X, S, k)$ clearly has $|X| = (2k-1)|X| - k^2 + 1 - \sum_{i=1}^{k-1} a_i$ hyperedges. For $k = 2$ and any X, $\mathcal{H}(X, S, k)$ is the well-known

Delaunay triangulation D_S with respect to the convex distance function defined by S. In particular, every inner face of D_S is a triangle and the outer face is the convex hull of X, i.e., has length a_1. Thus, by Euler's formula the number of hyperedges of $\mathcal{H}(X,S,2)$ (edges of D_S) is given by $3|X|-3-a_1 = (2k-1)|X|-k^2+1-\sum_{i=1}^{k-1} a_i$. So we may assume that $k \geq 3$. Moreover, by Lemma 7 we may assume that S is nice and X is in general position with respect to S.

By Proposition 12 we have

$$3|\mathcal{R}_1^k| + 2|\mathcal{R}_2^k| = 3|\mathcal{R}_1^{k+1}| + |\mathcal{A}_k| + 2\sum_{Y\in\mathcal{E}_1^k}(|\mathcal{R}(Y)|-1). \tag{7}$$

By Proposition 11 we have

$$|\mathcal{R}_1^k| = 2(k-2)|X| - \sum_{i=2}^{k-1}|\mathcal{A}_i| - (k-1)(k-2), \text{ and} \tag{8}$$

$$|\mathcal{R}_1^{k+1}| = 2(k-1)|X| - \sum_{i=2}^{k}|\mathcal{A}_i| - k(k-1). \tag{9}$$

Putting (3), (4), (7), (8) and (9) together, we conclude that

$$\begin{aligned}
2(|\mathcal{E}_1^k| + |\mathcal{E}_2^k|) &\overset{(3),(4)}{=} 2|\mathcal{R}_1^k| - 2\sum_{Y\in\mathcal{E}_1^k}(|\mathcal{R}(Y)|-1) + 2|\mathcal{R}_2^k| \\
&\overset{(7)}{=} 3|\mathcal{R}_1^{k+1}| + |\mathcal{A}_k| - |\mathcal{R}_1^k| \\
&\overset{(8),(9)}{=} 6(k-1)|X| - 3\sum_{i=2}^{k}|\mathcal{A}_i| - 3k(k-1) + |\mathcal{A}_k| \\
&\qquad - 2(k-2)|X| + \sum_{i=2}^{k-1}|\mathcal{A}_i| + (k-1)(k-2) \\
&= 2\left((2k-1)|X| - k^2 + 1 - \sum_{i=2}^{k}|\mathcal{A}_i|\right).
\end{aligned}$$

Thus we have with Lemma 5 that $|\mathcal{E}_1^k| + |\mathcal{E}_2^k| = (2k-1)|X| - k^2 + 1 - \sum_{i=2}^{k}|\mathcal{A}_i| = (2k-1)|X| - k^2 + 1 - \sum_{i=1}^{k-1} a_i$, as desired. □

5 Conclusions and remarks

In this paper we investigated k-uniform hypergraphs whose vertex set X is a set of points in the plane and whose hyperedges are exactly those k-subsets of X that can be captured by a homothetic copy of a fixed convex compact set S. These are so called k-uniform S-capturing hypergraphs. We have shown that every such hypergraph has at most $(2k-1)|X| - k^2 +$

$1 - \sum_{i=1}^{k-1} a_i$ hyperedges and that this is tight for every nice shape S. Here a_i denotes the number of i-subsets of X that can be separated with a straight line.

As an immediate corollary we obtain that if S is nice, then the total number of subsets of X captured by some homothet of S is given by the cake number $\binom{|X|}{3} + \binom{|X|}{2} + \binom{|X|}{1}$. Moreover, we obtain a bound on the number of hyperedges of size *at most* k: For every point set X, every convex set S and every $k \geq 1$ at most $k^2|X|$ different non-empty subsets of X of size at most k can be captured by a homothetic copy of S. This refines the recent $O(k^2|X|)$ bound by Buzaglo, Pinchasi and Rote [2].

Another interesting open problem concerns topological hypergraphs defined by a family of pseudodiscs. Here, the vertex set X is again a finite point set in the plane and every hyperedge is a subset of X surrounded by a closed Jordan curve such that any two such curves have at most two points in common. Buzaglo, Pinchasi and Rote [2] prove that every pseudodisc topological hypergraph has at most $O(k^2|X|)$ hyperedges of size *at most* k.

Question 1. What is the maximum number of hyperedges of size *exactly* k in a pseudodisc topological hypergraph?

As we learned after submission, Chevallier *et al.* [5,6] have an unpublished manuscript in which they prove that every inclusion-maximal k-uniform *convex* pseudodisc topological hypergraph with n vertices has exactly $(2k-1)|X| - k^2 + 1 - \sum_{i=1}^{k-1} a_i$ hyperedges. This independent result implies our Theorem 3. However, their proof is 40 pages long and involves higher-order Voronoi diagrams and higher-order centroid Delaunay triangulations. Our proof uses a completely different technique; it is short and completely self-contained, and hence is of independent interest.

Acknowledgments

We would like to thank Günter Rote and János Pach for helpful remarks and referring us first to [2] and later on to [6].

References

[1] Peter Brass, William Moser, and János Pach. *Research Problems in Discrete Geometry.* Springer, 2005.

[2] Sarit Buzaglo, Rom Pinchasi, and Günter Rote. Topological hypergraphs. In *Thirty Essays on Geometric Graph Theory*, pages 71–81. Springer, 2013.

[3] Jean Cardinal, Kolja B. Knauer, Piotr Micek, and Torsten Ueckerdt. Making triangles colorful. *Journal of Computational Geometry*, 4(1):240–246, 2013.

[4] Xiaomin Chen, János Pach, Mario Szegedy, and Gábor Tardos. Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles. *Random Struct. Algorithms*, 34(1):11–23, 2009.

[5] N. Chevallier, A. Fruchard, D. Schmitt, and J.-C. Spehner. Separation by Convex Pseudo-Circles. January 2015.

[6] Nicolas Chevallier, Augustin Fruchard, Dominique Schmitt, and Jean-Claude Spehner. Separation by convex pseudo-circles. In *Proc. Symposium on Computational Geometry*, (SOCG), pages 444–453, 2014.

[7] L. Paul Chew and Robert L. (Scot) Dyrsdale III. Voronoi diagrams based on convex distance functions. In *Proc. Symposium on Computational Geometry*, (SoCG), pages 235–244, 1985.

[8] Peter Frankl and János Pach. On disjointly representable sets. *Combinatorica*, 4(1):39–45, 1984.

[9] Balázs Keszegh and Dömötör Pálvölgyi. Octants are cover-decomposable into many coverings. *Computational Geometry*, 47(5):585–588, 2014.

[10] Seog-Jin Kim, Alexandr Kostochka, and Kittikorn Nakprasit. On the chromatic number of intersection graphs of convex sets in the plane. *Electron. J. Combin*, 11(1):R52, 2004. 12 pp.

[11] Seog-Jin Kim, Kittikorn Nakprasit, Michael J. Pelsmajer, and Jozef Skokan. Transversal numbers of translates of a convex body. *Discrete Mathematics*, 306(18):2166–2173, 2006.

[12] Jirí Matoušek. *Lectures on Discrete Geometry*. Springer, 2002.

[13] János Pach. Covering the plane with convex polygons. *Discrete & Computational Geometry*, 1:73–81, 1986.

[14] János Pach and Micha Sharir. *Combinatorial geometry and its algorithmic applications: The Alcalá lectures*. Number 152. American Mathematical Soc., 2009.

[15] János Pach and Gábor Tardos. Coloring axis-parallel rectangles. *J. Comb. Theory, Ser. A*, 117(6):776–782, 2010.

[16] János Pach and Gábor Tardos. Tight lower bounds for the size of epsilon-nets. In *Proc. Symposium on Computational Geometry*, (SoCG), pages 458–463, 2011.

[17] János Pach and Gerhard Woeginger. Some new bounds for ε-nets. In *Proc. Symposium on Computational Geometry*, (SoCG), pages 10–15, 1990.

[18] Dömötör Pálvölgyi. Indecomposable coverings with unit discs. *arXiv preprint arXiv:1310.6900*, 2013.

[19] Dömötör Pálvölgyi and Géza Tóth. Convex polygons are cover-decomposable. *Discrete & Computational Geometry*, 43(3):483–496, 2010.

[20] Norbert Sauer. On the density of families of sets. *J. Comb. Theory, Ser. A*, 13:145–147, 1972.

[21] Walter Schnyder. Embedding planar graphs on the grid. In *Proc. ACM-SIAM Symposium on Discrete Algorithms*, (SODA), pages 138–148, 1990.

[22] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific J. Math.*, 41:247–261, 1972.

[23] Shakhar Smorodinsky. Conflict-free coloring and its applications. In *Geometry–Intuitive, Discrete, and Convex*, pages 331–389. Springer, 2013.

[24] Gábor Tardos and Géza Tóth. Multiple coverings of the plane with triangles. *Discrete & Computational Geometry*, 38(2):443–450, 2007.

[25] Vladimir Vapnik and Alexey Červonenkis. The uniform convergence of frequencies of the appearance of events to their probabilities. *Teor. Verojatnost. i Primenen.*, 16:264–279, 1971.

[26] David R Wood. Hypergraph colouring and degeneracy. *arXiv preprint arXiv:1310.2972*, 2013.

Appendix

Here we give the omitted proofs of Lemma 6, Lemma 7, and Lemma 8.

For a set S, we call a line ℓ *touching* S if ℓ intersects S in exactly one point. So, S is nice if and only if for each point on its boundary there is exactly one touching line touching S at this point.

Lemma 6. *If S is a nice shape, S_1 and S_2 are distinct S-ranges, then each of the following holds.*

(i) $\partial S_1 \cap \partial S_2$ is a set of at most two points.

(ii) If $\partial S_1 \cap \partial S_2 = \{p, q\}$ and L and R are the two open halfplanes defined by $\overline{pq}$, then

- *$S_1 \cap L \subset S_2 \cap L$ and $S_1 \cap R \supset S_2 \cap R$ or*
- *$S_1 \cap L \supset S_2 \cap L$ and $S_1 \cap R \subset S_2 \cap R$.*

(iii) Any three non-collinear points lie on the boundary of a unique S-range.

(iv) For a subset of points $X \subset \mathbb{R}^2$ and any $Y \subset X$, $|Y| \geq 2$, that is captured by some S-range there exists at least one S-range S' with $Y = X \cap S'$ and $|\partial S' \cap X| \geq 2$.

Proof. (i) We show that for any two S-ranges S_1, S_2 such that $\{p, q, r\} \subseteq \partial S_1 \cap \partial S_2$, for distinct p, q, r, S_1 coincides with S_2. Assume not, and consider homothetic maps f_1, f_2 from S_1, S_2 to S, respectively. Let p_i, q_i, r_i be the images of p, q, r under f_i, $i = 1, 2$. Then we have that p_1, q_1, r_1 and p_2, q_2, r_2 form congruent triangles T_1, T_2 with vertices on ∂S. If these two triangles coincide, then $S_1 = S_2$. Otherwise consider two cases: a corner of one triangle is contained in the interior of the other triangle or not. If, without loss of generality, a corner of T_1 is in the interior of T_2, then by convexity, this corner can not be on ∂S. Otherwise, T_1 and T_2 are either disjoint or share a point on the corresponding side. In either case, one of the sides of T_1 is on the same line as the corresponding side of T_2, otherwise convexity of S is violated. Then, the boundary of S contains 3 collinear points, a contradiction to the fact that S is a nice shape.

(ii) We have, without loss of generality that $S_1 \cap L \subset S_2 \cap L$. If $S_1 \cap R \supset S_2 \cap R$, we are done. Otherwise, we have that $S_1 \cap R \subset S_2 \cap R$, and, in particular $S_1 \subseteq S_2$. Note that at p and q the S-ranges S_1 and S_2 have the same touching lines. Indeed, these lines are unique since S is nice. Consider maps f_1 and f_2 as before. We see that p, q are mapped into the same pair of points under both maps. Thus $S_1 = S_2$, a contradiction.

The remaining two items can be proven by considering two points p, q in the plane and the set $\mathcal{S}(p,q)$ of all S-ranges S' with $p, q \in \partial S'$. Indeed, given fixed p and q there is a bijection ϕ between the S-ranges in $\mathcal{S}(p,q)$ and the set $\mathscr{L}(p,q)$ of lines whose intersection with S is a non-trivial line segment parallel to the line $\overline{pq}$. We refer to Figure 7(a) for an illustration.

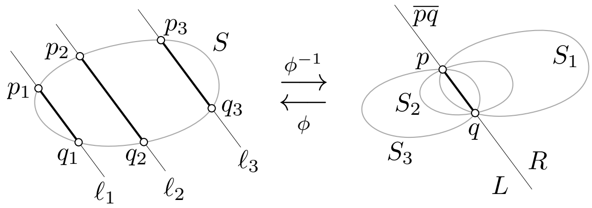

(a)

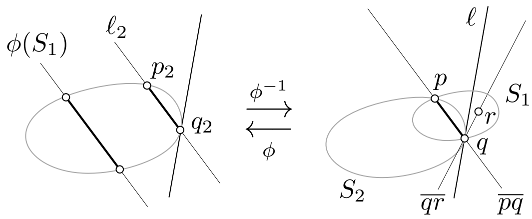

(b)

Figure 7: (a) A nice shape S, two points p, q in the plane, three lines $\ell_1, \ell_2, \ell_3 \in \mathscr{L}(p,q)$ and the corresponding S-ranges $S_1, S_2, S_3 \in \mathcal{S}(p,q)$. (b) Given an S-range $S_1 \in \mathcal{S}(p,q)$ with $r \in S_1$ we can find an S-range $S_2 \in \mathcal{S}(p,q)$ with $r \notin S_2$.

We verify that this bijection exists: For a line $\ell' \in \mathscr{L}(p,q)$ such that $\ell \cap \partial S = \{p_1, q_1\}$, let S' be an S-range obtained by contracting and translating S such that p_1 and q_1 are mapped into p and q, respectively. Let $\phi^{-1}(\ell) = S'$. Given an S'-range with $p, q \in \partial S$, consider a translation and contraction that maps S' to S. Let p_1 and q_1 be the images of p and q under this transform, then let $\phi(S') = \ell$, where ℓ is a line through p_1 and q_1. See again Figure 7(a) for an illustration.

(iii) Consider any three non-collinear points p, q, r in the plane. We shall first show that there is an S-range with p, q, r on its boundary. Let S_1 be an S-range of smallest area containing all three points. Clearly, $|\partial S_1 \cap \{p,q,r\}|$ is either 2 or 3. In the latter case we are done. So assume without loss of generality that $\partial S_1 \cap \{p,q,r\} = \{p,q\}$.

Now, we claim that there is another S-range S_2 that contains p and q but not r, with $p,q \in \partial S_2$. To find S_2, containing p,q on its boundary and not containing r, consider the triangle p,q,r and a line ℓ that goes through q having p and r on the different sides. Assume without loss of generality that q is the lowest point among p,q,r, and that $\overline{qr}$ and ℓ have positive slopes. Next, let q_2 be a point on ∂S whose touching line is parallel to ℓ and such that S is above this touching line. See Figure 7(b). Let ℓ_2 be the line parallel to $\overline{pq}$ and containing q_2 and let p_2 be the unique point in $\partial S \cap \ell_2$ different from q_2. Note that $|\ell_2 \cap \partial S| = 2$ follows from the fact that S is a nice shape and ℓ_2 and ℓ have different slopes. Let $S_2 = \phi^{-1}(\ell_2)$. We see that S_2 is above ℓ, but r is below ℓ. So, $r \notin S_2$.

Since ϕ is continuous and any two lines in $\mathscr{L}(p,q)$ can be continuously transformed into each other within $\mathscr{L}(p,q)$, we conclude that S_1 can be continuously transformed into S_2 within $\mathcal{S}(p,q)$. Thus, there is an S-range $S' \in \mathcal{S}(p,q)$ such that $r \in \partial S'$. This proves that any three non-collinear points p,q,r lie on the boundary of some S-range. The uniqueness of such a range follows from (i).

(iv) Let Y be a hyperedge in $\mathcal{H}(X,S,k)$ and S_1 be an S-range capturing Y. Contract S_1 until the resulting range S_2 contains at least one point, p of Y, on its boundary. If $|\partial S_2 \cap Y| \geq 2$, we are done. Otherwise, consider a small S-range S_3 containing p on its boundary and not containing any other points of X. Let q be the second point in $\partial S_2 \cap \partial S_3$. Similarly to the previous argumentation, S_2 can be continuously transformed into S_3 within $\mathcal{S}(p,q)$. Each intermediate S-range is contained in $S_2 \cup S_3$ and thus contains no points of $X \setminus Y$. One of the intermediate S-ranges will contain another point of Y on its boundary. □

Lemma 7. *For any point set X, positive integer k and a convex compact set S, there is a nice shape S' and a point set X' in general position with respect to S', such that $|X'| = |X|$ and the number of edges in $\mathcal{H}(X',S',k)$ is at least as large as the number of edges in $\mathcal{H}(X,S,k)$.*

Proof. First, we shall modify S slightly. Consider all hyperedges of $\mathcal{H}(X,S,k)$ and for each of them choose a single capturing S-range. Recall that two S-ranges are contraction/dilation of one-another if in their corresponding homothetic maps the origin is mapped to the same point. For each hyperedge consider two distinct capturing S-ranges S_1, S_2 with S_2 being a dilation of S_1. Among all hyperedges, consider the one for which these two S-ranges, $S_1 \subset S_2$ are such that the stretching factor between S_1 and S_2 is the smallest.

Let S' be a nice shape, $S_1 \subset S' \subset S_2$. Replace each of the other capturing ranges with an appropriate S'-range. Now, we have that $\mathcal{H}(X,S',k)$ has at least at many hyperedges as $\mathcal{H}(X,S,k)$. Next, we shall move the points from X slightly so that the new set X' is in a general position with respect to S' and contains as many hyperedges as $\mathcal{H}(X,S',k)$.

Observe first that since X is a finite point set, we can move each point of X by some small distance, call it ϵ in any direction such that the resulting hypergraph has the same set of hyperedges as $\mathcal{H}(X,S',k)$.

Call a point $x \in X$ *bad* if either x is on a vertical line together with some other point of X, x is on a line with two other points of X, or x is on the boundary of an S'-range together with at least three other vertices of X.

We shall move a bad x such that a new point set has smaller number of bad points and such that the resulting hypergraph has at least as many edges as $\mathcal{H}(X, S', k)$. From a ball $B(x, \epsilon)$ delete all vertical lines passing through a point of X, delete all lines that pass through at least two points of X and delete all boundaries of all S'-ranges containing at least 3 points of X. All together we have deleted at most $n + \binom{n}{2} + \binom{n}{3}$, where $n = |X|$, curves because there is one vertical line passing through each point, at most $\binom{n}{2}$ lines passing through some two points of X and at most $\binom{n}{3}$ S'-ranges having some three points of X on their boundary. So, there are points left in $B(x, \epsilon)$ after this deletion. Replace x with an available point x' in $B(x, \epsilon)$. Observe that x' is no longer bad in a new set $X - \{x\} \cup \{x'\}$. Moreover, if $z \in X$, $z \neq x$ was not a bad point, it is not a bad point in a new set $X - \{x\} \cup \{x'\}$. Indeed, since x' is not on a vertical line with any other point of X and not on any line containing two points of X, z is not on a bad line with x'. Moreover, since x' is not on the boundary of an S'-range together with at least three other points of X, z can not be together with x' on the boundary of an S'-range containing more than 3 points of X on its boundary. □

Lemma 8. *Let p, q be two points such that no four points in $X \cup \{p, q\}$ lie on the boundary of an S-range. Let L be a halfplane defined by $\overline{pq}$. Then the following holds.*

(i) The S-ranges in $\mathcal{S}(p, q)$ are linearly ordered, denoted by $\prec_{pq}$, by inclusion of their intersection with L:

$$S_1 \prec_{pq} S_2 \quad \Leftrightarrow \quad S_1 \cap L \subset S_2 \cap L \qquad \text{for all } S_1, S_2 \in \mathcal{S}(p, q)$$

(ii) For each $S_1 \in \mathcal{S}(p, q)$ there exists a $\prec_{pq}$-minimal $S_2 \in \mathcal{S}(p, q)$ with

$$(\partial S_2 \setminus \partial S_1) \cap X \neq \emptyset \qquad \text{and} \qquad S_1 \prec_{pq} S_2$$

if and only if $S_1 \triangle L$ contains a point from X in its interior.

Proof. (i) This follows immediately from Lemma 6 (ii).

(ii) First assume that $(S_1 \triangle L) \cap X \neq \emptyset$. For each point $r \in (S_1 \triangle L) \cap X$ consider the S-range S_r with $\{p, q, r\} \subset \partial S_r$. The existence and uniqueness of S_r is given by Lemma 6 (iii). By the first item, the S-ranges in $\{S_r \mid r \in (S_1 \triangle L) \cap X\}$ are linearly ordered by inclusion of their intersection with L. Hence there exists an S-range S_2 in this set, which is $\prec_{pq}$-minimal.

Let r be the point in ∂S_2 different from p and q. If $r \in L$, then $r \notin S_1$ and thus $L \cap S_1 \subset L \cap S_2$. If $r \notin L$, then $r \in S_1 \setminus \partial S_1$ and thus $S_2 \cap R \subset S_1 \cap R$, where R is the other halfplane defined by $\overline{pq}$. In any case we have $S_1 \prec_{pq} S_2$. Moreover, $r \in (\partial S_2 \setminus \partial S_1) \cap X$ and hence S_2 is the desired S-range.

Now assume that S_2 is a $\prec_{pq}$-minimal S-range in $\mathcal{S}(p, q)$ with $(\partial S_2 \setminus \partial S_1) \cap X \neq \emptyset$ and $S_1 \prec_{pq} S_2$. We claim that the point r in $(\partial S_2 \setminus \partial S_1) \cap X$ lies in the interior of $S_1 \triangle L$. If $r \in R$, then $r \in S_1$, because $\partial S_2 \cap R \subset S_1 \cap R$. On the other hand, if $r \in L$, then $r \notin S_1$, because $\partial S_1 \cap L \subset S_2 \cap L$. Hence r lies in the interior of $S_1 \triangle L$, as desired. □

Splitting Planar Graphs of Girth 6 into Two Linear Forests with Short Paths

Maria Axenovich,* Torsten Ueckerdt, and Pascal Weiner

KARLSRUHE INSTITUTE OF TECHNOLOGY
DEPARTMENT OF MATHEMATICS
KARLSRUHE, GERMANY
E-mail: maria.aksenovich@kit.edu; torsten.ueckerdt@kit.edu; Pascal.Weiner@gmx.de

Received July 9, 2015; Revised August 18, 2016

Published online in Wiley Online Library (wileyonlinelibrary.com).
DOI 10.1002/jgt.22093

Abstract: Recently, Borodin, Kostochka, and Yancey (Discrete Math 313(22) (2013), 2638–2649) showed that the vertices of each planar graph of girth at least 7 can be 2-colored so that each color class induces a subgraph of a matching. We prove that any planar graph of girth at least 6 admits a vertex coloring in two colors such that each monochromatic component is a path of length at most 14. Moreover, we show a list version of this result. On the other hand, for each positive integer $t \geq 3$, we construct a planar graph of girth 4 such that in any coloring of vertices in two colors there is a monochromatic path of length at least t. It remains open whether each planar graph of girth 5 admits a 2-coloring with no long monochromatic paths. © 2016 Wiley Periodicals, Inc. J. Graph Theory 00: 1–18, 2016

Keywords: *planar graph decomposition; defective coloring; fragmented coloring; path forests; 2-coloring*

1. INTRODUCTION

In this article, we consider the question of partitioning the vertex set of a planar graph into a small number of parts, also referred to as color classes, such that each part induces

*Author supported

a graph whose connected components are short paths. The length of a path is the number of its edges. The Four Color Theorem [3, 4] implies that four parts are sufficient to guarantee such a partition with paths of length 0, that is, on one vertex each. A result of Goddard [20] and Poh [26] shows that any planar graph can be vertex-colored with three colors such that each monochromatic component is a path. However, one cannot always restrict the lengths of monochromatic paths in 3-colorings of planar graph as was shown by a specific triangulation construction of Chartrand, Geller, and Hedetniemi [11]. Chappell, Gimbel, and Hartman [10] gave an explicit construction of a planar graph of girth 4 that can not be vertex colored in two colors such that each color class induces a path forest.

However, when the girth of a planar graph is sufficiently large, one can not only 3-color, but 2-color the vertices of the graph such that monochromatic components are short paths. Borodin, Kostochka, and Yancey [9] proved that the vertices of each planar graph of girth at least 7 can be 2-colored so that each monochromatic component has at most two vertices, that is, is a path of length at most 1. Note that the order of monochromatic components cannot be decreased to 1 as long as the graph is not bipartite.

Chappell, Gimbel, and Hartman [10] proved that any planar graph of girth at least 6 can be 2-colored such that each monochromatic component is a path, however no bound on the sizes of these paths was given. Borodin and Ivanova [8] conjectured that there is such a coloring with monochromatic components being paths of length at most 2.

Here, we show that planar graphs of girth at least 6 can be 2-colored such that each monochromatic component is a path of length at most 14. Moreover, we prove a list version of this result. On the other hand, for each positive integer $t \geq 3$, we construct a planar graph of girth 4 such that in any coloring of vertices in two colors there is a monochromatic path of length at least t. It remains open whether one can 2-color the vertices of a planar graph of girth 5 such that each monochromatic component is a short path.

Note that the problem we consider is a problem of *strong linear arboricity* or a *k-path chromatic number* introduced by Borodin et al. [8] and Akiyawa et al. [1], respectively. Here, a linear arboricity of a graph is the smallest number of parts in a vertex-partition of the graph such that each part induces a forest with path components. The k-strong linear arboricity or *k-path chromatic number* is the smallest number of colors in a vertex-coloring of the graph such that each monochromatic component is a path on at most k vertices.

Let L be a color list assignment for vertices of a graph G, that is, $L : V(G) \to 2^{\mathbb{Z}}$. We say that c is an L-coloring if $c : V \to \mathbb{Z}$ such that $c(v) \in L(v)$ for each $v \in V(G)$.

We prove the following theorems.

Theorem 1. *For any planar graph of girth at least 6 and any list assignment L with lists of size* 2 *there is an L-coloring so that each monochromatic component is a path of length at most* 14.

Theorem 2. *For every positive integer t there is a planar graph G_t of girth* 4 *such that any vertex coloring of G_t in two colors results in a monochromatic path of length $t - 1$.*

Our results are a contribution to the lively and active field of *improper vertex colorings of planar graphs*, where the number of colors is strictly less than four but various restrictions on the monochromatic components are imposed. For standard graph theoretic notions used here, we refer to [15].

1.1. Organization of the Article

In Section 2, we give a short survey of improper colorings of planar graphs, explain the relation to our results in the present article, and point out some open problems. In Section 3, we prove Theorem 1 and in Section 4 we prove Theorem 2. We conclude with some open questions in Section 5.

2. IMPROPER COLORINGS OF PLANAR GRAPHS

A proper vertex-coloring of a graph is a coloring in which each monochromatic component is a single vertex, or, equivalently, in which there are no two adjacent vertices of the same color. In this article, a c-coloring, $c \geq 1$, of a graph is a (not necessarily proper) vertex coloring using c colors. As every planar graph has a proper 4-coloring, we focus here on 2-colorings and 3-colorings. The most studied variants of improper colorings are *defective*, *fragmented*, and *P_k-free colorings*. A survey on the topic was done in the bachelor's thesis of Pascal Weiner [28].

2.1. Defective Colorings

For a nonnegative integer k, a vertex coloring is called *k-defective* if each monochromatic component has maximum degree at most k. We define $k_d(g, c)$ to be the smallest k such that every planar graph of girth at least g admits a k-defective c-coloring. Defective colorings were introduced in 1986 by Cowen, Cowen, and Woodall [13], who showed that $k_d(3, 3) = 2$, that is, every planar graph admits a 3-coloring in which every monochromatic component has maximum degree at most 2. In fact, there is a 3-coloring of any planar graph in which every monochromatic component is a path [26]. Eaton and Hull [16], and independently Škrekovski [27], proved that $k_d(3, 2) = \infty$, that is, there are planar graphs of girth 3 for which any 2-coloring results in a monochromatic component of arbitrarily high maximum degree. Cowen, Goddard, and Jerum [14] proved that every outerplanar graph admits a 2-defective 2-coloring. Havet and Sereni [21] showed that for $c \geq 2$, $k \geq 0$ every graph of maximum average degree less than $c + \frac{ck}{c+k}$ admits a k-defective c-coloring. By Euler's formula a planar graph of girth g has maximum average degree less than $\frac{2g}{g-2}$. Hence, the last result implies that $k_d(5, 2) \leq 4$ and $k_d(6, 2) \leq 2$. The result of Borodin et al. [9] shows that $k_d(7, 2) = 1$.

2.2. Fragmented Colorings

A c-coloring is *k-fragmented* if each monochromatic component has at most k vertices, and $k_f(g, c)$ denotes the smallest k such that every planar graph of girth at least g admits a k-fragmented c-coloring. Fragmented coloring were first introduced in 1997 by Kleinberg et al. in [23], where they showed that $k_f(3, 3) = \infty$, that is, there is no k such that every planar graph admits a k-fragmented 3-coloring, a result that has been independently proven by Alon et al. [2]. Esperet and Joret [17] recently proved that $k_f(4, 2) = \infty$,

TABLE I. Improper 2-coloring results for planar graphs of girth g

Girth g	3	4	5	6	7
$k_d(g, 2)$	∞	∞	≥ 2 Figure 7	≤ 2	1
	[16]	[27]	≤ 4 [21]	[21]	[9]
$k_f(g, 2)$	∞	∞	≥ 3	≤ 12	2
	[2, 23]	[17]	Figure 7	[18]	[9]
$k_p(g, 2)$	∞	∞	≥ 4	≤ 6	3
	[1, 7, 11]	Thm. 2	Figure 7	[19]	[9]

although this already follows from the fact that $k_d(4, 2) = \infty$ [27]. Esperet and Ochem [18] proved that $k_f(6, 2) \leq 12$.

2.3. P_k-free Colorings

Finally, a c-coloring is *P_k-free* if there is no monochromatic path on k vertices, and $k_p(g, c)$ denotes the smallest k such that every planar graph of girth at least g admits a P_k-free c-coloring. Such P_k-free colorings were already introduced in 1968 by Chartrand, Geller, and Hedetniemi [11], who showed that $k_p(3, 3) = \infty$, that is, there is no k such that every planar graph admits a P_k-free 3-coloring. In a different paper [12], the same authors showed that same holds for outerplanar graphs and two colors. More than 20 years later, the former result has been reproved by Akiyama et al. [1], as well as Berman and Paul [7]. Recently, Glebov and Zambalaeva [19] showed that every planar graph of girth at least 6 can be 2-colored such that every color class induces a P_6-free forest, that is, $k_p(6, 2) \leq 6$. We summarize the results for defective, fragmented, and P_k-free colorings using two colors.

Theorems 1 and 2 immediately imply the following (c.f. Table I).

Corollary 1. *We have that $k_f(6, 2) \leq 15$, $k_p(6, 2) \leq 16$, and $k_p(4, 2) = \infty$.*

Let us also mention that defective and fragmented colorings have also been considered for nonplanar graphs of bounded maximum degree [2, 6, 22], bounded number of vertices [24], and for minor-free graphs [29]. In natural generalizations, one allows different color classes to have different defect (see, e.g., [5, 25]), or considers list-coloring, which, in fact, is the case in many of the results above.

3. PROOF OF THEOREM 1

For a list assignment L, we call an L-coloring of a planar graph *good* if each monochromatic component is a path of length at most 14. Throughout this section we let, for the sake of contradiction, a graph G be a counterexample to Theorem 1, so that G is vertex-minimal, and among all such graphs has the largest number of edges. That is, G has no good L-coloring, an addition of any new edge to G creates a nonplanar graph or a cycle of length at most 5, and any subgraph of G with fewer vertices has a good L-coloring. To avoid a special treatment of an outer face we assume G to be embedded without crossings on the sphere and shall refer to the faces of the corresponding plane graph as faces of G.

Note also that if a graph has no good L-coloring, then any of its supergraphs has no good L-coloring.

3.1. Idea of the Proof

Our proof extends the ideas of Havet and Sereni [21]. We start by proving some structural properties of G, that is, that G has minimum degree 2, all faces of G are chordless cycles of length at most 9, and proving a statement about the distribution of vertices of degree 2 around every face F in G.

If G has a path P of length at most 14 with endpoints of degree 2 and all inner vertices of degree 3, then each vertex in P has exactly one neighbor not in P. Deleting the vertices of P from G gives a graph that has a good coloring. Color each vertex of P with a color different from the color of its neighbor not in P. This gives a good coloring of G contradicting the fact that G is a minimal counterexample.

We generalize this simple argument, which uses a single path, to path systems, that is, sets of (directed) facial paths in G with all inner vertices of degree 3. Next, we consider a charge of $\deg(v) - 3$ at every vertex v and define discharging rules shifting a charge of $1/2$ from the out-endpoint of every path in X_0 to its in-endpoint, based on a specific path system X_0. The total charge on all the vertices, before as well as after the discharging, is negative, giving some vertices ending up with negative charge. We consider such a vertex w_0, build another path system based on what is "outgoing" from this vertex, and show that the corresponding subgraph of G is a reducible configuration. Here, a subgraph H is reducible if any good L-coloring of $G - V(H)$ (which exists by the minimality of G) could be extended to a good L-coloring of the whole graph G. This contradicts the assumption that G is a counterexample and hence concludes the proof.

3.2. Structural Properties of *G*

Lemma 1. *G is connected and has minimum degree at least* 2.

Proof. Indeed, if G has a vertex v of degree 1, then a good coloring of $G - v$ can be extended to a good coloring of G by choosing the color of v to be different from its neighbor in $G - v$. If G is not connected, then one of its connected components is a smaller counterexample, contradicting the definition of G. ■

Lemma 2. *The boundary of each face of G forms a chordless cycle of length at most* 9.

Proof. First assume for the sake of contradiction there is a face F whose closed boundary walk $W = u_0, \ldots, u_m$ is not a cycle. Then there is a vertex u appearing at least twice on W, say $u = u_0 = u_j$ with $j \neq 0$. As G has minimum degree 2, each of the closed walks $W_1 = u_0, u_1, \ldots, u_j$ and $W_2 = u_j, u_{j+1}, \ldots, u_0$ contains at least one cycle, that is, has at least six vertices. Note that the vertices u_2 from W_1 and u_{j+3} from W_2 lie in distinct connected components of $G - u$. Moreover, as G has minimum degree 2, u_2 and u_{j+3} are at distance 2 and 3 from u along W, respectively. Hence any $u_2 - u_{j+3}$ path goes through u and, as there are no cycles of length at most 5, the distance between u_2 and u_{j+3} in G is 5. Thus we can add an edge u_2u_{j+3} into F, creating a planar graph with girth at least 6. A contradiction to edge-maximality of G.

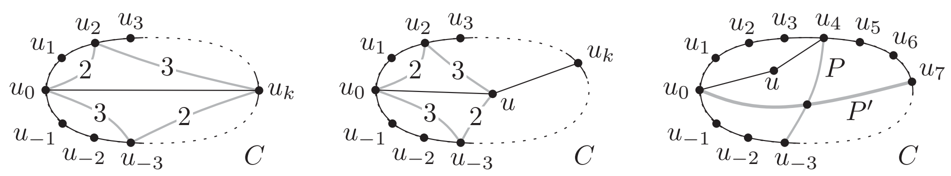

FIGURE 1. Illustration of Case 1 (left), Case 2 with $w \in \{u_0, u\}$ (middle), and Case 2 with $w = u_k$ (right). The face F bounded by C is shown as the outer face. The numbers indicate the minimum length of a path between the corresponding vertices.

Thus, the boundary of each face F forms a cycle, $C = u_0, \dots, u_m, u_0$. Assume that C has length at least 10, that is, that $m \geq 9$. Recall, that an *ear* E of a cycle C is a path that shares only its endpoints with the vertex set of the cycle. For $i = 0, \dots, m$ let $G'(i)$ be obtained from G by adding an edge u_i, u_{i+5} into the face F, addition of indices modulo $m + 1$. If $G'(i)$ has girth at least 6, this contradicts the edge-maximality of G. So, there is a cycle on at most five vertices containing edge u_iu_{i+5} in $G'(i)$, denote a shortest $u_i - u_{i+5}$ path in G by $P(i, i + 5)$. Its length is at most 4, less than the distance between i and $i + 5$ along C (as $m \geq 9$), so there is an ear of length ℓ, $\ell \leq 4$, and ℓ is less than the distance between its endpoints along C. The *width* of an ear is the smallest distance between its endpoints along the cycle. If Q is a path or a cycle and P is a path in Q with endpoints u and v, we write $P = uQv$. We denote the length of P as $||P||$. A *k-ear* is an ear of length k.

Case 1. C has a chord.

Assume that u_0u_k is a chord, $k \geq 5$. A path $P = P(-3, 2)$ must contain u_0 or u_k. If P contains u_0, then $||u_{-3}Pu_0|| \geq 3$ and $||u_0Pu_2|| \geq 2$, as otherwise $P \cup C$ contains a cycle of length at most 5. Similarly, if P contains u_k, then $||u_{-3}Pu_k|| \geq 2$ and $||u_kPu_2|| \geq 3$. In any case we have that $||P|| \geq 5$, a contradiction. See Figure 1 left.

Case 2. C has an ear of length 2 and no chords.

Let E be a 2-ear of smallest width, with vertices u_0, u, u_k, $4 \leq k \leq (m + 1)/2$. A path $P = P(-3, 2)$ contains $w \in \{u_0, u, u_k\}$. If $w = u_0$, see Figure 1 center, then (as in Case 1) $||u_{-3}Pu_0|| \geq 3$ and $||u_0Pu_2|| \geq 2$, and if $w = u$, then $||u_2Pu|| \geq 3$ and $||uPu_{-3}|| \geq 2$, as otherwise there is a cycle of length at most 5 in $P \cup C$. In both cases we have $||P|| \geq 5$, a contradiction. So $w = u_k$, see Figure 1 right, $||wPu_2|| \geq 2$, and $||u_{-3}Pw|| \geq 2$, otherwise there is a chord. Thus each of these segments has length 2. Since $||u_kPu_2|| = 2$, u_kPu_2 is a subpath of C, otherwise there is a 2-ear of a smaller width. So $k = 4$. Since $||u_{-3}Pu_k|| = 2$ and $m \geq 9$, $||u_{-3}Cu_k|| \geq 4$, so $||C|| \geq 11$. Looking at E in the other direction along C, and taking $P' = P(2, 7)$, we see symmetrically that $u_0P'u_7$ is an ear of length 2, that together with E and P creates a cycle of length 4.

Case 3. C has no ears of length 2 or chords.

Since each $P(i, i + 5)$ results in an ear whose length is smaller than its width, we see that either there is a u_i-u_{i+5} ear of length 4 or an ear of length 3 with width between 4 and 6. If all such ears are of length 4, then a $u_0 - u_5$ ear and $u_1 - u_6$ ear intersect and form, together with C, a cycle of length at most 5, a contradiction, see Figure 2 left. Assume that E is a 3-ear u_0, u, u', u_k, $4 \leq k \leq 6$, and all other 3-ears have width either at most 3 or at least k, see Figure 2 center and right. A path $P = P(2, -3)$ contains a 3- or a 4-ear. Let w be a point on P and E. We have that $||u_{-3}Pw|| \geq 2$ and $||wPu_2|| \geq 2$, otherwise there is either a chord, a 2-ear, or a cycle of length at most 5. It follows that

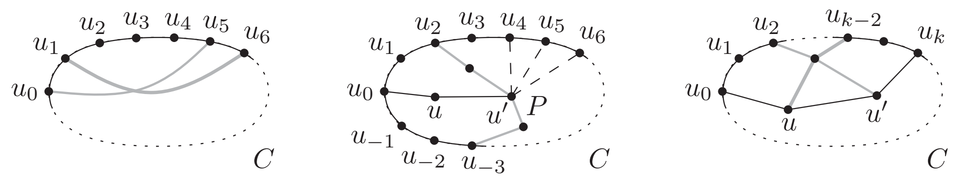

FIGURE 2. Illustration of Case 3 with only 4-ears (left) and Case 3 with a 3-ear (middle and right). The face F bounded by C is shown as the outer face.

$||u_{-3}Pw|| = 2 = ||wPu_2||$. Then $w = u'$, and $k \geq 5$. Looking at E in the other direction, we see symmetrically, that there is a path of length 2 between u_{k-2} and u, implying the existence of a triangle containing u and u', a contradiction.

Thus C has length at most 9. If C has a chord, then there is a cycle of length at most 5, a contradiction. So, C is a chordless cycle. This concludes the proof the Lemma. ■

Lemma 3. *Let F be any face of G incident to a vertex of degree 2. Then F is incident to a vertex of degree at least 4, and if F is incident to at least two vertices of degree 2, then there is a vertex of degree at least 4 between any two such vertices on both paths along F.*

Proof. Let C be the simple chordless cycle bounding F. First, assume for the sake of contradiction that C contains exactly one vertex v of degree 2 and all other vertices of degree 3. Consider a good L-coloring of $G' = G - V(C)$ and give each vertex u of C of degree 3 a color in $L(u)$ different from the color of its neighbor in G'. Give v a color in $L(v)$ such that C does not form a monochromatic cycle. As a result, the set of monochromatic components of G is formed by the monochromatic components of G', and paths on at most 8 vertices formed by vertices of C.

Second, assume that C contains two vertices u, v of degree 2 and a $u - v$ path P in C has no inner vertices or only inner vertices of degree 3. Consider a good L-coloring of $G' = G - V(P)$ and give the vertices of P colors from their lists, different from the colors of their unique neighbors in G'. This does not extend any connected monochromatic component of G' and every new monochromatic component is contained in P, that is, a path on at most eight vertices. That is, in both cases we have found a good L-coloring of G, a contradiction to G being a counterexample. ■

3.3. Path Systems

A *path system* is a set X of (not necessarily edge-disjoint) directed facial paths in G with all inner vertices being of degree 3, such that no vertex is an endpoint of one path in X and an inner vertex of another path in X. For a path $P \in X$ directed from vertex u to vertex v, we call u the *out-endvertex* and v the *in-endvertex* of P. For a path system X, the vertices that are the in-endvertices or out-endvertices of some path in X are called the *endvertices of* X, while the *inner vertices of* X are the inner vertices of some path in X. For any vertex v in G let out-$\deg_X(v)$ and in-$\deg_X(v)$ denote the number of paths in X with out-endvertex v and in-endvertex v, respectively. Note that for an inner vertex v of X we have out-$\deg(v) =$ in-$\deg(v) = 0$. A directed path P is *occupied* by a path system X if the first or last edge of P (incident to its out-endvertex or in-endvertex) is contained in some path in X. So, if $P \in X$, then P is occupied by X. Let us emphasize that throughout

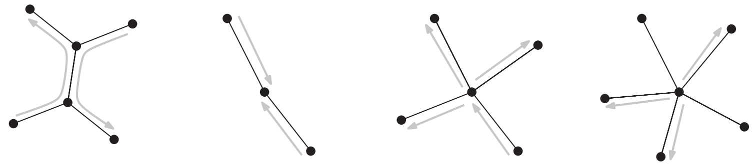

FIGURE 3. Illustration of properties (**D1**)–(**D5**).

the article $\deg(v)$ and $N(v)$ always refer to the degree and neighborhood of vertex v in G, even when we consider other subgraphs of G later.

For a path system X and any two vertices u, v in G we say that *u reaches v in X*, denoted by $u \to_X v$, if there is a sequence $u = v_1, \dots, v_k = v$ of vertices and a sequence $P_1, \dots, P_{k-1}$ of paths in X such that v_i and v_{i+1} are out-endvertex and in-endvertex of P_i, respectively, $i = 1, \dots, k-1$. Then X is *acyclic* if there are no two distinct vertices u, v with $u \to_X v$ and $v \to_X u$. For a vertex w of G, we define $X^+(w) \subseteq X$ to be the path system consisting of all paths in X whose out-endvertex is w or reachable from w in X.

A path system X is *nice* if each of the following properties (**D1**)–(**D5**) holds. A path system X with a distinguished vertex r, called root, is *almost nice* if the properties (**D1**)–(**D5**) hold for all vertices different from r.

See Figure 3 for an illustration.

(D1) Every edge that belongs to two paths in X joins two vertices of degree 3 each.
(D2) Every vertex of degree 2 has outdegree 0 in X.
(D3) Every vertex of degree 3 has indegree 0 and outdegree 0 in X.
(D4) Every vertex of degree 4 has positive indegree in X only if it has outdegree 3 in X.
(D5) Every vertex of degree at least 5 has in-degree 0 in X.

The following statements follow immediately from the definitions above.

Lemma 4. *For every path system X each of the following holds.*

(1) *If no path $P \in X$ is occupied by $X - \{P\}$, then X satisfies (**D1**).*
(2) *If $X' \subseteq X$ and X satisfies any of (**D1**)–(**D3**), (**D5**), then so does X'.*
(3) *If $X' \subseteq X$ and X is acyclic, then so is X'.*
(4) *If X is nice and w is a vertex, then $X^+(w)$ with root w is almost nice.*

3.4. Discharging with respect to a Path System X

Given a path system X, consider the following discharging: Put charge $ch(v) = \deg(v) - 3$ on each vertex of G. Note that $ch(v) = -1$ for a vertex of degree 2, and $ch(v) \geq 0$ for all other vertices. As all facial cycles have length at least 6, we have $6f \geq 2e$, where f denotes the number of faces of G. Together with Euler's formula $n - e + f = 2$ this implies $n - e + e/3 \geq 2$. Thus, the total charge is $\sum_{v \in V(G)} (\deg(v) - 3) = 2e - 3n \leq -6$.

Define $ch'(v) = ch(v) + \frac{1}{2}(\text{in-deg}_X(v) - \text{out-deg}_X(v))$. Intuitively, for every path in X a $1/2$-charge is sent from out-endvertex to in-endvertex. Thus, the total sum of charges in ch' is the same as in ch, that is, $\sum_v ch(v) = \sum_v ch'(v)$.

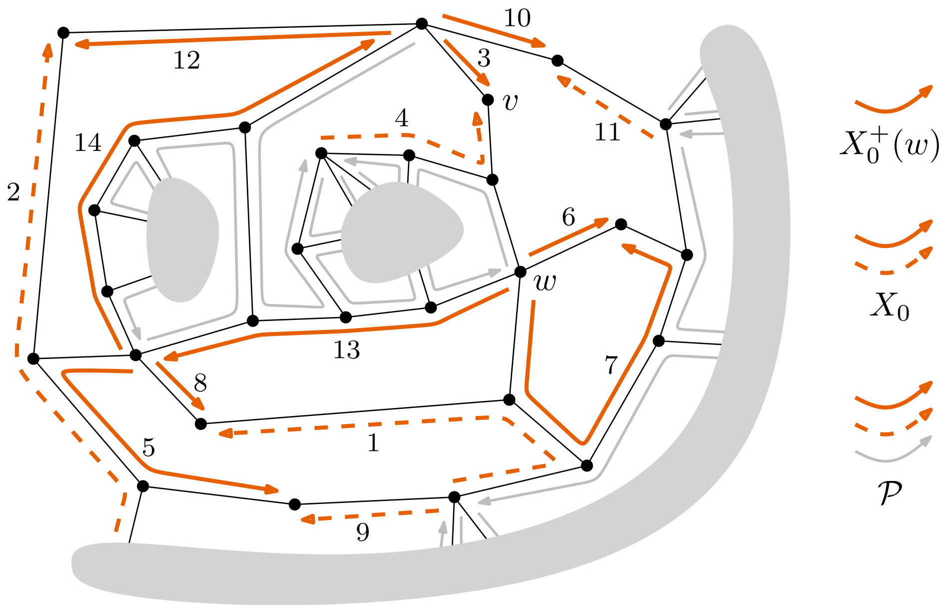

FIGURE 4. A (part of a) planar graph of girth 6 and the path systems in $\mathcal{P}$, X_0 and $X_0^+(w)$. The labels show the order in which paths of X_0 were selected. Paths 13 and 14 were selected in step 2.

3.5. Defining a Path System $\mathcal{P}$

As G has girth at least 6, there is a vertex v of degree 2 in G and by Lemma 3 both faces incident to v contain a vertex of degree at least 4. So there are faces with at least two vertices of degree different from 3. For each such face F the boundary of F can be uniquely partitioned into edge-disjoint counterclockwise oriented paths with all inner vertices of degree 3 and endpoints of degree different from 3. We denote by $\mathcal{P}$ the path system consisting of all such paths with in-endvertex of degree 2 or 4 and out-endvertex of degree at least 4, for all faces F with at least two vertices of degree different from 3. So for each path in $\mathcal{P}$ the degrees d_1, d_2 of its in-endvertex and out-endvertex, respectively, satisfy $(d_1, d_2) \in \{(2, 4), (2, \ell), (4, 4), (4, \ell) \mid \ell \geq 5\}$.

By Lemma 2 every face of G is bounded by a simple chordless cycle of length at most 9. Thus, every $P \in \mathcal{P}$ is a path on at most eight edges. As any two paths in $\mathcal{P}$ in the boundary of the same face F are edge-disjoint, every edge of G lies in at most two paths in $\mathcal{P}$, at most one for each face incident to the edge. If an edge lies in two paths in $\mathcal{P}$, these paths have the edge oriented in opposite directions. For a vertex v in G with $\deg(v) = 3$ we have $\text{out-deg}_{\mathcal{P}}(v) = \text{in-deg}_{\mathcal{P}}(v) = 0$ by definition. Note that by Lemma 3 for every vertex v with $\deg(v) = 2$ we have $\text{out-deg}_{\mathcal{P}}(v) = 0$ and $\text{in-deg}_{\mathcal{P}}(v) = 2$. For a vertex v with $\deg(v) \geq 5$ we have $\text{in-deg}_{\mathcal{P}}(v) = 0$, that is, $\mathcal{P}$ has properties (**D2**), (**D3**) and (**D5**). We provide an example illustrating these concepts in Figure 4.

3.6. Defining a Path System $X_0 \subseteq \mathcal{P}$

We define $X_0 \subseteq \mathcal{P}$ selecting paths one by one, using the following procedure, where we go through the vertices in question in an arbitrary but fixed order. At all times, let X_0 denote the set of already chosen paths, initially $X_0 = \emptyset$.

(1) For every vertex v with $\deg(v) = 2$ we put a path from $\mathcal{P}$ into X_0 if its in-endpoint is v and if it is not occupied by X_0.

After step 1 is done for all vertices of degree 2, we proceed as follows.

(2) For every vertex v with $\deg(v) = 4$ and out-$\deg_{X_0}(v) = 3$, put a path from $\mathcal{P}$ into X_0 if its in-endpoint is v and if it is not occupied by X_0.

Later, we shall show that the final path system X_0 is nice and acyclic. For now, we only need to observe that (**D2**) is satisfied and in-$\deg_{X_0}(u) = 2$ for every vertex u of degree 2. In fact, (**D2**) holds for $\mathcal{P}$ and thus by Lemma 4 (2) it also holds for $X_0 \subseteq \mathcal{P}$. Assume now that in-$\deg_{X_0}(u) < 2$. That is, P, one of the two paths in $\mathcal{P}$ with in-endvertex v, was occupied by during step 1. As all in-endvertices of paths chosen in step 1 are of degree 2, and the out-endvertex of P has degree at least 4, this is impossible.

3.7. Defining the Vertex w_0 Based on Discharging with respect to X_0

Let us apply discharging to X_0. For every vertex u with $\deg(u) = k$, we have in-$\deg_{X_0}(u) \geq 0$ and out-$\deg_{X_0}(u) \leq k$, that is, u looses a charge of at most $\frac{k}{2}$. Thus if $\deg(u) = k \geq 6$, the remaining charge $ch'(u)$ is at least $k - 3 - \frac{k}{2} \geq 0$. If $\deg(u) = 3$, then out-$\deg_{X_0}(u) =$ in-$\deg_{X_0}(u) = 0$ and hence $ch(u) = ch'(u) = 0$. If $\deg(u) = 2$, then in-$\deg_{X_0}(u) = 2$ and out-$\deg_{X_0}(u) = 0$ and hence $ch'(u) = \deg(u) - 3 + \frac{1}{2}(2 - 0) = 0$.

On the other hand, we have $\sum_v ch(v) = \sum_v ch'(v)$. As $\sum_v ch'(v) \leq -6$ there is a vertex w_0 in G with $ch'(w_0) < 0$. With the above considerations we conclude that $\deg(w_0) \in \{4, 5\}$.

If $\deg(w_0) = 5$, then $0 > ch'(w_0) \geq (5 - 3) - \frac{1}{2}$out-$\deg_{X_0}(w_0)$, so out-$\deg_{X_0}(w_0) \geq 5$. Since out-$\deg_{X_0}(w_0) \leq \deg(w_0)$, we have that out-$\deg_{X_0}(w_0) = 5$. If $\deg(w_0) = 4$, then $0 > ch'(w_0) = (4 - 3) + \frac{1}{2}($in-$\deg_{X_0}(w_0) -$ out-$\deg_{X_0}(w_0))$, so either out-$\deg_{X_0}(w_0) = 4$ or (out-$\deg_{X_0}(w_0) = 3$ and in-$\deg_{X_0}(w_0) = 0$). In particular, exactly one of the following must hold for the vertex w_0 with $ch'(w_0) < 0$:

Case 1: $\deg(w_0) \in \{4, 5\}$ and out-$\deg_{X_0}(w_0) = \deg(w_0)$.

Case 2: $\deg(w_0) = 4$, out-$\deg_{X_0}(w_0) = 3$ and in-$\deg_{X_0}(w_0) = 0$.

For example, in Figure 4 we see that **Case 2** applies to vertex w.

3.8. Defining Rooted Path Systems X_1, X_2, X_3, X_4 Based on w_0 and X_0

Depending on the structure of w_0 and X_0 we shall define one of four path systems X_1, X_2, X_3, X_4, each X_i with a specified vertex w_i, called the *root*, $i = 1, 2, 3, 4$. Path systems X_1, X_3 will be chosen as subsystems of X_0, X_2 as a subsystem of X_0 together with an additional path from $\mathcal{P}$, and X_4 as a subsystem of X_0 together with a subpath of a path from $\mathcal{P}$. Note that each of X_1, X_2, X_3, X_4 consists of paths of length at most 8.

Case 1: $\deg(w_0) \in \{4, 5\}$ and out-$\deg_{X_0}(w_0) = \deg(w_0)$.

In this case we define $X_1 = X_0^+(w_0)$ with root w_0.

Case 2: $\deg(w_0) = 4$, out-$\deg_{X_0}(w_0) = 3$ and in-$\deg_{X_0}(w_0) = 0$.

Consider the unique edge e at w_0 not contained in any path in X_0. As w_0 has outdegree 3, the clockwise next edge e' at w_0 after e is contained in some path in X_0 with out-endvertex w_0. The in-endvertex of this path is in the face F incident to w_0, e, and e'. See the middle

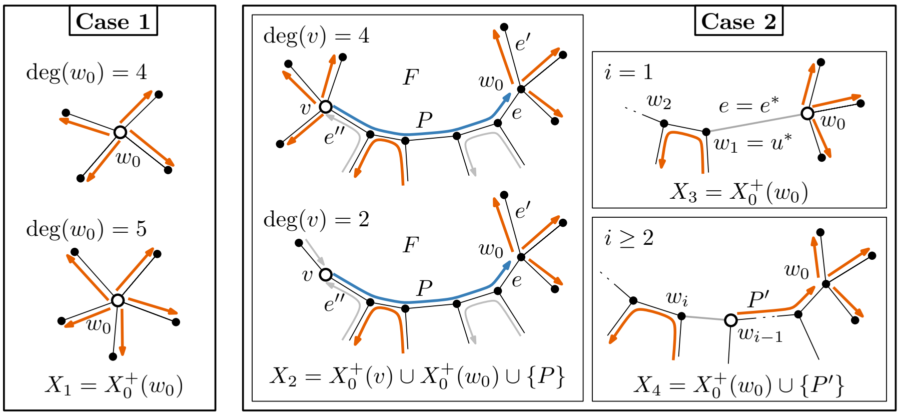

FIGURE 5. Illustrations of the rooted path systems X_1, X_2, X_3, X_4 with highlighted roots.

part of Figure 5 for an illustration. So F has at least two vertices of degree different from 3. Thus its boundary contains a counterclockwise path P with in-endvertex w_0, using the edge e, all inner vertices of degree 3 or no inner vertices at all and out-endvertex v with $\deg(v) \neq 3$. Let e'' be the edge of P incident to v. If $\deg(v) = 2$, then from the definition of X_0, in-$\deg_{X_0}(v) = 2$. Thus e'' belongs to a path in X_0 with in-endpoint v. If $\deg(v) \geq 4$ then $P \in \mathcal{P}$, and in step 2 of the construction of X_0 the path P must have been rejected because it was occupied, that is, e'' is contained in another path in X_0. As P has v as out-endvertex, the other path has v as in-endvertex and hence $\deg(v) = 4$. So, $\deg(v) \in \{2, 4\}$ and e'' lies in some path in X_0 with in-endvertex v. In particular it follows that $e \neq e''$, that is, P has at least one inner vertex.

Next, we distinguish the cases when v is reachable from w_0 in X_0 or not, corresponding to the right and middle part of Figure 5, respectively. In case v is not reachable from w_0 in X_0, we define $X_2 = X_0^+(v) \cup X_0^+(w_0) \cup \{P\}$ with root v.

When v is reachable from w_0 in X_0, let $w_0, w_1, \ldots, w_{k-1}, w_k = v$, $k \geq 2$, denote the vertices of P in their order along P from its in-endvertex w_0 to its out-endvertex v. Recall that P has at least one inner vertex. Let i be the smallest index such that $w_i \neq w_0$ and w_i is contained in a path in $X_0^+(w_0)$. See the right part of Figure 5. As $v = w_k$ is reachable from w_0 in X_0, this index is well defined. If $i = 1$, we define $X_3 = X_0^+(w_0)$ with root w_0. Otherwise we denote the directed w_{i-1}-to-w_0 subpath of P by P' and define $X_4 = X_0^+(w_0) \cup \{P'\}$ with root w_{i-1}. This is for example the case for vertex $w_0 = w$ in Figure 4.

Lemma 5.

(i) *Each of* X_0, X_1, X_2, X_3, X_4 *is acyclic.*
(ii) X_0 *and* X_1 *are nice.*
(iii) *If the root* v *of* X_2 *has degree* 4*, then* X_2 *is nice.*
(iv) *If the root* v *of* X_2 *has degree* 2*, then* X_2 *is almost nice with out-*$\deg_{X_2}(v) = 1$.
(v) X_3 *is almost nice with out-*$\deg_{X_3}(w_0) = 3$ *and in-*$\deg_{X_3}(w_0) = 0$.
(vi) X_4 *is almost nice with out-*$\deg_{X_4}(r) = 1$ *and in-*$\deg_{X_4}(r) = 0$ *for the root* r *of* X_4.
(vii) *If* $j \in \{1, 2, 3, 4\}$ *then each endvertex of* X_j*, different from the root, has degree* 2 *or* 4 *in* G*, the root has degree* 2, 3, 4*, or* 5*, and each inner vertex of* X_j *has degree* 3*. Moreover, each vertex of* X_j *has at most one neighbor that is not in* X_j.

Proof.

(i) First, we shall show that X_0 is acyclic. Assume for the sake of contradiction that $v_0, \dots, v_{k-1}$ and $P_0, \dots, P_{k-1}$ are two sequences of vertices and paths in X_0 such that for every $i \in \{0, \dots, k-1\}$ we have that v_i and v_{i+1} are out-endvertex and in-endvertex of P_i, respectively (all indices modulo k). For each $i \in \{0, \dots, k-1\}$ we have $\deg(v_i) \in \{2, 4\}$, as we add only paths with such in-endvertices to X_0 in step 1 and 2. Moreover, v_i is out-endvertex of P_{i-1} and thus we have that $\deg(v_i) \neq 2$. Hence for each $i \in \{0, \dots, k-1\}$, we have $\deg(v_i) = 4$ and P_i was put into X_0 in step 2 because v_{i+1} was the out-endvertex of exactly three already chosen paths. Assume without loss of generality that P_0 was the path that was put into X_0 in step 2 first among the paths $P_0, \dots, P_{k-1}$. This means that the path P_1, whose out-endvertex is v_1, was already put into X_0. This contradicts that P_0 was the first and proves that X_0 is acyclic.
Now, $X_1, X_3 \subseteq X_0$ are acyclic by Lemma 4 (3). Moreover, $X_2 = X_0^+(v) \cup X_0^+(w_0) \cup \{P\}$ is acyclic, because $X_0^+(v), X_0^+(w_0) \subseteq X_0$, P has in-endvertex w_0 and out-endvertex v, and (in this case) v is not reachable from w_0 in X_0. Finally, $X_4 = X_0^+(w_0) \cup \{P'\}$ is acyclic, because $X_0^+(w_0) \subseteq X_0$ and $V(P') \cap (\bigcup_{P \in X_0} V(P)) = \{w_0\}$.

(ii and v) Consider X_0. As mentioned earlier, $\mathcal{P}$ satisfies (**D2**), (**D3**), and (**D5**) and by Lemma 4 (2) so does X_0. Moreover, by definition X_0 satisfies (**D1**) and (**D4**), thus X_0 is nice. Consider X_1 and X_3. By Lemma 4 (4) we have that X_1 and X_3 are almost nice, as both are defined as $X_0^+(w_0)$ with root w_0. By construction, $\text{out-deg}_{X_3}(w_0) = 3$ and $\text{in-deg}_{X_3}(w_0) = 0$, which proves (v). For X_1 note that, if $\deg(w_0) = 4$, then (**D4**) holds for w_0 in X_1, and if $\deg(w_0) = 5$, then (**D5**) holds for w_0 since $X_1 \subseteq X_0 \subseteq \mathcal{P}$. Thus, X_1 is nice.

(iii and iv) Next consider $X_2 = X_0^+(v) \cup X_0^+(w_0) \cup \{P\}$ with root v and path P as defined above, see the middle part of Figure 5. Each of $X_0^+(v), X_0^+(w_0)$ is almost nice by Lemma 4 (4) and the niceness of X_0. We have neither $w_0 \to_{X_0} v$ (by assumption) nor $v \to_{X_0} w_0$ (as $\text{in-deg}_{X_0}(w_0) = 0$). So $X_0^+(v) \cup X_0^+(w_0)$ satisfies (**D1**)–(**D5**), except perhaps for w_0 and v. As X_2 additionally contains the path P from v to w_0, we have that (**D4**) is satisfied for w_0 and thus X_2 is nice when $\deg(v) = 2$ and almost nice when $\deg(v) = 4$. Because X_0 is nice, that is, satisfies (**D2**), we have $\text{out-deg}_{X_0}(v) = 0$ when $\deg(v) = 2$. As $X_2 - P \subseteq X_0$ and P is outgoing at v, we have $\text{out-deg}_{X_2}(v) = 1$.

(vi) The system $X_4 = X_0^+(w_0) \cup \{P'\}$ is almost nice, because P' and $X_0^+(w_0)$ share only vertex w_0, $X_0^+(w_0)$ is almost nice by Lemma 4 (4), and P' is incoming at w_0.

(vii) These properties are corollaries of the almost-niceness of X and the considerations for the root in the previous items. ■

3.9. Coloring Reducible Configurations based on X_1, X_2, X_3, X_4

Recall that a coloring is good if each monochromatic component is a path of length at most 14. A *reducible configuration* is a nonempty subgraph H of G, such that any good L-coloring of $G - V(H)$ (which exists by the minimality of G) can be extended to a good L-coloring of G in which every edge between a vertex in H and a vertex outside of H is colored properly. Showing that G has a reducible configuration will conclude the proof of Theorem 1. For convenience we say that X_i is reducible if the subgraph H of G induced by the vertices in X_i is a reducible configuration, $i = 1, 2, 3, 4$.

Lemma 6. *Each of X_1, X_2, X_3, X_4 is reducible, whenever it is defined.*

Proof. Consider $j \in \{1, 2, 3, 4\}$. Let r be the root of X_j, H be the subgraph of G consisting of all vertices and undirected edges in the path system X_j. Let V_1 be the set of vertices of H and H' be the subgraph of G induced by V_1, that is, $H \subseteq H'$. Let $W \subseteq V_1$ be the set of endvertices of X_j. Recall, that by Lemma 5(vii), if $w \in W - \{r\}$, then $\deg(w) \in \{2, 4\}$, $\deg(r) \in \{2, 3, 4, 5\}$, and if $u \in V_1 - W$, then $\deg(u) = 3$. In addition, the niceness or almost niceness of X_j and the degree conditions for r given in Lemma 5 imply that any vertex from V_1 has at most one neighbor not in V_1 and each vertex in $V_1 - \{r\}$ has at most one incident edge from $E(G) - E(H)$. In particular, $E(H') - E(H)$ is a matching, unless $j = 4$, in which case $E(H') - E(H)$ might contain two edges incident to r. In case $j \neq 3$, let $E_1 = E(H') - E(H)$. Otherwise (when $j = 3$) let $e^* = ru^*$ denote the unique edge in $E(H') - E(H)$ incident to the root r and let $E_1 = E(H') - (E(H) \cup e^*)$. In Figure 5 on the right we have $r = w_0$ and $u^* = w_1$. Note that if $j \in \{1, 2\}$, then there are no edges from $E(H') - E(H)$ incident to r.

We shall be coloring different sets of vertices of G one after another.

- First we make a good L-coloring c' of $G - V_1$, which exists by the minimality of G. Note that $G - V_1$ might be empty.

We shall color V_1 so that no vertex in V_1 has the same color as its neighbor (if exists) in $V(G) - V_1$ and such that each monochromatic path with vertices in V_1 is contained in the union of two paths from X_j.

- Consider $A \subseteq V_1$, the set of vertices that have a neighbor in $V(G) - V_1$. As X_j satisfies (**D4**) no vertex of degree 4 in H' is in A, except for possibly the root r. We color each vertex $v \in A$ such that its color is from $L(v)$ and differs from the color of its neighbor in $V(G) - V_1$.

Now, no matter how we color $V_1 - A$, each monochromatic path has all its vertices completely in V_1 or completely in $V(G) - V_1$. Since the coloring of $V(G) - V_1$ is good, each monochromatic path there has length at most 14. So, we only need to color $V_1 - A$ so that each monochromatic path with vertices in V_1 has length at most 14.

- Consider the vertices of E_1. First assume $r \in V(E_1)$, this could be only if $j = 2$ or $j = 4$. If $r \in A$, then r is already colored and if $r \notin A$, we give r any color from its list. Next, we color every neighbor of r in E_1 with a color from its respective list different from the color of r. Finally, we color the remaining vertices of E_1 from their lists such that each edge of E_1 has endpoints of different colors. If $r \notin V(E_1)$, that is, E_1 is a matching, color $V(E_1)$ such that each edge is colored properly.

This ensures that eventually every monochromatic component of H' is a subgraph of H or $H \cup e^*$ in case $j = 3$.

- Consider the set B of vertices from $V_1 - A$ not incident to E_1 and of degree 3. Note that $r \notin B$ because it is either of degree different from 3 or is incident to E_1 in case when $j = 3$. Hence B consists only of inner vertices of X_j, that is, $B = V_1 - (A \cup V(E_1) \cup W)$. For any $u \in B$ all three edges incident to u are in H, so u lies on at least two paths in X_j. We consider the paths in X_j in any order and when we process a path $P \in X_j$, we color the vertices in $B \cap V(P)$. For the current path P and the current vertex $u \in B \cap V(P)$, consider the neighbor u' of u not in P. If u' is not colored, color u arbitrarily from its list. Otherwise, color u with a color different from the color of u'.

This ensures that every monochromatic component of $H' - W$ is completely contained in some path in X_j. It remains to color the vertices in $W - A$ and in $e^* = ru^*$ (if it exists) in such a way that e^* is not monochromatic and at most two monochromatic components of $H' - W$ are part of the same monochromatic component of H'.

- Consider the vertices in $W - A$ and the vertex u^* (if it exists). Recall that u^* is an inner vertex of some path in X_j and hence $u^* \notin W$. For each $u \in W - A$, consider the paths in X with in-endvertex u and let $S(u)$ be the set of immediate neighbors of u on those paths, that is, $v \in S(u)$ if $uv \in E(H)$ and u is the in-endvertex of the path in X_j containing uv. In particular, $S(r) = \emptyset$, and for $u \neq r$ we have $|S(u)| = 1$ if $\deg(u) = 4$ and $|S(u)| = 2$ for $\deg(u) = 2$. Additionally let $S(u^*) = \{r\}$ when considering X_3. We apply the following rules to still uncolored vertices (initially the set $(W - A) \cup \{u^*\}$) as long as any of these is applicable:

Rule 1: If for some uncolored vertex u three of its neighbors have the same color a, we color u with a color in $L(u)$ different from a.

Rule 2: If Rule 1 does not apply, but for some uncolored vertex u some $u' \in S(u)$ is already colored, we color u with a color from its list different from the color of u'.

Rule 3: If neither Rule 1 nor Rule 2 applies, and the root r is uncolored, consider the set of colors appearing on $N(r)$ and a color a that is repeated the most in $N(r)$. Let $b \in L(r) - \{a\}$. Then b is repeated at most twice in $N(r)$ since $|N(r)| \leq 5$. Moreover, b is repeated at most once in $N(r)$ if $|N(r)| = 2$ or 3. Assign color b to r.

We claim that if none of the three rules applies, then all vertices are colored. Indeed, if neither Rule 1 nor Rule 2 applies and some vertex u_1 is uncolored, we have that $u_1 \neq r$ and $S(u_1)$ is uncolored, which implies $S(u_1) \subseteq (W - A) \cup \{u^*, r\}$. Let u_2 be any vertex in $S(u_1)$. So, $u_1, u_2 \in W$ and thus u_2u_1 is a path of length 1 in X with in-endvertex u_1 and out-endvertex u_2. As u_2 is uncolored and Rule 2 does not apply we have that $S(u_2)$ is uncolored. Continuing this way we obtain a sequence $u_1, u_2, \ldots,$ of uncolored vertices such that for each $i = 1, 2, \ldots,$ $u_{i+1} \in S(u_i)$ and $u_{i+1}u_i$ is a path of length 1 in X with in-endvertex u_i and out-endvertex u_{i+1}. As G is finite, we have $u_i = u_k$ for some $i < k$, which contradicts Lemma 5(i), stating that X_j is acyclic. This shows that if none of Rule 1, Rule 2, Rule 3 applies, then all vertices in H' are colored. So, applying Rule 1–Rule 3 as long as possible colors all the remaining vertices of G.

Next, we shall show that the produced coloring is good, or more specifically that each monochromatic components of H' is a subpath of the union of two paths from X_j. Rule 1 and Rule 3 ensure that every vertex $v \in W - A$ has at most two neighbors in the same color as v. If u^* exists, then $\deg(u^*) = 3$, and hence Rule 2 ensures that $e^* = ru^*$ is colored properly. Moreover, for every vertex $u \in W$ let $X(u)$ be the set of paths P in X_j containing u, for which the neighbor of u in P has the same color as u. Then Rule 1 and Rule 2 ensure that $X(u) = \emptyset$, or $X(u)$ consists of exactly one path with in-endvertex u, or $X(u)$ consists of at most two paths, both with out-endvertex u.

Recall that we colored the vertices in A so that no vertex in X_j has a neighbor outside of X_j in the same color. Moreover, we colored $V(E_1) \cup B \cup \{u^*\}$ in such a way that every monochromatic component of $X_j - W$ is completely contained in a path of X_j. Finally, we colored the vertices in W so that every monochromatic component of X_j is the union of at most two monochromatic components of $X_j - W$. Together this implies that every

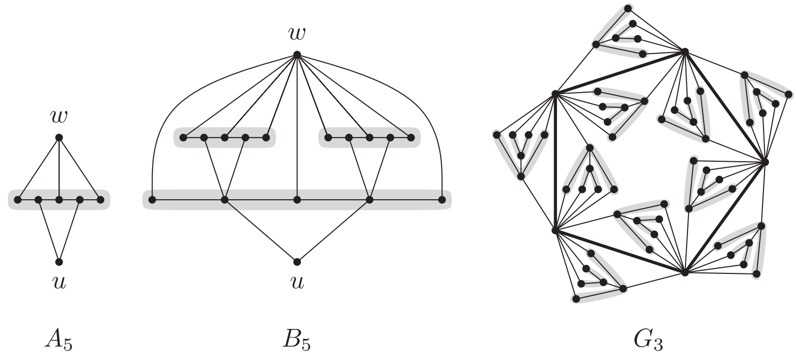

FIGURE 6. The graph A_5, B_5 and G_3.

monochromatic component is contained in the union of at most two paths in X_j. To summarize, we see that our coloring is good on $V(G) - V_1$. Now, each path of X_j is facial, that is, has at most eight edges by Lemma 2 and each monochromatic component in V_1 is a path contained in the union of some two paths from X_j. This monochromatic path has length at most 14, because it is induced and hence contains at most seven edges from each of the two paths. So, our coloring is good on V_1. Finally, since no vertex of V_1 has the same color as its neighbor (if exists) in $V(G) - V_1$, the vertices of each monochromatic component are completely contained in V_1 or in $V(G) - V_1$. Thus the coloring is good. This concludes the proof of Lemma 6 saying that X_j, $j = 1, 2, 3, 4$, is reducible.

To conclude the proof of Theorem 1, we see that Lemma 6 shows that G has a reducible configuration, contradicting the fact that G is a minimal counterexample. ■

4. PROOF OF THEOREM 2

For every integer $t \geq 2$ we define two planar graphs of girth 4, denoted by A_t and B_t, respectively. The graph A_t consists of a path P_t on t vertices and two special vertices u and w, such that the vertices along P_t are joined by an edge alternatingly to u and w. For example, A_2 is a path on four vertices and the left of Figure 6 shows A_5. The graph B_t consists of A_t with special vertices u and w, and for every neighbor v of u there is another copy of A_t, with special vertices being identified with v and w, respectively. See the middle of Figure 6.

Note that for every $t \geq 2$ the graph B_t has girth 4 and the two special vertices u and w are at distance 3 (counted by the number of edges) in B_t.

We construct G_t inductively. For $t = 2$, we define G_t to be the 5-cycle. Clearly, in any 2-vertex coloring of G_2 there is a monochromatic P_2.

For $t \geq 3$, let G be a copy of G_{t-1}. We obtain G_t from G by considering every edge xy in G, taking two copies B, B' of B_t with special vertices u, w and u', w', respectively, and identifying x, u' and w, as well as y, u and w'. Note that G_t has girth 4 and is indeed planar: We can embed B and B' on different "sides" of the edge xy, as in the right of Figure 6.

Now, fix any 2-vertex coloring of G_t and consider the inherited coloring of G_{t-1}. By induction hypothesis there is a monochromatic copy Q of P_{t-1} in G_{t-1}, say in color 1. Let x be an endpoint of Q and y be the neighbor of x in Q. Consider the copy B of H_t where

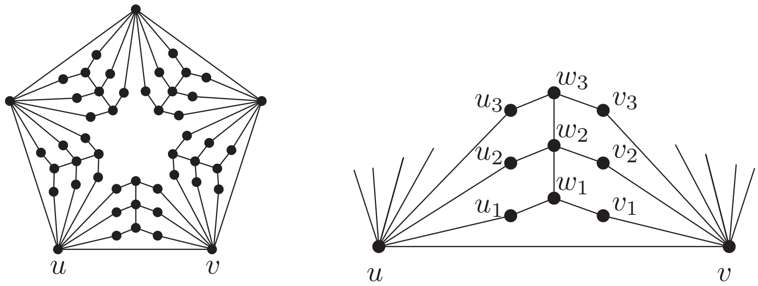

FIGURE 7. Every 2-coloring contains a monochromatic path on 3 vertices.

x is identified with w and y is identified with u and the copy A of A_t in B_t with special vertices $x = w$ and $u = y$.

If the copy P of P_t in A is not monochromatic, then at least one vertex v of P has color 1. If v is a neighbor of $x = w$, we have can extend Q by v into monochromatic P_t in color 1. Otherwise, v is a neighbor of $u = y$ and we consider the copy A' of A_t with special vertices $u' = v$ and $w' = x$. Again, if the copy P' of P_t in A' is not monochromatic, then at least one vertex v' of P' has color 1. If v' is a neighbor of x, then $Q \cup v'$ is a monochromatic P_t in color 1. Otherwise v' is a neighbor of v, and $Q \cup \{v, v'\}$ forms a monochromatic P_t in color 1. ■

5. CONCLUSIONS AND OPEN QUESTIONS

In this article, we proved that for any planar graph of girth 6 and any assignment of lists of two colors to each vertex, there is a coloring from these lists such that monochromatic components are paths of lengths at most 14. This extends a corresponding recent result of Borodin et al. [9] for planar graphs of girth 7. Our result can be interpreted as a statement about linear arboricity with short paths.

The proof uses discharging and reducible configurations. Compared to most of the previous discharging proofs, where the reducible configurations are small, here, the reducible configuration can be arbitrarily large. A similar approach was used by Havet and Sereni [21], who argued that every graph of maximum average degree less than 3 (which includes planar graphs of girth at least 6) has a 2-defective 2-list-coloring. The main difference between this proof and the proof of Theorem 1 is that Havet and Sereni can assume that in a minimal counterexample every edge is incident to a vertex of degree at least 4. Indeed, if there are two adjacent vertices u, v of degree at most 3 each, then a 2-defective 2-list-coloring of $G - \{u, v\}$ can easily be extended to a 2-defective 2-list-coloring of G. Such a reduction does not work in our case, since we can not bound the length of a longest monochromatic path. The reducible configurations of Havet and Sereni are not only simpler (they contain no vertices of degree 3), with their coloring of such a configuration one can get arbitrarily long monochromatic paths. Thus, our Lemma 6 requires less and proves more then the corresponding statement of Havet and Sereni [21, Lemma 2].

According to Table I the remaining open questions concern 2-colorings of planar graphs of girth 5 or 6. Figure 7 shows a planar graph of girth 5 that contains a monochromatic P_3 in every 2-coloring, that is, $k_d(5, 2) \geq 2$, $k_f(5, 2) \geq 3$, and $k_p(5, 2) \geq 4$. Indeed, we may assume without loss of generality, that in a given 2-coloring vertices u and v both have color 1. Then u_i and v_i, for $i = 1, 2, 3$, have color 2 or there is a monochromatic P_3 in color 1. Similarly, w_1, w_2, w_3 have color 1 or there is a monochromatic P_3 in color 2. But then these three vertices form a monochromatic P_3 in color 1.

This also follows from Montassier and Ochem [25] who provided an example of a planar graph of girth 5 such that in any red/blue coloring of its vertices there is a red P_3 or a vertex of degree at least 4 in the subgraph induced by blue vertices.

However, to the best of our knowledge, it is open whether $k_f(5, 2)$ and $k_p(5, 2)$ are finite. On the other hand, it is still possible that every planar graph of girth 5 and 6 admits a 2-coloring where every monochromatic component is a subgraph of P_3 and P_2, respectively.

ACKNOWLEDGMENTS

The authors thank the referees for the kind words and careful reading.

REFERENCES

[1] J. Akiyama, H. Era, S. V. Gervacio, and M. Watanabe, Path chromatic numbers of graphs, J Graph Theory 13(5) (1989), 569–575.

[2] N. Alon, G. Ding, B. Oporowski, and D. Vertigan, Partitioning into graphs with only small components, J Combin Theory Ser B 87(2) (2003), 231–243.

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

[4] K. Appel, W. Haken, and J. Koch, Every planar map is four colorable, Part II: Reducibility, Illinois J Math 21(3) (1977), 491–567.

[5] R. Berke and T. Szabó, Relaxed two-coloring of cubic graphs, J Combin Theory Ser B 97(4) (2007), 652–668.

[6] R. Berke and T. Szabó, Deciding relaxed two-colourability: A hardness jump, Comb Probabi Comput 18 (2009), 53–81.

[7] K. A. Berman and J. L. Paul, A 4-color theorem for surfaces of genus g, Proc Am Math Soc 105(2) (1989), 513–522.

[8] O. V. Borodin and A. O. Ivanova, List strong linear 2-arboricity of sparse graphs, J Graph Theory 67(2) (2011), 83–90.

[9] O. V. Borodin, A. Kostochka, and M. Yancey, On 1-improper 2-coloring of sparse graphs, Discrete Math 313(22) (2013), 2638–2649.

[10] G. G. Chappell, J. Gimbel, and C. Hartman, Thresholds for path colorings of planar graphs, In: Topics in Discrete Mathematics, Volume 26 of Algorithms and Combinatorics, Springer, Berlin, 2006, pp. 435–454.

[11] G. Chartrand, D. P. Geller, and S. Hedetniemi, A generalization of the chromatic number, Math Proc Cambridge Phil Soc 64 (1968), 265–271.

[12] G. Chartrand, D. P. Geller, and S. Hedetniemi, Graphs with forbidden subgraphs, J Combin Theory Ser B 10(1) (1971), 12–41.

[13] L. J. Cowen, R. H. Cowen, and D. R. Woodall, Defective colorings of graphs in surfaces: Partitions into subgraphs of bounded valency, J Graph Theory 10(2) (1986), 187–195.
[14] L. J. Cowen, W. Goddard, and C. E. Jesurum, Coloring with defect, In: Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '97, ACM, New York, NY, USA, 1997, pp. 548–557.
[15] R. Diestel, Graph Theory, 4th edn., Graduate Texts in Mathematics, Volume 173, Springer-Verlag, Heidelberg, 2010.
[16] N. Eaton and T. Hull, Defective list colorings of planar graphs, Bull Inst Combin Appl 25 (1999), 79–87.
[17] L. Esperet and G. Joret, Colouring planar graphs with three colours and no large monochromatic components, Comb Probab Comp 23 (2014), 551–570.
[18] L. Esperet and P. Ochem, Islands in graphs on surfaces, ArXiv e-prints, SIAM Journal on Discrete Mathematics, 30(1) (2016), 206–219.
[19] A. N. Glebov and D. Z. Zambalaeva, Partition of a planar graph with girth 6 into two forests with chain length at most 4, J Appl Indus Math 8(3) (2014), 317–328.
[20] W. Goddard, Acyclic colorings of planar graphs, Discrete Math 91(1) (1991), 91–94.
[21] F. Havet and J.-S. Sereni, Improper choosability of graphs and maximum average degree, J Graph Theory 52(3) (2006), 181–199.
[22] P. Haxell, T. Szabó, and G. Tardos, Bounded size components–partitions and transversals, J Combin Theory Ser B 88(2) (2003), 281–297.
[23] J. Kleinberg, R. Motwani, P. Raghavan, and S. Venkatasubramanian, Storage management for evolving databases, In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, FOCS '13, IEEE Computer Society, Washington, DC, USA, 1997, pp. 353–353.
[24] N. Linial, J. Matoušek, O. Sheffet, and G. Tardos, Graph colouring with no large monochromatic components, Comb Probab Comput 17 (2008), 577–589.
[25] M. Montassier and P. Ochem, Near-colorings: Non-colorable graphs and NP-completeness, Electron J Comb 22 (2015).
[26] K. S. Poh, On the linear vertex-arboricity of a planar graph, J Graph Theory 14(1) (1990), 73–75.
[27] R. Škrekovski, List improper colourings of planar graphs, Comb Probab Comput 8(3) (1999), 293–299.
[28] P. Weiner, Improper colourings of graphs, Bachelor's thesis, Karlsruhe Institute of Technology, October 2014.
[29] R. G. Wood and D. R. Woodall, Defective choosability of graphs without small minors, Electron J Comb 16(1) (2009), R92.

MAKING TRIANGLES COLORFUL*

Jean Cardinal,[†] *Kolja Knauer,*[‡] *Piotr Micek,*[§] *and Torsten Ueckerdt*[¶]

Abstract. We prove that for any finite point set P in the plane, a triangle T, and a positive integer k, there exists a coloring of P with k colors such that any homothetic copy of T containing at least $144k^8$ points of P contains at least one of each color. This is the first polynomial bound for range spaces induced by homothetic polygons. The only previously known bound for this problem applies to the more general case of octants in $\mathbb{R}^3$, but is doubly exponential.

1 Introduction

Covering and packing problems are ubiquitous in discrete geometry. In this context, the notion of ϵ-nets captures the idea of finding a small but representative sample of a data set (see for instance Chapter 10 in Matoušek's lectures [14]). Given a set system, or *range space*, on n elements, an ϵ-net for this system is a subset of the elements such that any set, or *range*, containing at least ϵn elements contains at least one element of the subset.

In this paper, we are interested in *coloring* the elements so that any range containing *sufficiently many* elements contains at least *one element of each color*. Hence instead of finding a single subset of representative elements, we wish to partition the elements into representative classes.

For a given class of range spaces, we define the function $p(k)$ as the minimum number p such that the following holds: the elements of every range space in that class can be colored with k colors so that any range containing at least p elements contains at least one of each color. It is not difficult to show that if $p(k) = O(k)$ for a class of range spaces, then this class admits ϵ-nets of size $O(1/\epsilon)$.

We are interested in range spaces defined by a collection $\mathcal{B}$ of subsets of $\mathbb{R}^d$. In what follows, we are mainly concerned with the case where $\mathcal{B}$ is a collection of *convex bodies*, that is, compact convex subsets of $\mathbb{R}^d$. For given $\mathcal{B}$ we obtain a range space whose ground set is a (countable or finite) point set $P \subseteq \mathbb{R}^d$ by considering all subsets of P formed by

*Research supported by the ESF EUROCORES programme EuroGIGA, CRP ComPoSe (http://www.eurogiga-compose.eu), and GraDR (http://kam.mff.cuni.cz/gradr/). Jean Cardinal is supported by the F.R.S-FNRS as grant convention no. R 70.01.11F within the ComPoSe project. Kolja Knauer is supported by DFG grant FE-340/8-1 as part of the project GraDR. Piotr Micek is supported by the Ministry of Science and Higher Education of Poland as grant no. 884/N-ESF-EuroGIGA/10/2011/0 within the GraDR project.

[†] *Université libre de Bruxelles, Belgium*, jcardin@ulb.ac.be
[‡] *Technische Universität Berlin, Germany*, knauer@math.tu-berlin.de
[§] *Jagiellonian University, Krakow, Poland*, Piotr.Micek@tcs.uj.edu.pl
[¶] *Karlsruhe Institute of Technology, Germany*, Torsten.Ueckerdt@kit.edu

intersecting P with a member of $\mathcal{B}$. This construction yields so-called *primal range spaces induced by* $\mathcal{B}$. For instance, if P is a set of points in the plane and $\mathcal{B}$ the set of all disks, then the ranges are all possible intersections of P with a disk. Such range spaces and their ϵ-nets appear frequently in discrete geometry and in applications such as sensor networks [8].

One can also consider *dual range spaces induced by* $\mathcal{B}$, where the ground set is a (countable or finite) subcollection $\mathcal{B}'$ of $\mathcal{B}$, and the ranges are all subsets X of $\mathcal{B}'$ such that there exists some $p \in \mathbb{R}^d$ with $X = \{B \in \mathcal{B}' \mid p \in B\}$. For instance, if $\mathcal{B}'$ is a set of disks in the plane, then the ranges are all maximal sets of disks containing a common point.

In general, those are also referred to as (primal and dual) *geometric hypergraphs.*

In the case of dual range spaces induced by a collection $\mathcal{B}$ of objects, the problem of bounding $p(k)$ is known as the *covering decomposition problem of* $\mathcal{B}$. In this setting, we are given a subcollection of these objects, and we wish to partition them into k color classes, so that whenever a point is contained in sufficiently many objects of the initial collection, it is contained in at least one object of each class.

We prove a polynomial upper bound on $p(k)$ for primal range spaces induced by homothetic triangles in the plane.

1.1 Previous Work

These questions were first studied by János Pach in the early eighties [15]. An account of early related results and conjectures can be found in Chapter 2 of the survey on open problems in discrete geometry by Brass, Moser, and Pach [4].

In the past five years, tremendous progress has been made in this area, for range spaces induced by various families of convex bodies. One of the most striking achievements is the recent proof that $p(k) = O(k)$ for translates of convex polygons, the culmination of a series of intermediate results for various special cases. We remark that convex bodies are considered because $p(k) = \infty$ for range spaces induced by translates of concave polygons [16]. We refer the reader to Table 1 for a summary of the known bounds.

The specific case of translates of a triangle with $k = 2$ was tackled by Tardos and Tóth in 2007 [21]. They proved that every point set can be colored red and blue so that every translate of a given triangle containing at least 43 points contains at least one red and one blue. We generalize this result in two ways: we consider *homothetic* triangles, and an arbitrary number of colors.

The only previously known results applying to our problem are due to Keszegh and Pálvölgyi [11, 12]. They actually apply to the more general case of translates of (say) the positive octant in a cartesian representation of $\mathbb{R}^3$. The special case of triangles homothetic to the triangle with vertices $(0, 0)$, $(1, 0)$ and $(0, 1)$ occurs when all points lie on a plane orthogonal to the vector $(1, 1, 1)$. The bound that was proven for arbitrary k is of the order of 12^{2^k}, and is most probably far from being tight.

Range spaces	primal	dual
halfplanes	$p(k) = 2k - 1$ [2, 10, 20]	$p(2) = 3$ [7] $p(k) \leqslant 3k - 2$ [2, 20]
translates of a convex polygon	$p(k) = O(k)$ [21, 19, 17, 1, 9]	
translates of an octant in $\mathbb{R}^3$	$p(2) \leqslant 12$ [11] $p(k) \leqslant 12^{2^k}$ [12]	
unit disks	∞ [18]	
bottomless rectangles	$p(2) = 4$ [10] $1.6k \leqslant p(k) \leqslant 3k - 2$ [3]	$p(2) = 3$ [10] $p(k) \leqslant 12^{2^k}$ [12] (from octants in $\mathbb{R}^3$)
axis-aligned rectangles	∞ [6]	∞ [16]
disks and halfspaces in $\mathbb{R}^3$	∞ [16]	∞ [18]

Table 1: Known results for various families of range spaces. For range spaces induced by translates of a set, the primal problem is the same as the dual. When more than one reference is given, they correspond to successive improvements, but only the best known bound is indicated. The symbol ∞ indicates that $p(k)$ does not exist.

1.2 Our Result

Theorem 1.1. *Given a finite point set $P \subseteq \mathbb{R}^2$, a triangle $T \subseteq \mathbb{R}^2$ and a positive integer k, there exists a coloring of P with k colors such that any homothetic copy of T containing at least $144 \cdot k^8$ points of P contains at least one of each color.*

The proof is elementary, and builds on the previous work by Keszegh and Pálvölgyi [11, 12]. The degree of the polynomial depends on $p(2)$. Hence any improvement on $p(2)$ would yield a polynomial improvement in the bound. For the same reason, it can be shown that the same coloring method cannot be used to prove any upper bound better than $O(k^4)$ (as $p(2) \geqslant 4$).

2 Proof

Let $\mathcal{B}$ be the collection of all homothetic copies of a fixed closed triangle T in the plane. We consider the class of primal range spaces induced by $\mathcal{B}$. From now on we denote by $p(k)$ the minimum p such that every finite set of points in the plane can be colored with k colors so that any homothetic copy of T containing at least p points contains at least one point of each color.

Lemma 2.1. *If $p(2) \leqslant c$, for some constant c, then $p(2k) \leqslant c^2 p(k)$, for all $k \geqslant 2$.*

Proof. It suffices to prove the lemma for any fixed triangle T and then argue for all others using an affine transformation of the plane. Let T be the triangle with vertices $(0, 0)$, $(1, 0)$ and $(0, 1)$.

Consider a finite point set P and a k-coloring $\phi : P \to \{1, \dots, k\}$ such that any homothetic copy of T containing at least $p(k)$ points contains one of each color. Note that $p(k) < \infty$ [12]. We suppose without loss of generality that no two points of P lie on a line of slope -1, otherwise we can slightly perturb the points, and a suitable coloring for the perturbed version will also work for P.

We now describe a simple procedure to double the number of colors. For $1 \leqslant i \leqslant k$ let $P_i = \phi^{-1}(i)$ that is the set of points with color i. Provided $p(2) \leqslant c$ there is a 2-coloring $\phi_i : P_i \to \{i', i''\}$ of P_i such that for any homothetic copy T' of T containing at least c points of P_i, T' contains at least one point of each color. We define ϕ' to be the disjoint union of all ϕ_i, and claim that ϕ' is a $2k$-coloring of P such that for any homothetic copy T' of T containing at least $c^2 p(k)$ points, T' contains at least one point of each of the $2k$ colors.

Consider a homothetic copy T' of a triangle T containing at least $c^2 p(k)$ points from P, and in order to get a contradiction suppose that one of the $2k$ colors used by ϕ' is missing in T'. Let i' be this color. Note that if there are at least c points in T' with color i then i' and i'' must be present in T', from the correctness of the 2-coloring ϕ_i. Hence we conclude that there are less than c points in T' with color i.

Order the points in $T' \cap P = \{p_1, p_2, \dots\}$ in such a way that the sum of their x- and y-coordinates is non-decreasing. Hence the order corresponds to a sweep of the points in $T' \cap P$ by a line of slope -1. By the pigeonhole principle, since there are less than c points colored with i, there must exist a subsequence $Q = (p_j, p_{j+1}, \dots, p_{j+\ell-1})$ of points of color distinct from i, of length $\ell := c^2 p(k)/c = cp(k)$.

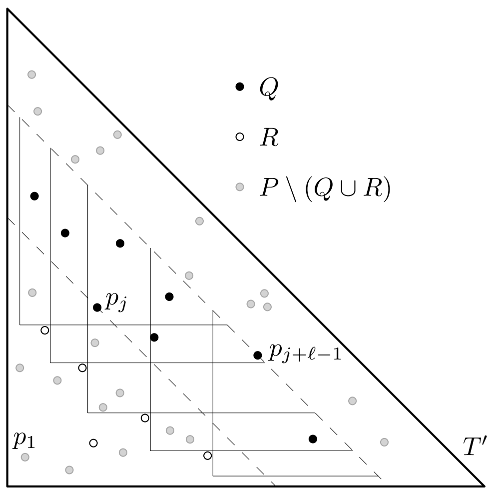

Figure 1: Illustration of the proof of Lemma 2.1.

Let $R := P_i \cap \{p_1, p_2, \dots, p_{j-1}\}$ be the set of points of color i that come before Q in the sweep order. By assumption, we have $|R| < c$. Hence the points of Q can be covered with c translates of the first quadrant, such that none of them intersects R; see Figure 1. For example, it is enough to consider all inclusion-wise maximal quadrants with apex in T' that avoid points in R. Applying the pigeonhole principle a second time, one of these quadrants must contain at least $|Q|/c = cp(k)/c = p(k)$ points, none of which is colored i.

This quadrant, together with the sweepline containing the last point $p_{j+\ell-1}$ of Q, forms a triangle that is homothetic to T, contains at least $p(k)$ points, none of which has color i. This is a contradiction with the correctness of the initial k-coloring ϕ. □

Proof of Theorem 1.1. It was shown by Keszegh and Pálvölgyi that $p(2) \leqslant 12$ [11]. Hence it remains to solve the recurrence of the previous lemma with $c = 12$. We look for an upper bound on $p(k)$ satisfying $p(2k) \leqslant 144 \cdot p(k)$ for any positive integer k, and $p(2) \leqslant 12$. This yields $p(2^i) \leqslant 144^i$ for any positive integer i, and $p(k) \leqslant 144^{\lceil \log_2 k \rceil} < 144 \cdot k^8$ for any positive integer k. □

3 Open Problems

The only lower bounds on $p(k)$ the authors are aware of is the bound $p(k) \geqslant 1.6k$ for bottomless rectangles [3] (which improves the bound $p(k) \geqslant 4k/3$ for translates of squares [17]) and the tight bound $p(k) \geqslant 2k - 1$ for halfplanes [20].

No bound at all is known for the primal range space induced by axis-aligned squares: does there exist a function $p(k)$ such that for any point set P there is a k-coloring of P such that any axis-aligned square containing at least $p(k)$ points of P contains at least one point of each color?

We remark that after this paper has been submitted the bound of $144k^8$ was improved to $O(k^6)$ even in the more general setting of translates of octants in $\mathbb{R}^3$ [5], and also by Keszegh and Pálvölgyi [13] to $O(k^{4.58})$ again only in the case of homothetic triangles. Both results rely on the same idea as the one in Lemma 2.1, namely defining a $2k$-coloring from a k-coloring by splitting each color class into two.

4 Acknowledgments

This work was carried out during a visit of Kolja, Piotr, and Torsten in Brussels, funded by a EUROCORES Short Term Visit grant, in the framework of the EuroGIGA programme. The authors wish to thank Stefan Langerman for insightful discussions, and the referees for their careful reading and useful comments.

References

[1] Greg Aloupis, Jean Cardinal, Sébastien Collette, Stefan Langerman, David Orden, and Pedro Ramos. Decomposition of multiple coverings into more parts. *Discrete & Computational Geometry*, 44(3):706–723, 2010.

[2] Greg Aloupis, Jean Cardinal, Sébastien Collette, Stefan Langerman, and Shakhar Smorodinsky. Coloring geometric range spaces. *Discrete & Computational Geometry*, 41(2):348–362, 2009.

[3] Andrei Asinowski, Jean Cardinal, Nathann Cohen, Sébastien Collette, Thomas Hackl, Michael Hoffmann, Kolja Knauer, Stefan Langerman, Michał Lasoń, Piotr Micek, Günter Rote, and Torsten Ueckerdt. Coloring hypergraphs induced by dynamic point sets and bottomless rectangles. In *Proc. Workshop on Data Structures and Algorithms (WADS)*, pages 73–84, 2013.

[4] Peter Brass, William Moser, and János Pach. *Research Problems in Discrete Geometry.* Springer, 2005.

[5] Jean Cardinal, Kolja Knauer, Piotr Micek, and Torsten Ueckerdt. Making octants colorful and related covering decomposition problems. In *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014. To appear.

[6] Xiaomin Chen, János Pach, Mario Szegedy, and Gábor Tardos. Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles. *Random Struct. Algorithms*, 34(1):11–23, 2009.

[7] Radoslav Fulek. Coloring geometric hypergraph defined by an arrangement of half-planes. In *Proc. Canadian Conference on Computational Geometry (CCCG)*, pages 71–74, 2010.

[8] Matt Gibson and Kasturi Varadarajan. Decomposing coverings and the planar sensor cover problem. In *50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 159–168, 2009.

[9] Matt Gibson and Kasturi Varadarajan. Optimally decomposing coverings with translates of a convex polygon. *Discrete & Computational Geometry*, 46(2):313–333, 2011.

[10] Balázs Keszegh. Coloring half-planes and bottomless rectangles. *Comput. Geom.*, 45(9):495–507, 2012.

[11] Balázs Keszegh and Dömötör Pálvölgyi. Octants are cover-decomposable. *Discrete & Computational Geometry*, 47(3):598–609, 2012.

[12] Balázs Keszegh and Dömötör Pálvölgyi. Octants are cover-decomposable into many coverings. *arXiv preprint arXiv:1207.0672*, 2012.

[13] Balázs Keszegh and Dömötör Pálvölgyi. Convex polygons are self-coverable. *arXiv preprint arXiv:1307.2411*, 2013.

[14] Jirí Matoušek. *Lectures on Discrete Geometry.* Springer, 2002.

[15] János Pach. Decomposition of multiple packing and covering. In *2. Kolloq. über Diskrete Geom.*, pages 169–178. Inst. Math. Univ. Salzburg, 1980.

[16] János Pach, Gábor Tardos, and Géza Tóth. Indecomposable coverings. In *Proc. 7th China-Japan Conference on Discrete Geometry, Combinatorics and Graph Theory (CJCDGCGT)*, pages 135–148, 2005.

[17] János Pach and Géza Tóth. Decomposition of multiple coverings into many parts. *Comput. Geom.*, 42(2):127–133, 2009.

[18] Dömötör Pálvölgyi. Indecomposable coverings with unit discs. *arXiv preprint arXiv:1310.6900*, 2013.

[19] Dömötör Pálvölgyi and Géza Tóth. Convex polygons are cover-decomposable. *Discrete & Computational Geometry*, 43(3):483–496, 2010.

[20] Shakhar Smorodinsky and Yelena Yuditsky. Polychromatic coloring for half-planes. In *Proc. Scandinavian Worskhop on Algorithms Theory (SWAT)*, pages 118–126, 2010.

[21] Gábor Tardos and Géza Tóth. Multiple coverings of the plane with triangles. *Discrete & Computational Geometry*, 38(2):443–450, 2007.

SIAM J. DISCRETE MATH.
Vol. 28, No. 4, pp. 1948–1959

MAKING OCTANTS COLORFUL AND RELATED COVERING DECOMPOSITION PROBLEMS*

JEAN CARDINAL†, KOLJA KNAUER‡, PIOTR MICEK§, AND TORSTEN UECKERDT¶

Abstract. We give new positive results on the long-standing open problem of geometric covering decomposition for homothetic polygons. In particular, we prove that for any positive integer k, every finite set of points in $\mathbb{R}^3$ can be colored with k colors so that every translate of the negative octant containing at least k^6 points contains at least one of each color. The best previously known bound was doubly exponential in k. This yields, among other corollaries, the first polynomial bound for the decomposability of multiple coverings by homothetic triangles. We also investigate related decomposition problems involving intervals appearing on a line. We prove that no algorithm can dynamically maintain a decomposition of a multiple covering by intervals under insertion of new intervals, even in a *semionline* model, in which some coloring decisions can be delayed. This implies that a wide range of sweeping plane algorithms cannot guarantee any bound even for special cases of the octant problem.

Key words. hypergraph coloring, geometric hypergraphs, covering decomposition, online algorithms

AMS subject classifications. 52C15, 05C15, 68W27

DOI. 10.1137/140955975

1. Introduction and main results.

We study coloring problems for hypergraphs induced by simple geometric objects. Given a family of convex bodies in $\mathbb{R}^d$, a natural colorability question that one may consider is the following: is it true that for any positive integer k, every collection of points $\mathcal{P} \subset \mathbb{R}^d$ can be colored with k colors so that any element of the family containing at least $p(k)$ of them, for some function $p(k)$, contains at least one of each color? This question has been investigated previously for convex bodies in the plane such as halfplanes and translates of a convex polygon.

Octants in three-space. In this paper, we give a polynomial upper bound on $p(k)$ when the family under consideration is the set of translates of the three-dimensional negative octant $\{(x, y, z) \in \mathbb{R}^3 : x \leqslant 0, y \leqslant 0, z \leqslant 0\}$. The best previously

*Received by the editors February 6, 2014; accepted for publication (in revised form) August 6, 2014; published electronically November 18, 2014. This research was supported by the ESF EUROCORES programme EuroGIGA, CRP ComPoSe, and GraDR. Parts of this research was carried out during a visit of Kolja, Piotr, and Torsten to ULB in Brussels, funded by a EUROCORES Short Term Visit grant, and during a visit of Jean at the Jagiellonian University in Kraków.
http://www.siam.org/journals/sidma/28-4/95597.html

†Université libre de Bruxelles, 1050 Brussels, Belgium (jcardin@ulb.ac.be). This author's research was supported by the F.R.S-FNRS as grant convention R 70.01.11F within the ComPoSe project.

‡LIRMM, Université Montpellier 2, 34095 Montpellier, France (kolja.knauer@gmail.com). This author's research was supported by the ANR TEOMATRO grant ANR-10-BLAN 0207 and the EGOS grant ANR-12-JS02-002-01.

§Theoretical Computer Science Department, Faculty of Mathematics and Computer Science, Jagiellonian University, PL30-348 Kraków, Poland (piotr.micek@tcs.uj.edu.pl). This author's research was supported by the Ministry of Science and Higher Education of Poland as grant 884/N-ESF-EuroGIGA/10/2011/0 within the GraDR project.

¶Department of Mathematics, Karlsruhe Institute of Technology, Karlsruhe 76128, Germany (torsten.ueckerdt@kit.edu).

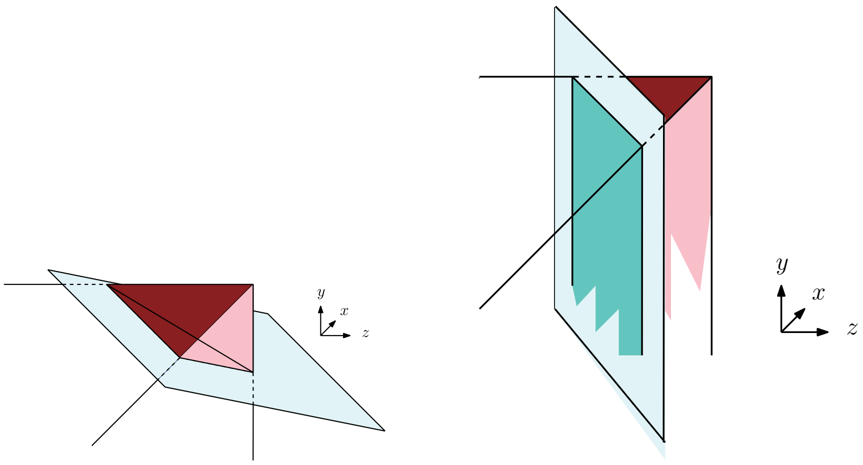

(a) From octants to triangles. (b) From octants to bottomless rectangles.

Fig. 1. *Special cases of the octant coloring problem.*

known bound is due to Keszegh and Pálvölgyi, and is doubly exponential in k [21].

Theorem 1.1. *There exists a constant $a < 6$ such that for any positive integer k, every finite set $\mathcal{P}$ of points in $\mathbb{R}^3$ can be colored with k colors so that every translate of the negative octant containing at least k^a points of $\mathcal{P}$ contains at least one of each color.*

A *dual* version of the above problem, sometimes referred to as *cover-decomposability*, can be stated as follows: Given a collection C of convex bodies, we wish to color them with k colors so that any point of $\mathbb{R}^d$ covered by at least $p(k)$ of them, for some function $p(k)$, is covered by at least one of each color. In the *primal* setting with respect to octants we can replace the point set $\mathcal{P}$ with a set C of *positive* octants with apices in $\mathcal{P}$. Then the primal value of $\mathcal{P}$ coincides with the dual value of C. Since clearly the dual problem is equivalent if we pick negative instead of positive octants, we have the following corollary.

Corollary 1.2. *There exists a constant $a < 6$ such that for any positive integer k, every finite set $\mathcal{P}$ of translates of the negative octant can be colored with k colors so that every point of $\mathbb{R}^3$ contained in at least k^a octants of $\mathcal{P}$ is contained in at least one of each color.*

The next corollary is obtained by observing that the intersections of a set of octants with a plane in $\mathbb{R}^3$ that is not parallel to any axis form a set of homothetic triangles (see Figure 1(a)).

Corollary 1.3. *There exists a constant $a < 6$ such that for any positive integer k, every finite set $\mathcal{P}$ of homothetic triangles in the plane can be colored with k colors so that every point contained in at least k^a triangles of $\mathcal{P}$ is contained in at least one of each color.*

Finally, using standard arguments, the latter result can be extended to infinite sets, and cast as a cover-decomposability statement. Here a covering is said to be *decomposable* into k coverings when the objects in the covering can be colored with k colors so that every color class is a covering by itself.

Corollary 1.4. *There exists a constant $a < 6$ such that for any positive in-*

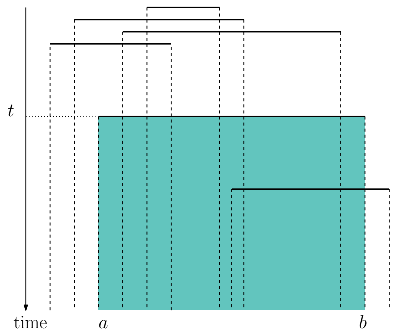

FIG. 2. *Intervals under insertion and bottomless rectangles.*

teger k, every locally finite k^a-fold covering of the plane by homothetic triangles is decomposable into k coverings.

The proof of Theorem 1.1 is given in section 3.

Intervals, bottomless rectangles, and sweeping algorithms. It is well-known that a theorem similar to Corollary 1.4 holds for the simpler case of intervals on the real line. Rado [32] observed that every k-fold covering of the real line by intervals can be decomposed into k coverings. Using this result, it is not difficult to prove that $p(k) = k$ for translates of the negative quadrant in $\mathbb{R}^2$.

In the second part of this paper, we study the problem of maintaining a decomposition of a set of intervals *under insertion*. This problem is similar in spirit, but distinct from the previous one.

We are now given a positive integer k, a collection of intervals on the real line, and for each such interval a real number representing an *insertion time*. This collection represents a set of intervals that evolves over time, in which the intervals present at time x are exactly those whose insertion time is at most x. We can now ask whether there exists a function $p(k)$ such that the following holds: there exists a k-coloring of the intervals in the collection $\mathcal{S}$ such that, *at any time*, any point that is covered by at least $p(k)$ intervals *present at that time* is covered by at least one of each color.

This can be conveniently represented in the plane by representing each interval $[a, b]$ with insertion time t as an axis-aligned rectangle with vertex coordinates $(a, -t), (b, -t), (b, -\infty), (a, -\infty)$; hence viewing time goes downward in the vertical direction. We refer to such rectangles, with a bottom edge at infinity, as *bottomless* rectangles. Now the k-coloring must be such that every point $p \in \mathbb{R}^2$ that is contained in at least $p(k)$ such rectangles must be contained in at least one of each color. Hence, the problem is actually about *decomposition of coverings by bottomless rectangles*. We illustrate this point of view in Figure 2. Also note that bottomless rectangles can be seen as *degenerate homothetic triangles*, which we will make use of for Corollary 1.7.

We now observe that bottomless rectangles can be formed by intersecting a negative octant with a vertical plane, as depicted in Figure 1(b). Hence we can formulate a new corollary of our main theorem.

COROLLARY 1.5. *There exists a constant $a < 6$ such that for any positive integer k, every finite set $\mathcal{P}$ of bottomless rectangles in the plane can be colored with k colors so that every point contained in at least k^a rectangles of $\mathcal{P}$ is contained in at least one of each color. Equivalently, every collection of intervals, each associated with an insertion time, can be k-colored so that at any time every point covered by at least k^a*

intervals present at this time is covered by at least one of each color.

With respect to the model of intervals with insertion times it is natural to ask whether it is possible to maintain a decomposition of a set of intervals under insertion *without knowing the future insertions* in advance. In section 4, we answer this question in the negative even if coloring decisions can be delayed.

More precisely, we rule out the existence of a semionline algorithm. A *semionline k-coloring algorithm* must consider the intervals in their order of insertion time. At any time, an interval in the sequence either has one of the k colors, or is left uncolored. Any interval can be colored at any time, but once an interval is assigned a color, it keeps this color forever.

A semionline k-coloring algorithm is said to be *colorful of value d* if it maintains at all times that the colors that are already assigned are such that any point contained in at least d intervals is contained in at least *one of each of the k colors*.

In order to obtain that there is no semionline colorful coloring algorithm of bounded value, we prove a stronger statement about the less restrictive *proper coloring* problem. We call a semionline k-coloring algorithm *proper of value d* if it maintains at all times that the colors that are already assigned are such that any point contained in at least d intervals is contained in at least *two of distinct colors*. O ur theorem says that for all natural numbers k, d, there is no semionline proper k-coloring algorithm of value d.

THEOREM 1.6. *For all natural numbers k, d, there is no semionline algorithm that k-colors intervals under the operation of inserting intervals, so that at any time, every point covered by at least d intervals is covered by at least two of distinct colors.*

Since any semionline colorful coloring algorithm is also proper, we obtain that there is no such algorithm of bounded value.

Note that in the bottomless rectangle model a semionline colorful coloring algorithm corresponds to sweeping the set of rectangles top to bottom with a line parallel to the x-axis and assigning colors irrevocably to already swept rectangles such that at any time every point contained in d of those already swept is contained in at least one of each color. Similarly, one can define sweeping line algorithms for coloring homothetic triangles, where the point set is swept top to bottom by a line parallel to one of the sides of the triangles. For octants a sweeping plane algorithm would sweep the point set from top to bottom with a plane parallel to the x, y-plane. Since bottomless rectangles can be viewed as a special case of both, we can summarize with the following corollary.

COROLLARY 1.7. *For all natural numbers k, d, there is no sweeping line (plane) coloring algorithm in the above sense such that for any set of bottomless rectangles, or triangles, or octants, at any time every point contained in d of the already swept ranges is contained in at least one of each color.*

Since for octants primal and dual problem are equivalent by Corollary 1.7, no such sweeping plane algorithm exists for the primal octant problem either.

We remark that Corollary 1.7 is in contrast with another recent result in [5], which deals with the primal version of the problem. It can be expressed as coloring points appearing on a line in such a way that at all times any interval containing $p(k)$ points contains one point of each color, or equivalently, coloring point sets in the plane such that every bottomless rectangle containing $p(k)$ points contains a point of each color. In [5] it is shown that in this case a linear upper bound on $p(k)$ can be achieved with a semionline coloring algorithm, or equivalently a sweeping line algorithm.

2. Previous results. The covering decomposition problem was first posed by Pach in 1980 and 1986 [24, 25]. This was originally motivated by the problem of determining the densities of the densest k-fold packings and the thinnest k-fold coverings of the plane with a given plane convex body (see section 2.1 in [8] for a complete historical account). In particular, he posed the following problem.

Is it true that for any plane convex polygon C and for any integer k, there exists an integer $p = p(C, k)$ such that every p-fold covering of the plane with homothetic copies of C can be decomposed into k coverings?

Our contribution shows that $p(C, k) = O(k^6)$, provided that C is a triangle and the covering is locally finite (Corollary 1.4).

Tremendous progress has been made recently in understanding the conditions for the existence of a function $p(k)$ for a given range space, that is, geometric hypergraphs induced by a family of bodies in $\mathbb{R}^d$. To the best of our knowledge, our result is the first polynomial bound for cover-decomposability of *homothetic* copies of a polygon. Linear upper bounds have been obtained for halfplanes [4, 34] and translates of a convex polygon in the plane [35, 29, 3, 16]. A restricted version of this problem involving unit balls is shown to be solvable using the probabilistic method in the eponymous book from Alon and Spencer [2]. The function $p(k)$ has been proved not to exist for range spaces induced by concave polygons [30], axis-aligned rectangles [12, 26], lines in $\mathbb{R}^2$, and disks [28]. In a remarkable recent preprint, Pálvölgyi proved the nonexistence of a function $p(k)$ even for *unit* disks [31], thereby invalidating earlier claims from Pach and Mani-Levitska in an unpublished manuscript. Note that the indecomposability results for axis-aligned rectangles imply the same for orthants in $\mathbb{R}^4$, since arbitrary such rectangles can be formed by intersecting four-dimensional orthants with a plane in $\mathbb{R}^4$.

Finally, in another recent preprint, Kovács [22] proved that there exists indecomposable coverings by homothets of any polygon with at least four edges, disproving the general conjecture above. Overall, this collection of results essentially closes most of Pach's questions on cover-decomposability of plane convex bodies.

Previous results on octants. Pálvölgyi proved the indecomposability of coverings by translates of a convex polyhedron in $\mathbb{R}^3$ [30]. His proof, however, does not hold for unbounded polyhedra with three facets. This prompted the first author of the current paper to pose the problem of decomposability of coverings by octants. This was solved by Keszegh and Pálvölgyi, who showed that $p(2) \leqslant 12$ in this case [19]. Since we will reuse this theorem in our proof, it is worth reproducing it here.

THEOREM 2.1 (see [19]). *There exists a constant $c \leqslant 12$ such that every finite collection $\mathcal{P} \subset \mathbb{R}^3$ of points can be 2-colored so that every negative octant containing at least c points of $\mathcal{P}$ contains at least one of each color.*

In the past two years, the above result was improved and generalized. First, Keszegh and Pálvölgyi proved that Theorem 2.1 implies that $p(k)$ is bounded for every k [21]. Note that this is not obvious, as one could well imagine that for some range spaces, $p(2)$ is bounded, but not $p(k)$ for some $k > 2$. Their upper bound on $p(k)$, however, is doubly exponential in k. In particular, their proof implies $p(k) \leqslant 12^{2^k}$.

Later, the current authors gave a polynomial upper bound on $p(k)$, but restricted themselves to the special case of homothetic triangles in the plane, where points are to be colored [9]. The proof uses a new technique involving recoloring each color class of a k-coloring with two colors in order to obtain a $2k$-coloring.

Finally, in May 2013, a manuscript from Keszegh and Pálvölgyi was communicated to us by Pach, in which an improved polynomial upper bound was given for the same special case of homothetic triangles [20]. This improvement makes use of a lemma stating the so-called *self-coverability* property of triangles.

We managed to harness the power of these observations for the general case of octants. In particular, we reuse the recoloring algorithm given in [9] in Lemma 3.2, and also give a three-dimensional generalization of the self-coverability lemma of [20] in the form of Lemma 3.1. Our proof of Theorem 1.1 is longer than that of the doubly exponential upper bound in [21], but not significantly more involved.

Previous results on online coloring problems and proper colorings of geometric hypergraphs. Semionline algorithms have proved to be useful in an interesting special case of the problem with octants, in which all points considered in Theorem 1.1 lie on a vertical plane. This setting can be thought of as points appearing on a line, and we want to color the points with k colors such that at any time, any set of $p(k)$ consecutive points contains at least one of each color. This problem has been studied by a number of authors, whose results were compiled in a joint paper [5]. In particular, they showed that under this restriction, we have $1.6k \leqslant p(k) \leqslant 3k-2$. The upper bound is achieved using a semionline algorithm, which does not require the knowledge of the future point insertions, and never recolors a point. This also amounts to coloring *primal* range spaces induced by bottomless rectangles with a sweeping line algorithm, i.e., coloring points such that bottomless rectangles containing many of them contain all colors.

In contrast to our negative result about semionline algorithms, a larger class of algorithms called *quasi-online* has led to a short proof that $p(2) = 3$ in the setting corresponding to our Corollary 1.5 (see [18]) and is indeed also used to obtain Theorem 2.1 in [19].

Clearly, colorful 2-colorings and proper 2-colorings coincide, but also for a larger number of colors proper colorings of geometric hypergraphs have been considered in the primal and dual setting. There are results for bottomless rectangles [17], halfplanes [15, 17], octants [10], rectangles [12, 1, 26], and disks [28, 33].

Similarly to our Theorem 1.6, Keszegh, Lemons, and Pálvölgyi [18] consider *online* proper coloring algorithms (points must be colored on arrival). While it is easy to see that there is an optimal online algorithm to color points such that quadrants are colorful, they show that there is no online proper coloring algorithm of bounded value in the primal setting of bottomless rectangles and octants. This is implied by Theorem 1.6 and indeed the proof methods have similarities. In [18] the quality of online algorithms is then measured as a function of the input size.

In another vein, Bar-Noy et al. [7, 6] considered *conflict-free* colorings in an online setting. There, the problem is to maintain that every d-covered point p is covered by one interval whose color is unique among all intervals covering p.

Other related results. In 2010, Varadarajan gave a feasibility result for the *fractional set cover packing* problem with fat triangles (Corollary 2 in [36]). This problem can be seen as a fractional variant of the covering decomposition problem. This result involves the construction of so-called *quasi-uniform ε-nets*. This construction was recently improved by Chan et al. [11]. These results are essentially motivated by the design of improved approximation algorithms for geometric versions of the weighted set cover problem. However, they can also be seen as an intermediate step between the problem of finding small ε-nets and the covering decomposition problem, which involves partitioning a set into ε-nets (see our conclusion for a discussion on this relation).

3. Proof of Theorem 1.1. In what follows, we will use the shorthand notation $[n] = \{1, 2, \ldots, n\}$ for a positive integer n. We will refer to the three coordinates of a point p as p_x, p_y, and p_z, respectively. The *negative octant* with *apex* $(p_x, p_y, p_z) \in \mathbb{R}^3$ is the set $\{(x, y, z) \in \mathbb{R}^3 : x \leqslant p_x, y \leqslant p_y, z \leqslant p_z\}$. Similarly the *positive octant* of (p_x, p_y, p_z) is $\{(x, y, z) \in \mathbb{R}^3 : x \geqslant p_x, y \geqslant p_y, z \geqslant p_z\}$. For convenience we also allow the coordinates of an apex to be equal to ∞. In what follows, an octant will generally be considered to be negative, unless explicitly stated otherwise. For two points $p, q \in \mathbb{R}^3$, we say that p *dominates* q whenever the negative octant with apex p contains q, or, equivalently, whenever p is greater than q coordinatewise. We say that a set of points $\mathcal{P} \subset \mathbb{R}^3$ is *independent* whenever no point in $\mathcal{P}$ is dominated by another. Finally, we say that a point set is in *general position* whenever no two points have the same x, y, or z-coordinates. By a standard perturbation argument it suffices to prove Theorem 1.1 for point sets in general position.

LEMMA 3.1. *For every finite independent set $\mathcal{P} \subset \mathbb{R}^3$ in general position, there exists a collection $\mathcal{N}$ of negative octants such that*

(i) $|\mathcal{N}| = 2|\mathcal{P}| + 1$,
(ii) *the octants in $\mathcal{N}$ do not contain any point of $\mathcal{P}$ in their interior,*
(iii) *all points of $\mathbb{R}^3$ that do not dominate any point in $\mathcal{P}$ are contained in $\bigcup \mathcal{N}$.*

Proof. Let $n = |\mathcal{P}|$. We prove the lemma by induction on n. For $n = 0$ we take the negative octant covering the whole space with apex (∞, ∞, ∞). If $\mathcal{P} = \{p\}$, then we take the octants with apices (p_x, ∞, ∞), (∞, p_y, ∞), and (∞, ∞, p_z). For $n \geqslant 2$ we consider the points of $\mathcal{N}$ in order of increasing z-coordinates. Let us denote them by $p_1, p_2, \ldots, p_n$ in this order. Note that since $\mathcal{P}$ is independent, we have $p_{i,x} < p_{j,x}$ or $p_{i,y} < p_{j,y}$ for every $i, j \in [n]$ such that $j < i$.

Suppose, for the sake of induction, that there exists such a collection $\mathcal{N}_{n-1}$ for the first $n-1$ points of $\mathcal{P}$. We then consider the next point p_n and construct a new collection $\mathcal{N}_n$. We do this in three steps. First, we include in $\mathcal{N}_n$ all the octants of $\mathcal{N}_{n-1}$ that do not contain p_n. Then for each octant $Q' \in \mathcal{N}_{n-1}$ such that $p_n \in Q'$, we let Q be the octant having the same apex as Q', but with its z-coordinate changed to $p_{n,z}$. We add each such octant Q to $\mathcal{N}_n$. Finally, we add two new octants to $\mathcal{N}_n$. The first octant, L_n (for *left*), will have the point $(p_{n,x}, y, \infty)$ as apex, where $y = \min(\{p_{j,y} : 1 \leqslant j < i, p_{j,x} < p_{n,x}\} \cup \{\infty\})$. The second, B_n (for *bottom*), will have the point $(x, p_{n,y}, \infty)$ as apex, where $x = \min(\{p_{j,x} : 1 \leqslant j < i, p_{j,y} < p_{n,y}\} \cup \{\infty\})$. See Figure 3 for an illustration.

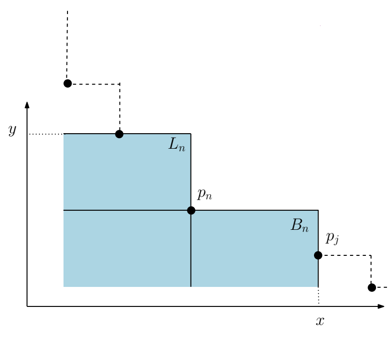

FIG. 3. *Octants L_n and B_n in the proof of Lemma* 3.1.

The first property on the cardinality of $\mathcal{N}_n$ holds by construction, as we add exactly two octants at each iteration. The second property can be checked as follows. First, by the induction hypothesis, octants in $\mathcal{N}_{n-1}$ avoid $p_1, \ldots, p_{n-1}$. Those octants from $\mathcal{N}_{n-1}$ which avoid p_n were copied to $\mathcal{N}_n$ and others have their z-coordinate modified in a way to avoid p_n. Finally, the two new octants L_n and B_n have their interiors disjoint from $\mathcal{P}$ by definition and the fact that $\mathcal{P}$ is independent.

In order to verify the third property, let us consider a point p' that is not dominating any point of $\mathcal{P}$. First, suppose that $p'_z < p_{n,z}$. By induction, there exists an octant in $\mathcal{N}_{n-1}$ containing p'. This octant is either contained in $\mathcal{N}_n$, or has its counterpart in $\mathcal{N}_n$ with a modified z-coordinate. In both cases, p' is covered by this octant in $\mathcal{N}_n$. Now suppose that $p'_z \geqslant p_{n,z}$. We can further suppose that p' neither belongs to L_n nor to B_n. Then either $p'_x > x$, or $p'_y > y$, where x and y are the two values used to define L_n and B_n. Let us suppose that $p'_x > x$, the other case being symmetric. Let $p_j, j < n$, be the point realizing the minimum in the definition of x. We must have $p'_y < p_{j,y}$, as otherwise p' would dominate p_j. Then p' must be covered by an octant $Q \in \mathcal{N}_{n-1}$ whose y-coordinate is smaller than $p_{j,y}$, as otherwise Q would contain p_j. But by definition $p_{j,y} < p_{n,y}$; hence Q does not contain p_n and therefore also belongs to $\mathcal{N}_n$. In all cases, p' is contained in an octant of $\mathcal{N}_n$ and the third property holds. ∎

Note that the upper bound on the size of $\mathcal{N}$ in Lemma 3.1 is tight. For example, consider the point sets $\mathcal{P}_n = \{(i, -i, -i) \mid i = 1, \ldots, n\}$. Indeed, Lemma 3.1 and the fact that it is tight for all point sets that are in general position and do not lie in a plane containing the all-ones vector can be deduced from a more general theorem of Scarf [14].

In order to prove our main theorem, we will use Theorem 2.1, due to Keszegh and Pálvölgyi [19]. We proceed to describe a coloring algorithm that achieves the bound of Theorem 1.1. We do this in two steps. First, we consider the case where the points to color form an independent set.

LEMMA 3.2. *Let c be a constant satisfying the property in Theorem* 2.1. *For any positive integer k, every finite independent set $\mathcal{P} \subset \mathbb{R}^3$ in general position can be colored with k colors so that every negative octant containing at least $ck^{\log_2(2c-1)}$ points of $\mathcal{P}$ contains at least one of each color.*

Proof. For $k = 2$, we know there exists a 2-coloring of $\mathcal{P}$ satisfying the property of Theorem 2.1. Suppose now, as an induction hypothesis, that we have a k-coloring ϕ of $\mathcal{P}$ such that every octant containing at least $p(k)$ points contains at least one of each color. Label the colors of ϕ by $1, \ldots, k$.

We now describe a $2k$-coloring ϕ'. For $i \in [k]$, let $\mathcal{P}_i = \phi^{-1}(i)$ be the set of points with color i. We know from Theorem 2.1 that there exists a 2-coloring $\phi_i : \mathcal{P}_i \to \{i', i''\}$ of $\mathcal{P}_i$ such that every octant containing at least c points of $\mathcal{P}_i$ contains at least one of each color i' and i''. We now define ϕ' as the $2k$-coloring obtained by partitioning each color class in this way. We now claim that the coloring ϕ' is such that any octant containing at least $(2c-1)p(k)$ points contains at least one of each of the $2k$ colors.

For the sake of contradiction, let Q be an octant containing at least $(2c-1)p(k)$ points of $\mathcal{P}$, but not any point of color i' in ϕ'. Let $\mathcal{P}_Q \subseteq \mathcal{P}$ be the set of points contained in Q. If Q does not contain any point of color i', it means that it contains at most $c-1$ points of $\phi^{-1}(i)$. Let $\mathcal{P}_i = \phi^{-1}(i) \cap \mathcal{P}_Q$ be the points of color i in ϕ contained in Q.

From Lemma 3.1 and the fact that $\mathcal{P}_Q \subset \mathcal{P}$ is an independent set, we know that there exists a collection $\mathcal{N}$ of at most $2(c-1)+1 = 2c-1$ octants whose interiors do not contain any point of $\mathcal{P}_i$, but that collectively cover all points of $\mathcal{P}_Q \setminus \mathcal{P}_i$. Indeed, we can assume that after intersecting with Q, $\mathcal{N}$ covers precisely $\mathcal{P}_Q \setminus \mathcal{P}_i$ and no other point of $\mathcal{P}$.

Hence from the pigeonhole principle, one of the octants $N \in \mathcal{N}$ contains at least $\lceil ((2c-1)p(k) - (c-1))/(2c-1) \rceil = p(k)$ points of $\mathcal{P}_Q$ in its interior, but no point of $\mathcal{P}_i$. From the general position assumption, we can find an octant contained in N that contains *exactly* $p(k)$ points of $\mathcal{P}_Q$, but no point of $\mathcal{P}_i$. But this is a contradiction with the induction hypothesis, since this octant should have contained a point of color i in ϕ.

It remains to solve the following recurrence, with starting value $p(2) = c$:

$$\begin{aligned} p(2k) &\leqslant (2c-1)p(k), \\ p(k) &\leqslant c(2c-1)^{\lceil \log_2 k \rceil - 1} \\ &< ck^{\log_2(2c-1)}. \qquad \square \end{aligned}$$

We now describe our algorithm for coloring an arbitrary set of points in general position. This requires a new definition.

Given a set $\mathcal{P}$ of points in general position in $\mathbb{R}^3$, the *minimal points* of $\mathcal{P}$ is the subcollection of points of $\mathcal{P}$ that are not dominating any other point of $\mathcal{P}$. In general, we define the *ith layer* $\mathcal{L}_i$ of $\mathcal{P}$ as its minimal points for $i = 1$, and as the minimal points of $\mathcal{P} \setminus \bigcup_{1 \leqslant j < i} \mathcal{L}_j$ for $i > 1$. By definition each layer is an independent set of points.

LEMMA 3.3. *Let $f(k) = ck^{\log_2(2c-1)}$ be the function derived in Lemma* 3.2*, where c is a constant satisfying the property in Theorem* 2.1*. For any positive integer k, every finite set $\mathcal{P} \subset \mathbb{R}^3$ in general position can be colored with k colors so that every negative octant containing at least $(k-1)f(k)$ points of $\mathcal{P}$ contains at least one of each color.*

Proof. We will color the points of $\mathcal{P}$ by considering the successive layers one by one, starting with the minimal points. For each layer $\mathcal{L}_i$, we do the following:

- precolor the points of $\mathcal{L}_i$ with colors in $[k]$, as is done in Lemma 3.2;
- for each point $p \in \mathcal{L}_i$:
  - consider the set of points $D_p = \{q \in \mathcal{P} : q \text{ dominated by } p\}$;
  - if p is precolored with a color that is not used for any point in D_p, then this color is the final color of p;
  - otherwise pick any color not present on points in D_p and color p with it; if all k colors are used within D_p, leave p uncolored.

The main observation here is that although the recoloring step harms the validity of the coloring within a single layer, it is globally innocuous, since any octant containing the point p in the ith layer also contains all the points in D_p, from the previous layers. Thus, any octant containing p contains a point colored by the same color as the precolor of p. Note that each point in the ith layer must dominate at least one point from each $i-1$ earlier layers. This forces the invariant that any octant containing a point of the ith layer contains points with at least i distinct colors. In particular, any octant containing a point of the kth layer will contain all the colors.

The analysis is now straightforward. Suppose that an octant contains at least $(k-1)f(k)$ points. If it contains a point of the kth layer, then it contains all k colors. Otherwise, it must contain points of at most $k-1$ layers, and from the pigeonhole principle, it contains at least $(k-1)f(k)/(k-1) = f(k)$ points in a single layer. Then the precoloring of this single layer guarantees each octant of size at least $f(k)$ to be colorful. $\square$

Now Theorem 1.1 follows by replacing c by 12 in the expression of Lemma 3.3, yielding $a \simeq 5.58$.

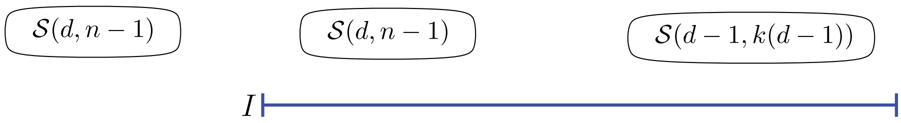

FIG. 4. *Defining strategy $\mathcal{S}(d,n)$ once $\mathcal{S}(d-1,k(d-1))$ and $\mathcal{S}(d,n-1)$ are defined, in the case where $t_i = t'_i$ for all $i \in [k]$.*

4. Proof of Theorem 1.6. We say that a point of the real line is *d-covered*, if it is contained in exactly d intervals presented so far. We shall define for every d and n an adversarial strategy $\mathcal{S}(d,n)$ for presenting intervals such that the following is true:

(i) Every semionline proper k-coloring algorithm of value at most d executed against $\mathcal{S}(d,n)$ yields k points $p_1, \dots, p_k$ such that for every $i \in [k]$, the point p_i is eventually covered by exactly t_i intervals, all of which have color i, and

(ii) $t_1 + \cdots + t_k \geqslant n$.

Clearly, if for some semionline k-coloring algorithm $\mathcal{A}$ there is a point eventually covered by at least d intervals, all of which have the same color, then the value of $\mathcal{A}$ is at least $d+1$. Thus if $S(d,kd)$ exists and satisfies (i) and (ii), then there is no semionline k-coloring algorithm of value at most d, which proves the theorem.

We prove the existence of $\mathcal{S}(d,n)$ by a double induction on d and n. Strategies $\mathcal{S}(d,0)$ are vacuous as (i) and (ii) for $n=0$ hold for the empty set of intervals and any set of k distinct points $p_1, \dots, p_k$. We define $\mathcal{S}(d,n)$, for $n > 0$, once we have defined $\mathcal{S}(d-1,k(d-1))$ and $\mathcal{S}(d,n-1)$.

Before continuing let us present the following useful claim.

Claim. Consider a set $\mathcal{I}$ of intervals already presented, $I \in \mathcal{I}$ and $I' \notin \mathcal{I}$ such that $I' \subset I$ and $I' \cap J = \emptyset$ for all $J \in \mathcal{I} \setminus I$. If $\mathcal{S}(d-1,k(d-1))$ exists, then we can present the intervals of $\mathcal{S}(d-1,k(d-1))$ inside I' forcing any semionline algorithm of value at most d to color I.

Proof. We present the intervals for $\mathcal{S}(d-1,k(d-1))$ completely inside I'. If the algorithm does not color I, then it can be seen as a k-coloring algorithm of value at most $d-1$ executed against $\mathcal{S}(d-1,k(d-1))$. We already know that there is no such algorithm and therefore every k-coloring algorithm of value at most d has to color interval I. $\square$

Now, we are ready to define $S(d,n)$ for $n > 0$. First, present two families of intervals, both realizing strategy $S(d,n-1)$, disjointly next to each other. By (i) there exist two sets of k points each, $p_1, p_2, \dots, p_k$ and $p'_1, \dots, p'_k$, and nonnegative integers $t_1, \dots, t_k, t'_1, \dots, t'_k$ such that p_i is t_i-covered and all its intervals are colored with i, and also p'_i is t'_i-covered and all its intervals are colored with i, for every $i \in [k]$. Moreover, by (ii) we have $t_1 + \ldots + t_k \geqslant n$ and $t'_1 + \ldots + t'_k \geqslant n$.

If there exists some $i \in \{1, \dots, k\}$ with $t_i \neq t'_i$, then the sequence of maxima $m_i = \max(t_i, t'_i)$ satisfies $m_1 + \cdots + m_k \geqslant n+1$. Thus, taking for each $i \in \{1, \dots, k\}$ the point from $\{p_i, p'_i\}$ that corresponds to the larger value of t_i, t'_i, we obtain a set of k points satisfying (i) and (ii).

Hence we assume without loss of generality that $t_i = t'_i$ for all $i \in \{1, \ldots, k\}$. Then we present one additional interval I that contains all the points $p'_1, \ldots, p'_k$ but none of the points $p_1, \ldots, p_k$. Moreover, I is chosen big enough so that there exists some $I' \subset I$ that is disjoint from all the other intervals presented so far. We present the intervals realizing strategy $S(d-1, k(d-1))$ inside I', forcing I to be colored (see Figure 4). Let j be the color of I. Then p'_j is now contained in exactly $t'_j + 1$ intervals all of which are colored with j. Thus $(\{p_1, \ldots, p_k\} \setminus p_j) \cup \{p'_j\}$ is a set of k points satisfying (i) and (ii), which concludes the proof. ∎

Discussion and open problems. A well-studied problem in discrete geometry is to identify properties of range spaces, or geometric hypergraphs, that allow one to find small *ε-nets*. It is known, for instance, that range spaces of bounded VC-dimension have ε-nets of size $O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$. (See, for instance, Chapter 10 in Matoušek's lectures [23].)

The coloring problem that we consider can be cast as the problem of partitioning a point set into k ε-nets for $\varepsilon = p(k)/n$. In fact, it is one of the negative results for covering decomposition that formed the basis of a construction of Pach and Tardos for proving superlinear lower bounds on the size of ε-nets [27]. One can realize that if $p(k) = O(k)$ for a given range space, then it implies that this range space also has ε-nets of size $O(1/\varepsilon)$. The latter is known to hold for range spaces induced by octants [13]. Whether $p(k) = O(k)$ for octants is therefore an interesting open problem. In general, giving improved upper or lower bounds on $p(k)$ for octants is the major remaining open question.

Another interesting open question concerns the primal problem, in which points are colored with k colors so that every region containing $p(k)$ points contains a point of each color. The existence of such a function $p(k)$ is still open for homothetic copies of a square, for instance.

Acknowledgments. The authors wish to thank Balázs Keszegh, János Pach, and Dömötör Pálvölgyi for stimulating discussions on this topic.

REFERENCES

[1] D. Ajwani, K. Elbassioni, S. Govindarajan, and S. Ray, *Conflict-free coloring for rectangle ranges using $O(n^{.382})$ colors*, Discrete Comput. Geom., 48 (2012), pp. 39–52.

[2] N. Alon and J. Spencer, *The Probabilistic Method*, 2nd ed., Wiley-Intersci. Ser. Discrete Math. Optim., John Wiley & Sons Inc., New York, 2000.

[3] G. Aloupis, J. Cardinal, S. Collette, S. Langerman, D. Orden, and P. Ramos, *Decomposition of multiple coverings into more parts*, Discrete Comput. Geom., 44 (2010), pp. 706–723.

[4] G. Aloupis, J. Cardinal, S. Collette, S. Langerman, and S. Smorodinsky, *Coloring geometric range spaces*, Discrete Comput. Geom., 41 (2009), pp. 348–362.

[5] A. Asinowski, J. Cardinal, N. Cohen, S. Collette, T. Hackl, M. Hoffmann, K. Knauer, S. Langerman, M. Lasoń, P. Micek, G. Rote, and T. Ueckerdt, *Coloring hypergraphs induced by dynamic point sets and bottomless rectangles*, in Proceedings of the Workshop on Data Structures and Algorithms (WADS), 2013, pp. 73–84.

[6] A. Bar-Noy, P. Cheilaris, S. Olonetsky, and S. Smorodinsky, *Online conflict-free colouring for hypergraphs*, Combin. Probab. Comput., 19 (2010), pp. 493–516.

[7] A. Bar-Noy, P. Cheilaris, and S. Smorodinsky, *Deterministic conflict-free coloring for intervals: From offline to online*, ACM Trans. Algorithms, 4 (2008), article 44.

[8] P. Brass, W. Moser, and J. Pach, *Research Problems in Discrete Geometry*, Springer, New York, 2005.

[9] J. Cardinal, K. Knauer, P. Micek, and T. Ueckerdt, *Making triangles colorful*, J. Comput. Geom., 4 (2013), pp. 240–246.

[10] J. Cardinal and M. Korman, *Coloring planar homothets and three-dimensional hypergraphs*, Comput. Geom., 46 (2013), pp. 1027–1035.

[11] T. M. Chan, E. Grant, J. Könemann, and M. Sharpe, *Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling*, in Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, Philadelphia, 2012, pp. 1576–1585.

[12] X. Chen, J. Pach, M. Szegedy, and G. Tardos, *Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles*, Random Structures Algorithms, 34 (2009), pp. 11–23.

[13] K. L. Clarkson and K. R. Varadarajan, *Improved approximation algorithms for geometric set cover*, Discrete Comput. Geom., 37 (2007), pp. 43–58.

[14] S. Felsner and S. Kappes, *Orthogonal surfaces and their CP-orders*, Order, 25 (2008), pp. 19–47.

[15] R. Fulek, *Coloring geometric hypergraph defined by an arrangement of half-planes*, in Proceedings of the Canadian Conference on Computational Geometry (CCCG), 2010, pp. 71–74.

[16] M. Gibson and K. R. Varadarajan, *Optimally decomposing coverings with translates of a convex polygon*, Discrete Comput. Geom., 46 (2011), pp. 313–333.

[17] B. Keszegh, *Coloring half-planes and bottomless rectangles*, Comput. Geom., 45 (2012), pp. 495–507.

[18] B. Keszegh, N. Lemons, and D. Pálvölgyi, *Online and quasi-online colorings of wedges and intervals*, in Proceedings of the Theory and Practice of Computer Science (SOFSEM), Lecture Notes in Comput. Sci. 7741, Springer, Heidelberg, 2013, pp. 292–306.

[19] B. Keszegh and D. Pálvölgyi, *Octants are cover-decomposable*, Discrete Comput. Geom., 47 (2012), pp. 598–609.

[20] B. Keszegh and D. Pálvölgyi, *Convex polygons are self-coverable*, Discrete Comput. Geom., 51 (2014), pp. 885–895.

[21] B. Keszegh and D. Pálvölgyi, *Octants are cover-decomposable into many coverings*, Comput. Geom., 47 (2014), pp. 585–588.

[22] I. Kovács, *Indecomposable coverings with homothetic polygons*, CoRR, abs/1312.4597 (2013).

[23] J. Matoušek, *Lectures on Discrete Geometry*, Grad. Texts in Math. 212. Springer-Verlag, New York, 2002.

[24] J. Pach, *Decomposition of multiple packing and covering*, in 2. Kolloq. über Diskrete Geom., Inst. Math. Univ. Salzburg, 1980, pp. 169–178.

[25] J. Pach, *Covering the plane with convex polygons*, Discrete Comput. Geom., 1 (1986), pp. 73–81.

[26] J. Pach and G. Tardos, *Coloring axis-parallel rectangles*, J. Combin. Theory Ser. A, 117 (2010), pp. 776–782.

[27] J. Pach and G. Tardos, *Tight lower bounds for the size of epsilon-nets*, in Proceedings of the Symposium on Computational Geometry (SoCG), ACM, New York, 2011, pp. 458–463.

[28] J. Pach, G. Tardos, and G. Tóth, *Indecomposable coverings*, in Proceedings of the 7th China-Japan Conference on Discrete Geometry, Combinatorics and Graph Theory (CJCDGCGT), 2005, Lecture Notes in Comput. Sci. 4381, Springer, Berlin, 2007, pp. 135–148.

[29] J. Pach and G. Tóth, *Decomposition of multiple coverings into many parts*, Comput. Geom., 42 (2009), pp. 127–133.

[30] D. Pálvölgyi, *Indecomposable coverings with concave polygons*, Discrete Comput. Geom., 44 (2010), pp. 577–588.

[31] D. Pálvölgyi, *Indecomposable coverings with unit discs*, CoRR, abs/1310.6900 (2013).

[32] R. Rado, *Covering theorems for ordered sets*, Proc. London Math. Soc. (2), 50 (1949), pp. 509–535.

[33] S. Smorodinsky, *On the chromatic number of geometric hypergraphs*, SIAM J. Discrete Math., 21 (2007), pp. 676–687.

[34] S. Smorodinsky and Y. Yuditsky, *Polychromatic coloring for half-planes*, in Proceedings of the Scandinavian Worskhop on Algorithms Theory (SWAT), Lecture Notes in Comput. Sci. 6139, Springer, Berlin, 2010, pp. 118–126.

[35] G. Tardos and G. Tóth, *Multiple coverings of the plane with triangles*, Discrete Comput. Geom., 38 (2007), pp. 443–450.

[36] K. Varadarajan, *Weighted geometric set cover via quasi-uniform sampling*, in Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC), ACM, New York, 2010, pp. 641–648.

The Density of Fan-Planar Graphs

Michael Kaufmann[1] and Torsten Ueckerdt[2]

March 26, 2014

Abstract

A topological drawing of a graph is fan-planar if for each edge e the edges crossing e have a common endpoint on the same side of e, and a fan-planar graph is a graph admitting such a drawing. Equivalently, this can be formulated by two forbidden patterns, one of which is the configuration where e is crossed by two independent edges and the other where e is crossed by incident edges with the common endpoint on different sides of e. In particular every edge of a fan-planar graph is crossed only by the edges of a star. A topological drawing is simple if any two edges have at most one point in common.

The class of fan-planar graphs is a natural variant of other classes defined by forbidden intersection patterns in a topological drawing of the graph. So every 1-planar graph is also fan-planar, and every fan-planar graph is also quasiplanar, where both inclusions are strict. Fan-planar graphs also fit perfectly in a recent series of work on nearly-planar graphs from the area of graph drawing and combinatorial embeddings.

For topologically defined graph classes, one of the most fundamental questions asks for the maximum number of edges in any such graph with n vertices. We prove that every n-vertex graph without loops and parallel edges that admits a simple fan-planar drawig has at most $5n - 10$ edges and that this bound is tight for every $n \geq 20$.

Furthermore we discuss possible extensions and generalizations of these new concepts.

Keywords: Topological drawing, quasiplanar, 1-planar, intersection pattern, density.

1 Introduction

Planarity of a graph is a well-studied concept in graph theory, computational geometry and graph drawing. The famous Euler formula characterizes for a certain embedding the relation between vertices, edges and faces, and many different algorithms [28, 23, 11] following different objectives have been developed to compute appropriate embeddings in the plane.

Because of the importance of the concepts, a series of generalizations have been developed in the past. Topological graphs and topological drawings respectively are being considered, i.e., the vertices are drawn as points in the plane and the edges drawn as Jordan curves between corresponding points without any other vertex as an interior point. In [16], the authors state "Finding the maximum number of edges in a topological graph with a forbidden crossing pattern is a fundamental problem in extremal topological graph theory" together with 9 citations from a large group of authors. Most of the existent literature considers topological drawings that are *simple*, i.e., where any two edges have at most one point in common. In particular, two edges may not cross more than once and incident edges may not cross at all. Throughout this paper we shall consider simple topological graphs only. Indeed, we shall argue in Section 4 that if we drop this assumptions and allow non-homeomorphic parallel edges, then even 3-vertex fan-planar graphs have arbitrarily many edges.

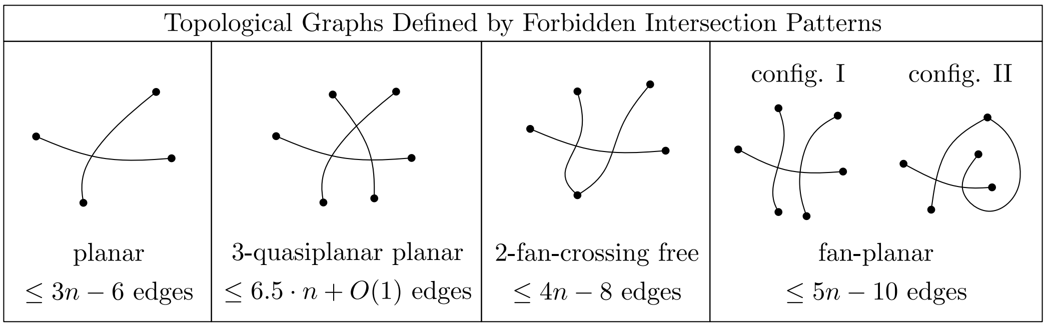

Figure 1: Topological graphs defined by forbidden patterns and the corresponding maximum number of edges in an n-vertex such graph.

Related work. Most notably are the k-planar graphs and the k-quasiplanar graphs [4]. A k-planar graph admits a topological drawing in which no edge is crossed more than k times by other edges, while a k-quasiplanar graph admits a drawing in which no k edges pairwise cross each other.

The topic of k-quasiplanar graphs is almost classical [9]. A famous conjecture [9] states that for constant k the maximal number of edges in k-quasiplanar graphs is linear in the number of vertices. Note that 2-quasiplanar graphs correspond to planar graphs. A first linear bound for $k = 3$, i.e. 3-quasiplanar graphs, has been shown in [4] and subsequently improved in [21]. For 4-quasiplanar graphs the current best bound is $76(n-2)$ [1]. For the general case, the bounds have been gradually improved from $O(n(\log n)^{O(\log k)})$ [21], and $O(n \log n \cdot 2^{\alpha(n)^c})$.

In case of simple topological drawings, where each pair of edges intersects at most once, a bound of $6.5n+O(1)$ has been proven for 3-quasiplanar graphs [3] and recently $O(n \log n)$ for k-quasiplanar graphs with any fixed $k \geq 2$ [24]. It is still open, if the conjecture holds for general k.

A k-planar graph admits a topological drawing in which each edge has at most k crossings. The special case of 1-planar graphs have been introduced by Ringel [22], who considered the chromatic number of these graphs. Important work about the characterization on 1-planar graphs has been performed by Suzuki [25], Thomassen [27] and Hong *et al.* [19]. Related questions on testing 1-planarity have been explored, where NP-completeness has been shown for the general case [17] while efficient algorithms have been found for testing 1-planarity for a given rotation system [14] and for the case of outer-planarity [7, 18]. Additionally aspects like straight-line embeddings [5] and maximality [8] etc. have been explored in the past.

Closely related to 1-planar graphs are RAC-drawable graphs [13, 6], that is graphs that can be drawn in the plane with straight-line edges and right-angle crossings. For the maximum number of edges in such a graph with n vertices, a bound of $4n - 10$ could be proven [15], which is remarkably close to the $4n - 8$ bound for the class of 1-planar graphs [21]. A necessary condition for RAC-drawable graph is the absence of fan-crossings. An edge has a k-fan-crossing if it crosses k edges that have a common endpoint, cf. Figure 1. RAC-drawings do not allow 2-fan-crossings. In a recent paper [10], Cheong *et al.* considered k-fan-crossing free graphs and gave bounds for their maximum number of edges. They obtain a tight bound of $4n - 8$ for n-vertex 2-fan-crossing free graphs, and a tight $4n - 9$ when edges are required to be straight-line segments. For $k > 2$, they prove an upper bound of $3(k - 1)(n - 2)$ edges, while all known examples of k-fan-crossing free graphs on n vertices have no more than kn edges.

Our results and more related work. Throughout this paper we consider only simple topological drawings, i.e., any two edges have at most one point in common, and only simple graphs, i.e., graphs without loops and parallel edges. We consider here another variant of sparse non-planar graphs, somehow halfway between 1-planar graphs and quasiplanar graphs, where we allow more than one crossing on an edge e, but only if the crossing edges have a common endpoint on the same side of e. We call this a **fan-crossing** and the class of topological graphs obtained this way **fan-planar graphs**. Note that we do not differentiate on k-fan-crossings as it has been done by Cheong *et al.* [10].

The requirement that every edge in G is crossed by a fan-crossing can be stated in terms of forbidden configurations. We define *configuration I* to be one edge that is crossed by two independent edges, and *configuration II* to be an edge e that is crossed by incident edges, which however have their common endpoint on different sides of e, see Figure 2. Note that since we consider only simple topological drawings, configuration II is well-defined. Now a simple topological graph is fan-planar if and only if neither configuration I nor II occurs. Note that if we forbid only configuration I, then an edge may be crossed by the three edges of a triangle, which is actually not a star, nor a fan-crossing. However, if every edge is drawn as a straight-line segment, then configuration II can not occur and hence in this case it is enough to forbid configuration I.

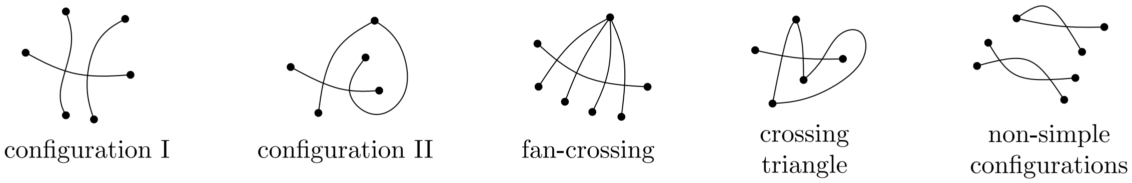

Figure 2: Crossing configurations

Obviously, 1-planar graphs are also fan-planar. Furthermore, fan-planar graphs are 3-quasiplanar since there are no three independent edges that mutually cross. So, we know already that the maximum number of edges in an n-vertex fan-planar graph is approximately between $4n$ and $6.5n$. In the following, we will explore the exact bound.

Theorem 1. *Every simple topological graph G on $n \geq 3$ vertices with neither configuration I nor configuration II has at most $5n - 10$ edges. This bound is tight for $n \geq 20$.*

We remark that fan-planar drawings graphs may have $\Omega(n^2)$ crossings, e.g., a straight-line drawing of $K_{2,n}$ with the bipartition classes places on two parallel lines.

Very closely related to our approach is the research on forbidden grids in topological graphs, where a (k, l) grid denotes a k-subset of the edges pairwise intersected by an l-subset of the edges, see [20] and [26]. It is known that topological graphs without (k, l) grids have a linear number of edges if k and l are fixed. Note that configuration I, but also a 2-fan-crossing, are $(2, 1)$ grids. Subsequently [2], "natural" (k, l) grids have been considered, which have the additional requirement that the k edges, as well as the l edges, forming the grid are pairwise disjoint. For natural grids, the achieved bounds are superlinear. Linear bounds on the number of edges have been found for the special case of forbidden natural $(k, 1)$ grids where the leading constant heavily depends on the parameter k. In particular, the authors give a bound of $65n$ for the case of forbidden natural $(2, 1)$ grids, which correspond to our forbidden configuration I. Additionally, the case of geometric graphs, that is, graphs with straight-line edges, has been explored. For details and differences let us refer to [2]. We remark that many arguments in this field of research are based on the probabilistic method, while in this paper we use a direct approach aiming on tight upper and lower bounds.

2 Examples of Fan-Planar Graphs with Many Edges

The following examples have approximately $5n$ edges. The first one is a $K_{4,n-4}$, where the $n-4$ edges are connected by a path, see Figure 3(a). An easy calculation shown that this graph has $4(n-4)+(n-5) = 5n-21$ edges. Indeed, one can add 10 edges to the graph, keeping fan-planarity, as well as additionally one vertex with 6 more incident edges and obtain a graph on $n+1$ vertices and $5(n+1)-10$ edges. We remark that this graph has parallel edges; however every pair of parallel edges is non-homeomorphic, that is, it surrounds at least one vertex of G.

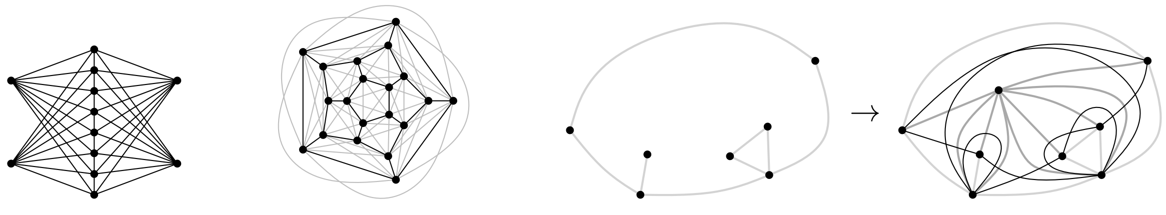

Figure 3: (a) $K_{4,n-4}$ with $n-4$ vertices on a path. (b) The dodecahedral graph with a pentagram in each face. (c) Adding 2-hops and spokes into a face.

The second example is the (planar) dodecahedral graph where in each 5-face, we draw 5 additional edges as a pentagram, see Figure 3(b). This graph has $n = 20$ vertices and $5n - 10 = 90$ edges, and has already served as a tight example for 2-planar graphs [21].

Proposition 1. *Every connected planar embedded graph H on $n \geq 3$ vertices can be extended to a fan-planar graph G with $5|V(G)|-10$ edges by adding an independent set of vertices and sufficiently many edges, such that the uncrossed edges of G are precisely the edges of H.*

Moreover, if H is 3-connected and each face has length at least 5, then G is a simple topological graph without loops or parallel edges.

Proof. Let n and m be the number of vertices and edges of H, respectively, and F be the set of all faces of H. We construct the fan-planar graph G by adding one vertex and two sets of edges into each face $f \in F$. So let f be any face of H. Since H is connected, f corresponds to a single closed walk $v_1, \ldots, v_s$ in H around f, where vertices and edges may be repeated. We do the following, which is illustrated in Figure 3(c).

(1) Add a new vertex v_f into f.
(2) For $i = 1, \ldots, s$ add a new edge $v_f v_i$ drawn in the interior of f.
(3) For $i = 1, \ldots, s$ add a new edge $v_{i-1}v_{i+1}$ (with indices modulo s) crossing the edge $v_f v_i$.

In (1) we added $|F|$ new vertices. In (2) we added $\deg(f)$ many "spoke edges" inside face f, in total $\sum_f \deg(f) = 2m$ new edges. And in (3) we added again $\deg(f)$ many "2-hop edges" inside face f, in total $\sum_f \deg(f) = 2m$ new edges. Thus we calculate

$$\begin{aligned} |V(G)| &= n + |F| \\ |E(G)| &= m + 2m + 2m = 5m, \end{aligned}$$

which together with Euler's formula $m = n + |F| - 2$ gives $|E(G)| = 5|V(G)| - 10$. It remains to see that no two edges in G are homeomorphic, and that G is fan-planar. The "2-hop edges" form shortcuts for paths of length 2. Since $s \geq 4$ by assumption, none of these s are edges is already in the facial walk for f. Each "spoke edge" $v_f v_i$ crosses only one 2-hop edge, and each 2-hop edge $v_{i-1}v_{i+1}$ crosses only three edges $v_{i-2}v_i$, $v_f v_i$ and $v_i v_{i+2}$, which have v_i as a common endpoint. Hence the resulting graph G is fan-planar.

Finally, note that if the planar graph H is 3-connected and each face has length at least 5, then the fan-planar graph G has no loops, nor parallel edges, nor crossing incident edges. Examples for such planar graphs are fullerene graphs. □

3 The $5n - 10$ Upper Bound For the Number of Edges

In this section we prove Theorem 1. We shall fix a fan-planar embedding of G and split the edges of G into three sets. The first set contains all edges that are uncrossed. We denote by H the subgraph of G with all vertices in V and all uncrossed edges of G. Sometimes we may refer to H as the *planar subgraph of* G. Note that H might be disconnected even if G is connected. In the second set we consider every crossed edge whose endpoints lie in the same connected component of H. And the third set contains all remaining edges, i.e., every crossed edge with endpoints in different components of H. We show how to count the edges in each of the three sets and derive the upper bound.

To prove Theorem 1 it clearly suffices to consider simple topological graphs G that do not contain configuration I nor II and additionally satisfy the following properties.

(i) The chosen embedding of G has the maximum number of uncrossed edges.
(ii) The addition of any edge to the given embedding violates the fan-planarity of G, that is, G is maximal fan-planar with respect to the given embedding.

So for the remainder of this paper let G be a maximal fan-planar graph with a fixed fan-planar embedding that has the maximum number of uncrossed edges. Recall that the embedding of G is simple, i.e., any two edges have at most one point in common.

3.1 Notation, Definitions and Preliminaries Results

We call a connected component of the plane after the removal of all vertices and edges of G a *cell of* G. Whenever we consider a subgraph of G we consider it together with its fan-planar embedding, which is inherited from the embedding of G. We will sometimes consider cells of a subgraph G' of G, even though those might contain vertices and edges of $G - G'$. The boundary of each cell c is composed of a number of edge segments and some (possibly none) vertices of G'. With slight abuse of notation we call the cyclic order of vertices and edge segments along c the *boundary of* c, denoted by ∂c. Note that vertices and edges may appear more than once in the boundary of a single cell. We define the *size of a cell* c, denoted by $||c||$, as the total number of vertices and edge segments in ∂c counted with multiplicity.

Note that from the additional assumptions (i) and (ii) on G it follows that if two vertices are in the same cell c of G then they are connected by an uncrossed edge of G. However, this uncrossed edge does not necessarily bound the cell c.

Lemma 1. *If two edges vw and ux cross in a point p, no edge at v crosses ux between p and u, and no edge at x crosses vw between p and w, then u and w are contained in the same cell of G.*

Proof. Let $e_0 = ux$ and $e_1 = vw$ be two edges that cross in point $p = p_1$ such that no edge at v crosses e_0 between p_1 and u, and no edge at x crosses e_1 between p_1 and w. If no edge of G crosses e_0 nor e_1 between p_1 and u, respectively w, then clearly u and w are bounding the same cell. So assume without loss of generality that some edge of G crosses e_1 between p_1 and w. By fan-planarity such edges are incident to u. Let e_2 be the edge whose crossings with e_1 is closest to w, and let p_2 be the crossing point. See Figure 4(a) for an illustration.

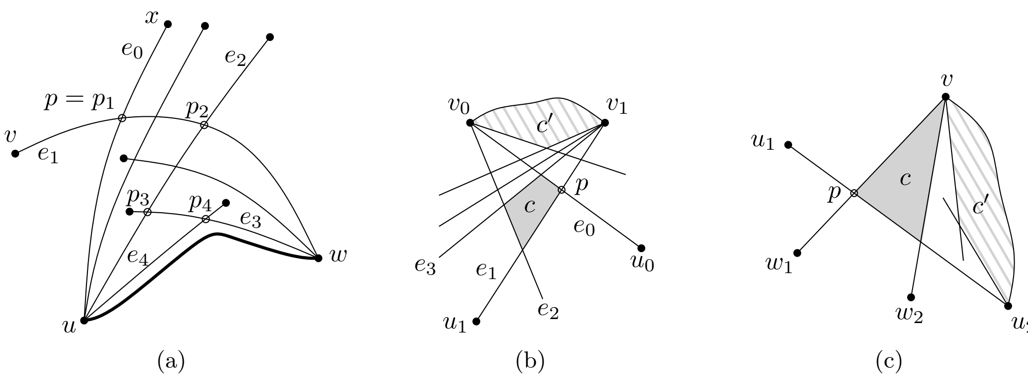

Figure 4: Illustration of the proofs of Lemma 1 (a) and Corollary 2 (b),(c).

No edge crosses e_1 between w and p_2. If e_2 is not crossed between u and p_2, then u and w are bounding the same cell and we are done. Otherwise let e_3 be the edge whose crossing with e_2 is closest to u, and let p_3 be the crossing point. By fan-planarity e_3 and e_1 have a common endpoint, and it is not v since e_3 does not cross e_0 between p_1 and u. So e_3 endpoints at w and we have that e_2 is not crossed between u and p_3. Again, if u and w are not on the same cell then some edge crosses e_3 between p_3 and w. By fan-planarity any such edge has a common endpoint with e_2, and if it would not be u then e_1 would be crossed by two independent edges – a contradiction to the fan-planarity of G. So all edges crossing e_3 between w and p_3 are incident to u. Let e_4 be such edge whose crossing with e_3 is closest to w, and let p_4 be the crossing point. Let us again refer to Figure 4(a) for an illustration.

Iterating this procedure until no edge crosses e_i nor e_{i-1} between p_i and u, w we see that u and w lie indeed on the same cell, which concludes the proof. □

Lemma 1 has a couple of nice consequences.

Corollary 1. *Any two crossing edges in G are connected by an uncrossed edge.*

Proof. Let ux and vw be the two crossing edges. By fan-planarity either no other edge at x or no other edge at u crosses the edge vw, say there is no such edge at x. Similarly, we may assume without loss of generality, that no edge at v crosses the edge ux. However, this implies that ux and vw satisfy the requirements of Lemma 1 and we have that u and w are on the same cell. In particular, we can draw an uncrossed edge between u and w in this cell. Because G is maximally fan-planar, uw is indeed an edge of G. And since G is embedded with the maximum number of uncrossed edges, uw is also drawn uncrossed. □

Corollary 2. *If c is a cell of any subgraph of G, and $||c|| = 4$, then c contains no vertex of G in its interior.*

Proof. Let c be a cell of $G' \subseteq G$ with $||c|| = 4$. Then ∂c consists either of four edge segments or one vertex and three edge segments. Let us assume for the sake of contradiction that c contains a set $S \neq \emptyset$ of vertices in its interior.

Case 1. ∂c consists of four edge segments. Let e_0, e_1, e_2, e_3 be the edges bounding c is this cyclic order. From the fan-planarity of G follows that e_0 and e_2 have a common endpoint v_0. Similarly

e_1 and e_3 have a common endpoint v_1. See Figure 4(b) for an illustration. If p denotes the crossing point of $e_0 = v_0u_0$ and $e_1 = v_1u_1$, then by fan-planarity no edge at u_i crosses e_{i+1} between p and v_{i+1}, where $i \in \{0, 1\}$ and indices are taken modulo 2. Hence by Lemma 1 there exists a cell c' of G that contains both v_0 and v_1.

Now consider the subgraph $G[S]$ of G on the vertices inside c. From the fan-planarity follows that every edge between $G[S]$ and $G[V \setminus S]$ has as one endpoint v_0 or v_1. We now change the embedding of G by placing the subgraph $G[S]$ (keeping its inherited embedding) into the cell c' that contains v_0 and v_1. The resulting embedding of G is still fan-planar and moreover at least one edge between $G[S]$ and $\{v_0, v_1\}$ is now uncrossed – a contradiction to our assumption (i) that the embedding of G has the maximum number of uncrossed edges.

Case 2. ∂c consists of one vertex and three edge segments. Let v be the vertex and vw_1, vw_2, u_1u_2 be the edges bounding c. See Figure 4(c) for an illustration. If p denotes the crossing point of vw_1 and u_1u_2, then by fan-planarity either no edge at u_1 crosses vw_1 between p and v or no edge at u_2 crosses vw_1 between p and v. Moreover, for $i = 1, 2$ the edge vw_i is the only edge at w_i that crosses u_1u_2. Hence by Lemma 1 we have that either v and u_1 or v and u_2 are contained in the same cell of G – say cell c' contains v and u_2.

Now, similarly to the previous case, consider the subgraph $G[S]$ of G on the vertices inside c. From the fan-planarity, it follows that every edge between $G[S]$ and $G[V \setminus S]$ has as one endpoint v, u_1 or u_2. Moreover, every edge between a vertex in $G[S]$ and u_1 or u_2 is crossed only by edges incident to v, as otherwise u_1u_2 would be crossed by two independent edges. We now change the embedding of G by placing the subgraph $G[S]$ (keeping its inherited embedding) into the cell c' that contains v and u_2. The resulting embedding of G is still fan-planar and moreover at least one edge between $G[S]$ and u_2 is now uncrossed – a contradiction to (i). □

Corollary 3. *If $e_0 = u_0v_0$ and $e_1 = u_1v_1$ are two crossing edges of G such that every edge of G crossing e_i is crossed only by edges incident to u_{i+1}, where $i \in \{0, 1\}$ and indices are taken modulo 2, then v_0 and v_1 are in the same connected component of H.*

Proof. Let p be the point in which e_0 and e_1 cross. For $i = 0, 1$ let S_i be the set of all edges crossing e_{i+1} between p and v_{i+1}. (All indices are taken modulo 2.) By assumption S_i is a star centered at u_i. Consider the embedding of the graph $S_0 \cup S_1$ inherited from G. By fan-planarity u_0 and u_1 are contained in the outer cell of $S_0 \cup S_1$. Moreover, every inner cell c of $S_0 \cup S_1$ has $||c|| = 4$ and thus by Corollary 2 all leaves of S_0 and S_1 are also contained in the outer cell c^* of $S_0 \cup S_1$.

We claim that no edge segment in the boundary ∂c^* of the outer cell is crossed by another edge in G. Indeed, if e' is an edge crossing some edge $e \in S_0 \cup S_1$ between the crossing of e and e_0 or e_1 and the endpoint of e different from u_0, u_1, then by assumption one endpoint of e' is u_0 or u_1 – say u_1. Moreover, since by Corollary 2 no cell c with $||c|| = 4$ contains any vertex, we have that e' crosses e_0 between p and v_0 and thus $e \in S_1$. See Figure 5(b).

We conclude that if we label the vertices of $S_0 \cup S_1$ such that their cyclic order around c^* is $u_0, u_1, v_0 = w_1, w_2, \ldots, w_k = v_1$, then for each $j \in \{1, \ldots, k-1\}$ the vertices w_j and w_{j+1} are contained in the same cell of G and hence by maximality of G joint by an uncrossed edge. See Figure 5(a) for an illustration. □

Recall that H denotes the planar subgraph of G. For convenience we refer to the closure of cells of H as the *faces of* G. The boundary of a face f is a disjoint set of (not necessarily simple) cycles of H, which we call *facial walks*. The *length of a facial walk* W, denoted by $|W|$, is the number of its edges counted with multiplicity. We remark that a facial walk may consist of only a single vertex, in which case its length is 0. See Figure 6(a) for an example.

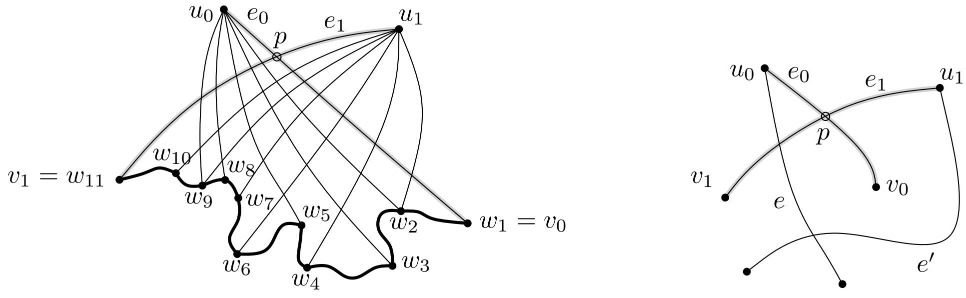

(a) (b)

Figure 5: (a) The stars S_0 and S_1 in the proof of Corollary 3 (b) If an edge e' crosses $e \in S_0$ between the crossing of e and e_1 and the endpoint of e different from u_0, and $e' \notin S_1$, then v_0 is contained in a cell c bounded by e, e' and e_1 with $||c|| = 4$.

For a face f and a facial walk W of f, we define $G(W)$ to be the subgraph of G consisting of the walk W and all edges that are drawn entirely inside f and have both endpoints on W. The set of cells of $G(W)$ that lie inside f is denoted by $C(W)$. Finally, the graph $G(W)$ is called a *sunflower* if $|W| \geq 5$ and $G(W)$ has exactly $|W|$ inner edges each of which connects two vertices at distance 2 on W. See Figure 6(b) for an example of a sunflower. We remark that for convenience we depict facial walks in our figures as simple cycles, even when there are repeated vertices or edges.

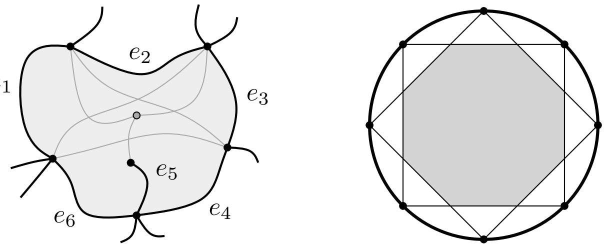

(a) (b)

Figure 6: (a) A cell of H (drawn black) is shown in gray. The boundary of the cell is the cycle $e_1, e_2, e_3, e_4, e_5, e_5, e_6$. (b) A sunflower on 8 vertices. The facial walk W is drawn thick. A cell bounded by 8 edge segments and no vertex is highlighted.

3.2 Counting the Number of Edges

We shall count the number of edges of G in three sets:

- Edges in H, that is all uncrossed edges.
- Edges in $E(G(W)) \setminus E(W)$ for every facial walk W.
- Edges between different facial walks of the same face f of G.

The edges in H will be counted in the final proof of Theorem 1 below. We start by counting the crossed edges, first within the same facial walk and afterwards between different facial walks. For convenience, let us call for a facial walk W the edges in $E(G(W)) \setminus E(W)$ and their edge segments *inner edges* and *inner edge segments* of $G(W)$, respectively.

Lemma 2. *Let W be any facial walk. If every inner edge segment of $G(W)$ bounds a cell of $G(W)$ of size 4 and no cell of $G(W)$ contains two vertices on its boundary that are not consecutive in W, then $G(W)$ is a sunflower.*

Proof. Let $v_0, \dots, v_k$ be the clockwise order of vertices around W. (In the following, indices are considered modulo $k+1$.) For any vertex v_i we consider the set of inner edges incident to v_i. Since no two non-consecutive vertices of W lie on the same cell, every v_i has at least one such edge. Moreover, note that for each edge $v_i v_{i+1}$ of W the unique cell c_i with $v_i v_{i+1}$ on its boundary has size at least 5. This implies that every v_i has indeed at least two incident inner edges. Finally, note that every inner edge is crossed, since otherwise there would be two non-consecutive vertices of W bounding the same cell of $G(W)$.

Now let us consider the clockwise first inner edge incident to v_i, denoted by e_i^1. Since an edge segment of e_i^1 bounds the cell c_i on the other side of this segment is a cell of size 4. This means that e_i^1 and the clockwise next inner edge at v_i are crossed by some edge e. By fan-planarity e crosses only edges incident to v_i. Thus each endpoint of v bounds together with v_i some cell of $G(W)$. Since only consecutive vertices of W bound the same cell of $G(W)$, this implies that $e = v_{i-1}v_{i+1}$. Since this is true for every $i \in \{0, \dots, k\}$, we conclude that $G(W)$ is a sunflower. □

Recall that $C(W)$ denotes the set of all bounded cells of $G(W)$.

Lemma 3. *For every facial walk W with $|W| \geq 3$ we have*

$$|E(G(W)) \setminus E(W)| \leq 2|W| - 5 - \sum_{c \in C(W)} \max\{0, ||c|| - 5\}.$$

Proof. Without loss of generality we may assume that W is a simple cycle. We proceed by induction on $|E(G(W))|$. As induction base we consider the case that W is a triangle. Then $G(W) = W$ and $C(W)$ consists of a single cell c with $||c|| = 6$. Thus

$$|E(G(W)) \setminus E(W)| = 0 = 2|W| - 5 - (||c|| - 5).$$

First, consider any inner edge segment e^* and the two cells $c_1, c_2 \in C(W)$ containing e^* on their boundary. If c^* denotes the set $c_1 \cup c_2$ in $G(W) \setminus e$, then

$$||c^*|| = ||c_1|| + ||c_2|| - 4$$

and thus

$$\max\{0, ||c^*|| - 5\} = \max\{0, ||c_1|| - 5\} + \max\{0, ||c_2|| - 5\} + x, \tag{1}$$

where $x = 1$ if $||c_1|| \geq 5$ and $||c_2|| \geq 5$ and $x = 0$ otherwise.

Now, we shall distinguish three cases: $G(W)$ is a sunflower, some inner edge segment is not bounded by a cell of size 4, and some cell of $G(W)$ contains two vertices on its boundary that are not consecutive in W. By Lemma 2 this is a complete case distinction.

Case 1. $G(W)$ is a sunflower. Then by definition, $G(W)$ has exactly $|W|$ inner edges. Moreover, $C(W)$ contains exactly one cell c of size greater than 4 and for that cell we have $||c|| = |W|$. Thus

$$|E(G(W)) \setminus E(W)| = |W| = 2|W| - 5 - (|W| - 5).$$

Case 2. Some edge segment e^* of some inner edge e bounds two cells c_1, c_2 of size at least 5 each. Then applying induction to the graph $G' = G(W) \setminus e$ we get

$$|E(G(W)) \setminus E(W)| = 1 + |E(G') \setminus E(W)| \leq 1 + 2|W| - 5 - \sum_{c \in C(G')} \max\{0, ||c|| - 5\}$$
$$\stackrel{(1)}{=} 1 + 2|W| - 5 - \sum_{c \in C(W)} \max\{0, ||c|| - 5\} - 1.$$

Case 3. Some cell of $G(W)$ contains two vertices u, w on its boundary that are not consecutive on W. Note that uw may or may not be an inner edge of $G(W)$. In the latter case we denote by c^* the unique cell that is bounded by u and w. In any case exactly two cells c_1, c_2 of $G(W) \cup uw$ are bounded by u and w and we have $||c^*|| = ||c_1|| + ||c_2|| - 4$, provided c^* exists.

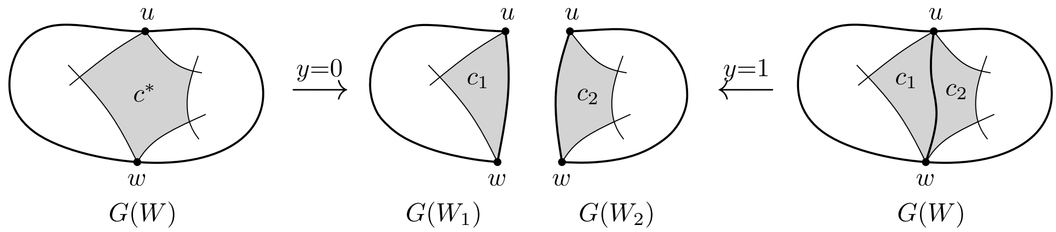

Figure 7: The graph $G(W)$ is split into two graphs $G(W_1)$ and $G(W_2)$ along two vertices u, w that are not consecutive on W but bound the same cell of $G(W)$.

We consider the two cycles W_1, W_2 in $W \cup uw$ that are different from W, such that W_1 surrounds c_1 and W_2 surrounds c_2. For $i = 1, 2$ consider $G(W_i)$, i.e., the subgraph of $G(W) \cup uw$ induced by W_i, see Figure 7. We have

$$|W| = |W_1| + |W_2| - 2,$$
$$|E(G(W)) \setminus E(W)| = |(E(G(W_1)) \setminus E(W_1))| + |(E(G(W_2)) \setminus E(W_2))| + y,$$
$$\sum_{c \in C(W)} \max\{0, ||c|| - 5\} \stackrel{(1)}{=} \sum_{c \in C(W_1)} \max\{0, ||c|| - 5\} + \sum_{c \in C(W_2)} \max\{0, ||c|| - 5\} + (1 - y),$$

where $y = 1$ if uw already was an inner edge of $G(W)$ and $y = 0$ otherwise. Now, applying induction to $G(W_1)$ and $G(W_2)$ gives the claimed bound.

□

Let us define by $C(f)$ the union of $C(W)$ for all facial walks W of f. Moreover, we partition $C(f)$ into $C_\emptyset(f)$ and $C_*(f)$, where a cell $c \in C(f)$ lies in $C_\emptyset$ if and only if $(c \setminus \partial c) \cap V(G) = \emptyset$. I.e., cells in $C_\emptyset(f)$ do not have any vertex of G in their open interior, whereas cells in $C_*(f)$ contain some vertex of G in their interior. Without loss of generality we have that for each $C_*(f)$ is either empty or contains at least one bounded cell. This can be achieved by picking a cell of G that has the maximum number of surrounding Jordan curves of the form ∂c for $c \in \bigcup_f C_*(f)$, and defining it to be in the unbounded cell of G.

Before we bound the number of edges between different facial walks of f we need one more lemma. Consider a face f of G with at least two facial walks and a cell $c \in C_*(f)$ that is inclusion-minimal. Let W_1 be the facial walk with $c \in C(W_1)$ and $W_2, \ldots, W_k$ be the facial walks that are contained in c. For $i = 1, \ldots, k$ let c_i be the cell of $G(W_i)$ that contains all walks W_j with $j \neq i$.

In particular, we have $c_1 = c$. Moreover, we call an edge between two distinct facial walks W_i and W_j a *W_iW_j-edge.*

Lemma 4. *Exactly one of $c_1, \dots, c_k$ has a vertex on its boundary.*

Proof. We proceed by proving a series of claims first.

Claim 1. *If a W_iW_j-edge and a $W_{i'}W_{j'}$-edge cross, then $\{i, j\} = \{i', j'\}$.*

Proof of Claim. Consider a W_iW_j-edge $e_0 = u_0v_0$ crossing a $W_{i'}W_{j'}$-edge $e_1 = u_1v_1$. By Corollary 1 one endpoint of e_0, say $u_0 \in W_i$, and one endpoint of e_1, say $u_1 \in W_{i'}$, are joint by an uncrossed edge. In particular, $W_i = W_{i'}$.

If, Case 1, e_0 is crossed by a second edge incident to v_1, then applying Lemma 1 gives an uncrossed edge u_0v_1, which is a contradiction to the fact that $W_{j'} \neq W_{i'}$, or an uncrossed edge v_0v_1, which implies $W_j = W_{j'}$ as desired.

Otherwise, Case 2, e_0 is crossed only by edges at u_1, and by symmetry e_1 is crossed only by edges at u_0. Applying Corollary 3 we get that v_0 and v_1 are in the same connected component of H and hence $W_j = W_{j'}$, as desired. △

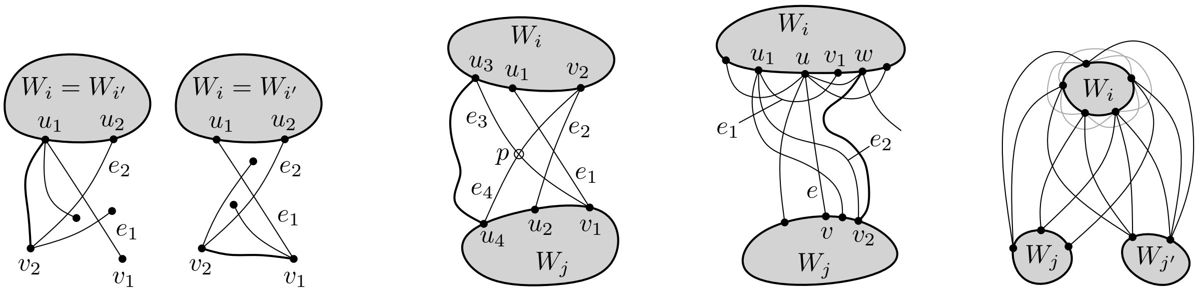

(a) (b) (c) (d)

Figure 8: (a) Case 1 in the proof of Claim 1. Illustrations of the proofs of Claim 2 (b), Claim 3 (c) and Claim 4 (d).

For a facial walk W_i a vertex $v \in W_i$ is called *open* if v lies on ∂c_i. Moreover, a vertex $v \in W_i$ is called *closed* if v is not open but there is at least one edge between v and another facial walk $W_j \neq W_i$. So every edge between distinct facial walks has endpoints that are open or closed, and by fan-planarity at least one endpoint is open.

Claim 2. *If two W_iW_j-edges cross then both have exactly one open end, which moreover are in the same facial walk.*

Proof of Claim. Let $e_1 = u_1v_1$ and $e_2 = u_2v_2$ be two crossing W_iW_j-edges. Assume for the sake of contradiction that e_1 has an open endpoint $u_1 \in W_i$ and e_2 has an open endpoint $u_2 \in W_j$. We consider the edges $e_3 = u_3v_1$ and $e_4 = u_4v_2$ that are incident to v_1 and v_2 respectively, cross each other and whose crossing point p is furthest away from v_1 and v_2. See Figure 8(b) for an illustration. Note that possibly $e_1 = e_3$ and/or $e_4 = e_2$.

Now u_3 is not in W_j because u_1 is an open endpoint and u_4 is not in W_i because u_2 is an open end. Hence by Claim 1 $u_3 \in W_i$ and $u_4 \in W_j$. Moreover, by Lemma 1 u_3u_4 is an uncrossed edge of G – a contradiction to the fact that W_i and W_j are distinct facial walks. △

Claim 2 implies that every edge between different facial walks has exactly one open endpoint and one closed end, because every such edge with two open endpoints would be crossed by some other edge between two facial walks.

Claim 3. *If a W_iW_j-edge has a closed endpoint $u \in W_i$ and w is the counterclockwise next open or closed vertex of W_i after u, then there exists a W_iW_j-edge incident to w with open endpoint in W_j.*

Proof of Claim. Let $e = uv$ be a W_iW_j-edge that has a closed endpoint $u \in W_i$. By fan-planarity v is an open vertex of W_j. Consider the crossing of e closest to u and let $e_1 = u_1v_1$ be the crossing edge. Clearly, e_1 is an edge from $G(W_i)$, where without loss of generality v_1 comes counterclockwise after u in W_i. Further assume without loss of generality that e is the W_iW_j-edge at u whose crossing with e_1 is closest to v_1. If e is not crossed between v and its crossing with e_1 then we can draw a W_iW_j-edge between v and w that is not crossed by any edge between facial walks and we are done.

Otherwise, if e is crossed by some edge e_2 between its crossing with e_1 and v, then by fan-planarity e_2 is incident to u_1 or v_1. Moreover, by Claim 1 and Claim 2 e_2 has a closed endpoint in W_i and an open endpoint in W_j. Thus if e_2 is incident to v_1, then we have found the desired W_iW_j-edge. So assume that $e_2 = u_1v_2$ for some $v_2 \in W_j$. Moreover, let e_2 be the W_iW_j-edge whose crossing with e is closest to u. We refer to Figure 8(c) for an illustration.

Because e_2 has a closed endpoint $u_1 \in W_i$ every edge crossing e_1 or e_2 endpoints in u. Thus by the choice of e we conclude that e_2 is not crossed between v_2 and its crossing with e and that e_1 is not crossed between its crossing with e and the next vertex or edge in $G(W_i)$. Moreover, by the choice of e_2 the edge e is not crossed between its crossings with e_2 and e_1. Thus we can draw a W_iW_j-edge from v_2 to w. △

Claim 3 together with Claim 2 implies that on each facial walk every closed vertex is followed by another closed vertex. In particular, the facial walks come in two kinds, one with open vertices only and one with closed vertices only. We remark one can show that, if W_i has only closed vertices, then $G(W_i)$ is a sunflower.

Claim 4. *Every facial walk with only closed vertices has edges to exactly one facial walk with only open vertices.*

Proof of Claim. Assume for the sake of contradiction that facial walk W_i with only closed vertices has edges to two different facials walks $W_j, W_{j'}$ with only open vertices. Claim 3 implies that if some closed vertex of W_i has an edge to W_j, then every closed vertex of W_i has an edge to W_j, and the same is true for $W_{j'}$. Hence, each of the at least three closed vertices in W_i has edge to W_j and $W_{j'}$, which implies that some W_iW_j-edge and some $W_iW_{j'}$-edge must cross, see Figure 8(d). (Indeed, if any two such edges would not cross, then contracting W_j and $W_{j'}$ into a single point each and placing a new vertex in the middle of W_i with an edge to every closed vertex in W_i would give a planar drawing of $K_{3,3}$.) Thus by Claim 1 we have $W_j = W_{j'}$ – a contradiction to our assumption. △

We are now ready to prove that at most one facial walk has open vertices. Recall that by Claim 3 every facial walk is of one of two kinds: only open vertices or only closed vertices. Moreover, by fan-planarity and Claim 2 no edge runs between two facial walks of the same kind. We consider a bipartite graph F whose black and white vertices correspond to facial walks of the first and second kind, respectively, and whose edges correspond to pairs W_i, W_j of facial walks for which there is at least one W_iW_j-edge in G. Since G is connected, F is connected, and by Claim 4 every white vertex is adjacent to exactly one black vertex. This means that F is a star and has exactly one black vertex, which concludes the proof. □

Having Lemma 4 we can now bound the number of W_iW_j-edges. Recall that $W_1, \dots, W_k$ denote the facial walks for the fixed face f of G, and that for $i = 1, \dots, k$ we denote by c_i the cell of $G(W_i)$ containing all W_j with $j \neq i$.

Lemma 5. *The number of edges between $W_1, \dots, W_k$ is at most*

$$4(k-2) + \sum_{i=1}^{k} ||c_i||.$$

Proof. By Lemma 4 exactly one of $c_1, \dots, c_k$ has vertices on its boundary, say W_1. Let U be the set of vertices on the boundary of c_1. For a vertex $u \in U$ and an index $i \in \{2, \dots, k\}$ we call an edge between u and W_i a *uW_i-edge.* We define a bipartite graph J as follows. One bipartition class is formed by the vertices in U. In the second bipartition class there is one vertex w_i for each facial walk W_i, $i = 1, \dots, k$. A vertex $u \in U$ is connected by an edge to w_i if and only if $i = 1$ or $i \geq 2$ and there is a uW_i-edge.

Claim 5. *The graph J is planar.*

Proof of Claim. We consider the following embedding of J. Afterwards we shall argue that this embedding is indeed a plane embedding. So take the position of every vertex $u \in U$ from the fan-planar embedding of G. For $i \geq 2$, we consider the drawing of W_i in the embedding of G, for each edge between a vertex $u \in U$ and the vertex w_i in J we take the drawing of one uW_i-edge in G, and then contract the drawing of W_i into a single point – the position for vertex w_i. Finally, we place the last vertex w_1 outside the cell c_1 and connect w_1 to each $u \in U$ in such a way that these edges do not cross any other edge in J. See Figure 9(a) for an illustrating example.

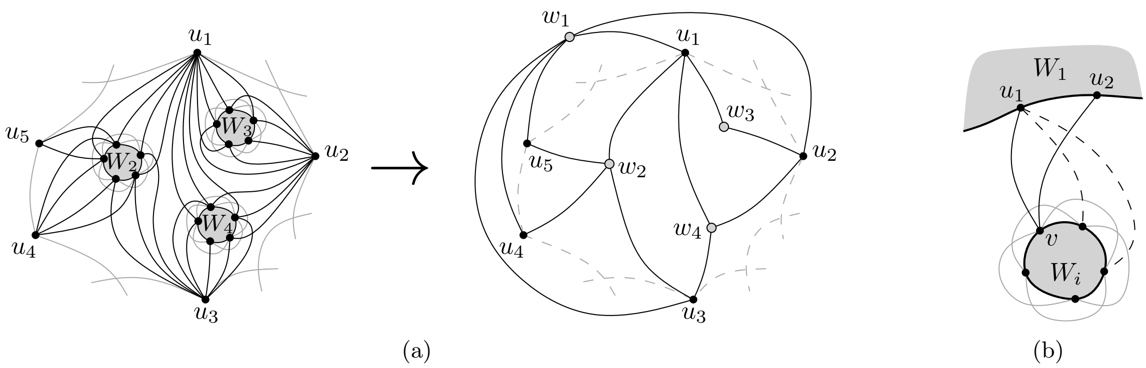

Figure 9: (a) Obtaining the graph J. (b) The contradiction in Claim 6.

Now the resulting drawing of J contains crossing edges only if a uW_i-edge crosses a $u'W_{i'}$-edge in G. However, by Lemma 4 the cells $c_2, \dots, c_k$ have no vertices on their boundary. Consequently, for each $i = 2, \dots, k$ every uW_i-edge crosses an edge of $G(W_i)$. Now if a uW_i-edge e would cross a $u'W_{i'}$-edge with $u \neq u'$ and $i \neq i'$, then e would be crossed by two independent edges, contradicting the fan-planarity of G. △

Since J is a planar bipartite graph with bipartition classes of size $|U|$ and k we have

$$|E(J)| = \sum_{i=1}^{k} \deg_J(w_i) \leq 2(|U| + k) - 4.$$

Claim 6. *For each $i = 2, \ldots, k$ the number of uW_i-edges is at most*

$$||c_i|| + 2\deg_J(w_i).$$

Proof of Claim. Consider the vertices on W_i and the set $U' \subseteq U$ of vertices on W_1 that have a neighbor on W_i. For each $u \in U'$ the uW_i-edges form a consecutive set in the cyclic ordering of edges around u. Since not every edge at u is a uW_i-edge (at least one edge endpoints in W_1) we obtain a linear order on the uW_i-edges going counterclockwise around u.

Now we claim that when we remove for each $u \in U'$ the last two uW_i-edges in the linear order for u, then every vertex v in W_i is the endpoint of at most one uW_i-edge. Indeed, if after the edges have been removed two vertices $u_1, u_2 \in U'$ have a common neighbor v on W_i, then at least two u_1W_i-edges cross the edge u_2v (or the other way around). However, not both these edges can endpoint at the same vertex on W_i, and thus u_2v is crossed by two independent edges, one u_1W_i-edge and one edge in $G(W_i)$ – a contradiction to the fan-planarity of G. So the number of uW_i-edges is at most $2|U'| + |W_i| = ||c_i|| + 2\deg_J(w_i)$. △

We can now bound the total number of uW_i-edges with $i \geq 2$ as follows.

$$\begin{aligned}
\sum_{i=2}^{k} \#uW_i\text{-edges} &\leq \sum_{i=2}^{k} (||c_i|| + 2\deg_J(w_i)) \\
&= \sum_{i=2}^{k} ||c_i|| + 2|E(J)| - 2\deg_J(w_1) \\
&\leq \sum_{i=2}^{k} ||c_i|| + 4(|U| + k) - 8 - 2|U| \\
&= \sum_{i=2}^{k} ||c_i|| + 2|U| + 4(k-2) \leq \sum_{i=2}^{k} ||c_i|| + ||c_1|| + 4(k-2)
\end{aligned}$$

□

We continue by bounding the total number of crossed edges of G that are drawn inside a fixed face f of G. To this end let k_f be the number of distinct facial walks of f and $|f|$ be the sum of lengths of facial walks of f, i.e., $|f| = \sum_{W \text{ facial walk of } f} |W|$.

Lemma 6. *The number of edges inside f is at most*

$$2|f| + 5(k_f - 2) - \sum_{c \in C_\emptyset(f)} \max\{0, ||c|| - 5\}.$$

Proof. We do induction on k_f.

First let $k_f = 1$, i.e., the face f is bounded by a unique facial walk W. Then by Lemma 3 there are at most $2|W| - 5 - \sum_{c \in C(W)} \max\{0, ||c|| - 5\}$ edges inside f. With $|W| = |f|$ and $C_\emptyset(f) = C(W)$ this gives the claimed bound.

Now assume that $k_f \geq 2$, i.e., the face f has $k = k_f$ distinct facial walks $W_1, \ldots, W_k$. Let c_1 be an inclusion-minimal cell in $(C(W_1) \cup \cdots \cup C(W_k)) \setminus C_\emptyset(f)$. Without loss of generality let W_1 be the facial walk with $c_1 \in C(W_1)$ and $W_2, \ldots, W_j$ be the facial walks of f that lie inside c_1. In particular we have $2 \leq j \leq k$. Let G' be the graph that is obtained from G after removing all vertices that lie inside c_1. We consider G' with its fan-planar embedding inherited from G. Clearly, the face f' in G' corresponding to f in G has exactly $k-(j-1)<k$ facial walks and we have

$$|f| = |f'| + |W_2| + \cdots + |W_j|.$$

For $i = 2, \ldots, j$ we denote by c_i the cell of $G(W_i)$ containing W_1. Moreover, let $C = C(W_2) \cup \cdots \cup C(W_j)$. Then

$$C_\emptyset(f) = (C_\emptyset(f') \cup C) \setminus \{c_1, c_2, \ldots, c_j\}.$$

Further we partition the edges inside f into three disjoint sets E_1, E_2, E_3 as follows:

- The edges in E_1 are precisely the edges of G' inside f'.
- The edges in E_2 are precisely the edges of G between W_1 and $W_2 \cup \cdots \cup W_j$.
- $E_3 = (E(G(W_2)) \setminus E(W_2)) \cup \cdots \cup (E(G(W_j)) \setminus E(W_j))$.

Now by induction hypothesis we have

$$|E_1| \leq 2|f'| + 5(k-j-1) - \sum_{c \in C_\emptyset(f')} \max\{0, ||c|| - 5\}.$$

By Lemma 5 we have

$$|E_2| \leq \sum_{i=1}^{j} ||c_i|| + 4(j-2) \leq \sum_{i=1}^{j} \max\{0, ||c_i|| - 5\} + 9j - 8.$$

By Lemma 3 we have

$$|E_3| \leq 2(|W_2| + \cdots + |W_j|) - 5(j-1) - \sum_{c \in C} \max\{0, ||c|| - 5\}.$$

Plugging everything together we conclude that the number of edges of G inside f is at most

$$\begin{aligned}
|E_1 \dot\cup E_2 \dot\cup E_3| &\leq 2|f'| + 5(k-j-1) - \sum_{c \in C_\emptyset(f')} \max\{0, ||c|| - 5\} \\
&\quad + \sum_{i=1}^{j} \max\{0, ||c_i|| - 5\} + 9j - 8 \\
&\quad + 2(|W_2| + \cdots + |W_j|) - 5(j-1) - \sum_{c \in C} \max\{0, ||c|| - 5\} \\
&= 2|f| + 5(k-2) - (j-2) - \sum_{c \in C_\emptyset(f)} \max\{0, ||c|| - 5\} \\
&\leq 2|f| + 5(k_f - 2) - \sum_{c \in C_\emptyset(f)} \max\{0, ||c|| - 5\},
\end{aligned}$$

which concludes the proof. □

Note that Lemma 6 implies that inside a face f of H there are at most $2|f| + 5(k_f - 2)$ edges. Having this, we are now ready to prove our main theorem. For convenience we restate it here.

Theorem 1. *Every simple topological graph G on $n \geq 3$ vertices with neither configuration I nor configuration II has at most $5n - 10$ edges. This bound is tight for $n \geq 20$.*

Proof. Consider a fan-planar graph $G = (V, E)$ on n vertices with properties (i) and (ii). Let H be the spanning subgraph of G on all uncrossed edges. In particular

$$V(H) = V(G).$$

Let us denote by $F(H)$ the set of all faces of H. Since every edge $e \in E(H)$ appears either exactly once in two distinct facial walks or exactly twice in the same facial walk, we have

$$\sum_{f \in F(H)} |f| = 2|E(H)|.$$

Further we denote by k_f the number of facial walks for a given face f, and by $CC(H)$ the number of connected components of H. Since a face with k facial walks gives rise to k connected components of H, we have

$$\sum_{f \in F(H)} (k_f - 1) = CC(H) - 1.$$

Hence we conclude

$$\begin{aligned}
|E(G)| &\overset{\text{Lemma 6}}{\leq} |E(H)| + \sum_{f \in F(H)} (2|f| + 5(k_f - 2)) \\
&= |E(H)| + 2 \sum_{f \in F(H)} |f| + 5 \sum_{f \in F(H)} (k_f - 1) - 5|F(H)| \\
&= 5|E(H)| + 5CC(H) - 5|F(H)| - 5 = 5|V(H)| - 10,
\end{aligned}$$

where the last equation is Euler's formula for the plane embedded graph H. With $|V(H)| = |V(G)| = n$ this concludes the proof. ☐

4 Discussion

We have shown that every simple n-vertex graph without configurations I and II has at most $5n-10$ edges. Of course, if we allow G to have parallel edges or loops, there could be arbirarily many edges, even if the drawing of G is planar. However, if we allow only *non-homeomorphic parallel edges* and only *non-trivial loops*, then G has a bounded number of edges. Here, two parallel edges are non-homeomorphic and a loop is non-trivial if the bounded component of the plane after the removal of both parallel edges, respectively the loop, contains at least one vertex of G. Note for instance that Euler's formula still holds for plane graphs with non-homeomorphic parallel edges and non-trivial loops, and that in this case every face still has length at least 3. Therefore any such plane graph with n vertices still has at most $3n - 6$ edges. We strongly conjecture that our $5n - 10$ bound also holds if non-homeomorphic parallel edges and non-trivial loops are allowed.

Another relaxation would be to allow non-simple topological graphs, i.e., to allow edges to cross more than once and incident edges to cross. It would be interesting to see whether there is an n-vertex non-simple fan-planar graph with strictly more than $5n - 10$ edges. However, let us remark that if we allow both, non-simple drawings and non-homeomorphic parallel edges, then there are 3-vertex topological graph with arbitrarily many edges. Let us simply refer to Figure 10(a) for such an example. The idea is to start with an edge e_1 from u to v, and edge e_i starts clockwise next to e_{i-1} at u goes in parallel with e_{i-1} until e_{i-1} endpoints at v, where e_i goes a little further surrounding vertex w once and then ending at v. This way no two such parallel edges are homeomorphic.

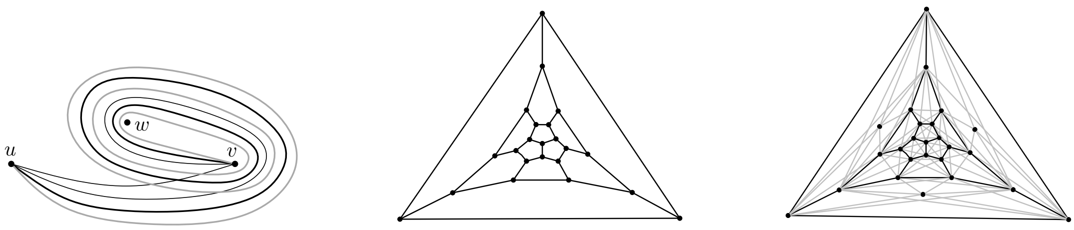

(a) (b) (c)

Figure 10: (a) A topological non-simple fan-planar graph with arbitrarily many edges. (b) The modified dodecahedral graph without the extensions and (c) fully extended to obtain $5n - 11$ straight-line edges.

Also, one can relax fan-planarity to *k-fan-planarity* for some $k \geq 1$, where every edge may only be crossed by k fan-crossings. We remark that a simple probabilistic argument shows that for fixed k every n-vertex k-fan-planar graph has only linearly many edges, see Lemma 2.9 in [2]. However, exact bounds are not known.

It can also be interesting to consider strengthenings of fan-planar graphs, e.g., to consider straight-line fan-planar embeddings. Note that the dodecahedral graph with pentagrams which was a tight example of the $5n - 10$ bound, can be extended as follows to obtain a straight-line fan-planar graph with $5n - 11$ edges: Replace one vertex of the dodecahedron by a single triangle, which is used as the outer face. Draw the planar graph with convex faces such that all (additional) edges can be drawn straightline without producing unnecessary crossings, cf. Figure 10(b). The 3 adjacent pentagons now converted to hexagons are extended by 2-hops and spokes as explained in Proposition 1, i.e., by one additional vertex and 12 edges each. We do not suspect that an n-vertex straight-line fan-planar graph can have $5n - 10$ edges.

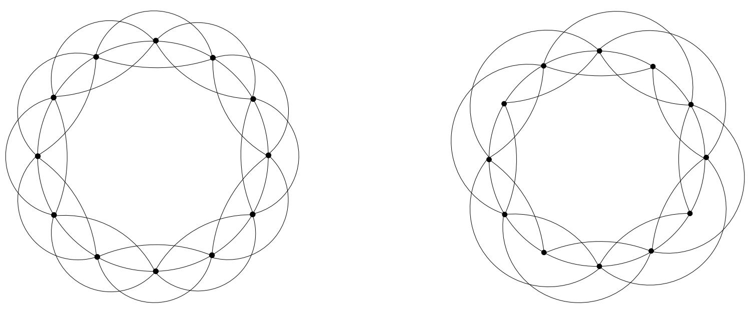

(a) (b)

Figure 11: (a) An edge-maximal fan-planar graph with non-homeomorphic parallel edges on $3n$ edges. (b) An edge-maximal simple fan-planar graph on $\frac{8}{3}n$ edges.

Finally, one is usually also interested in edge-maximal topological graphs with as *few* edges as possible. In our case we can construct edge-maximal fan-planar graphs on no more than $3n$ edges if parallel edges are allowed (Figure 11(a)) and no more than $\frac{8}{3}n$ edges if parallel edges are not allowed (Figure 11(b)). We suspect these examples to be best-possible.

Let us summarize some possible research directions.

Problems. *Each of the following is open.*

P1: *What is the maximum number of edges in a simple topological graph G with forbidded configuration I, but where configuration II is allowed?*

P2: *Is there an n-vertex simple fan-planar graph with non-homeomorphic parallel edges and/or non-trivial loops with strictly more than* $5n - 10$ *edges?*

P3: *Does the* $5n - 10$ *upper bound also hold for non-simple fan-planar graphs?*

P4: *For* $k \geq 2$ *what is the largest number of edges in an n-vertex k-fan-planar graph?*

P5: *Prove that the* $5n - 11$ *bound is tight for straight-line fan-planar embeddings similar to the* $4n - 9$ *bound for straight-line embedded* 1*-planar graphs [12].*

P6: *How many edges has an n-vertex edge-maximal graph without configurations I and II at least?*

Acknowledgements
This work has been started at the Bertinoro Workshop on Graph Drawing 2013. The authors want to thank some referees for valuable hints especially concerning the relation to the previous work on grids in topological graphs.

References

[1] E. Ackerman. On the maximum number of edges in topological graphs with no four pairwise crossing edges. *Discrete & Computational Geometry*, 41(3):365–375, 2009.

[2] E. Ackerman, J. Fox, J. Pach, and A. Suk. On grids in topological graphs. In *Proceedings of the 25th Annual Symposium on Computational Geometry*, pages 403–412. ACM, 2009.

[3] E. Ackerman and G. Tardos. On the maximum number of edges in quasi-planar graphs. *J. Comb. Theory, Ser. A*, 114(3):563–571, 2007.

[4] P. K. Agarwal, B. Aronov, J. Pach, R. Pollack, and M. Sharir. Quasi-planar graphs have a linear number of edges. *Combinatorica*, 17(1):1–9, 1997.

[5] M. J. Alam, F.-J. Brandenburg, and S. Kobourov. Straight-line drawings of 3-connected 1-planar graphs. In *Graph Drawing*, pages 83–94, 2013.

[6] E. N. Argyriou, M. A. Bekos, M. Kaufmann, and A. Symvonis. Geometric rac simultaneous drawings of graphs. *J. Graph Algorithms Appl.*, 17(1):11–34, 2013.

[7] C. Auer, C. Bachmaier, F.-J. Brandenburg, A. Gleißner, M. T. Goodrich, K. Hanauer, D. Neuwirth, and J. Reislhuber. Recognizing outer 1-planar graphs in linear time. In *Graph Drawing*, pages 107–118, 2013.

[8] F.-J. Brandenburg, D. Eppstein, A. Gleißner, M. T. Goodrich, K. Hanauer, and J. Reislhuber. On the density of maximal 1-planar graphs. In *Graph Drawing*, pages 327–338, 2012.

[9] P. Brass, W. O. J. Moser, and J. Pach. *Research problems in discrete geometry.* Springer, 2005.

[10] O. Cheong, S. Har-Peled, H. Kim, and H.-S. Kim. On the number of edges of fan-crossing free graphs. In *ISAAC*, pages 163–173, 2013.

[11] H. De Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.

[12] W. Didimo. Density of straight-line 1-planar graph drawings. *Inf. Process. Lett.*, 113(7):236–240, 2013.

[13] W. Didimo, P. Eades, and G. Liotta. Drawing graphs with right angle crossings. *Theor. Comput. Sci.*, 412(39):5156–5166, 2011.

[14] P. Eades, S.-H. Hong, N. Katoh, G. Liotta, P. Schweitzer, and Y. Suzuki. A linear time algorithm for testing maximal 1-planarity of graphs with a rotation system. *Theoretical Computer Science*, 513:65–76, 2013.

[15] P. Eades and G. Liotta. Right angle crossing graphs and 1-planarity. *Discrete Applied Mathematics*, 161(7-8):961–969, 2013.

[16] J. Fox, J. Pach, and A. Suk. The number of edges in k-quasi-planar graphs. *SIAM J. Discrete Math.*, 27(1):550–561, 2013.

[17] A. Grigoriev and H. L. Bodlaender. Algorithms for graphs embeddable with few crossings per edge. *Algorithmica*, 49(1):1–11, 2007.

[18] S.-H. Hong, P. Eades, N. Katoh, G. Liotta, P. Schweizer, and Y. Suzuki. A linear-time algorithm for testing outer-1-planarity. pages 71–82, 2013.

[19] S.-H. Hong, P. Eades, G. Liotta, and S.-H. Poon. Fárys theorem for 1-planar graphs. *Computing and Combinatorics*, page 335, 2012.

[20] J. Pach, R. Pinchasi, M. Sharir, and G. Tóth. Topological graphs with no large grids. *Graphs and Combinatorics*, 21(3):355–364, 2005.

[21] J. Pach and G. Tóth. Graphs drawn with few crossings per edge. *Combinatorica*, 17(3):427–439, 1997.

[22] G. Ringel. Ein Sechsfarbenproblem auf der Kugel. *Abh. des Mathematischen Seminars Hamburg*, 29:107–117, 1965.

[23] W. Schnyder. Embedding planar graphs on the grid. In *Proceedings of the first Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 138–148. Society for Industrial and Applied Mathematics, 1990.

[24] A. Suk and B. Walczak. New bounds on the maximum number of edges in k-quasi-planar graphs. *CoRR*, abs/1309.0395, 2013.

[25] Y. Suzuki. Re-embeddings of maximum 1-planar graphs. *SIAM J. Discrete Math.*, 24(4):1527–1540, 2010.

[26] G. Tardos and G. Tóth. Crossing stars in topological graphs. *SIAM Journal on Discrete Mathematics*, 21(3):737–749, 2007.

[27] C. Thomassen. Rectilinear drawings of graphs. *Journal of Graph Theory*, 12(3):335–341, 1988.

[28] W. T. Tutte. How to draw a graph. *Proc. London Math. Soc*, 13(3):743–768, 1963.

Articles on Intersection Representations

Journal of Graph Algorithms and Applications
http://jgaa.info/ vol. 17, no. 4, pp. 475–494 (2013)
DOI: 10.7155/jgaa.00300

Planar Graphs as VPG-Graphs

Steven Chaplick[1] *Torsten Ueckerdt*[2]

[1]Department of Applied Mathematics, Charles University
[2]Department of Mathematics, Karlsruhe Institute of Technology

Abstract

A graph is B_k-VPG when it has an intersection representation by paths in a rectangular grid with at most k bends (turns). It is known that all planar graphs are B_3-VPG and this was conjectured to be tight. We disprove this conjecture by showing that all planar graphs are B_2-VPG. We also show that the 4-connected planar graphs constitute a subclass of the intersection graphs of Z-shapes (i.e., a special case of B_2-VPG). Additionally, we demonstrate that a B_2-VPG representation of a planar graph can be constructed in $O(n^{3/2})$ time. We further show that the triangle-free planar graphs are contact graphs of: L-shapes, Γ-shapes, vertical segments, and horizontal segments (i.e., a special case of contact B_1-VPG). From this proof we obtain a new proof that bipartite planar graphs are a subclass of 2-DIR.

Submitted:	Reviewed:	Revised:	Accepted:	Final:
December 2012	March 2013	April 2013	May 2013	July 2013
		Published: July 2013		
	Article type: Regular paper		Communicated by: W. Didimo and M. Patrignani	

An extended abstract of this paper was presented at the *20th International Symposium on Graph Drawing*, in Redmond, USA, in September 2012 [7].

The research of the first author was supported by NSERC and partially by GraDR EUROGIGA project No. GIG/11/E023. The research of the second author was supported by GraDR EUROGIGA project No. GIG/11/E023.

E-mail addresses: chaplick@kam.mff.cuni.cz (Steven Chaplick) torsten.ueckerdt@kit.edu (Torsten Ueckerdt)

1 Introduction

Planar graphs have a long history of being described as geometric intersection (and contact) graphs; i.e., for a planar graph G, each vertex can be mapped to a geometric object O_v such that (u, v) is an edge of G if and only if O_v and O_u intersect.[1] Two well-known results of this variety are that: every planar graph is an intersection graph of curves in the plane [12] (1978), and every planar graph is a contact graph of discs in the plane [21] (1936).

In this paper we consider representations of planar graphs as the intersection and contact graphs of restricted families of curves in the plane. The most general class of intersection graphs of curves in the plane is the class of *string graphs*. Formally, a graph $G = (V, E)$ is *STRING* if and only if each $v \in V$ can be associated with a curve c_v in the plane such that for every pair $u, v \in V$, $(u, v) \in E$ if and only if c_u and c_v intersect. STRING was first considered regarding thin film RC-circuits [27].

Perhaps the most significant result describing planar graphs as intersection graphs of curves is the recent proof of Scheinerman's conjecture that all planar graphs are *segment graphs (SEG)*; i.e., the intersection graphs of line segments in the plane. Scheinerman conjectured this in his Ph.D. thesis (1984) [26], and it was proven in 2009 by Chalopin and Gonçalves [5]. Leading up to this result were several partial results. Bipartite planar graphs were the first subclass shown to be intersection graphs of line segments having two distinct slopes (2-DIR) [10, 4]. This was followed by triangle-free planar graphs being shown to be intersection graphs of line segments having three distinct slopes (3-DIR) [8]. It has also been proven that segment graphs include every planar graph that can be 4-colored so that no separating cycle uses all four colors [9]. Planar graphs were also shown to be representable by curves in the plane where each pair of curves intersect in at most one point (i.e., only "simple" intersections are allowed) [6] – the proof of Scheinerman's conjecture was a strengthening of this result. The early work on this topic led West to conjecture that every planar graph is an intersection graph of line segments in four directions (4-DIR) [31].

Segment graphs have been generalized to k-segment graphs (k-SEG) where each vertex is represented by a piecewise linear curve consisting of at most k segments [23]. Interestingly, a very recent result is that all planar graphs are contact 2-SEG [1]. In this context one may now consider k-SEG where the segments of the piecewise linear curves have a bounded number of slopes. Recently, Asinowski et al. [3] introduced the class of vertex intersection graphs of paths in a rectangular grid (VPG); equivalently, VPG is the set of intersection graphs of axis-aligned rectilinear curves in the plane (or $\bigcup_{k\geq 1}$ k-SEG where each segment is either vertical or horizontal). They prove that VPG and STRING are the same graph class (this was known previously as a folklore result). Also, they focus on the subclasses which are obtained when each path in the representation has at most k bends (turns) and they refer to such a subclass as B_k-VPG (i.e., this is $(k + 1)$-SEG with two slopes). Several relationships between existing

[1] In the case of contact representations, objects may only "touch" each other, but not "cross over" each other.

graph classes and the B_k-VPG classes were observed. For example, every planar graph is B_3-VPG (this was also conjectured to be tight) and every circle graph is B_1-VPG. In other words, planar graphs are 4-SEG where the segments only have two distinct slopes. This result follows from the fact that every planar graph has a representation by a T-contact system [11] and each T-shape can be simulated by a rectilinear curve with three bends.

In this paper we present the following results. Our main contribution is that every planar graph is B_2-VPG (disproving the conjecture of Asinowski et al. [3]). This result consists of the following interesting components. We first demonstrate that every 4-connected planar graph is the intersection graph of Z-shapes (i.e., a special case of B_2-VPG). This result is extended to show that every planar graph is B_2-VPG (this extension involves the additional use of C-shapes – i.e., it uses the full capability of B_2-VPG) and that a B_2-VPG representation of a planar graph can be constructed in $O(n^{3/2})$ time. The secondary contribution of this paper is that every triangle-free planar graph is a contact graph of: L-shapes, Γ-shapes, vertical segments, and horizontal segments (i.e., it is a specialized contact B_1-VPG graph). We show how to construct such a contact representation in linear time. Moreover, if the input is bipartite then each path is a horizontal or vertical segment. In particular, we obtain a new proof that planar bipartite graphs are 2-DIR. Interestingly, the class of contact segment graphs has recently been shown to be the same as the class of contact B_1-VPG graphs [20].

2 Preliminaries

A grid path (a path in the plane square grid) consists of *horizontal and vertical segments* that appear alternatingly along the path. Every horizontal segment has a *left endpoint* and a *right endpoint*, and every vertical segment an *upper endpoint* and a *lower endpoint* in the obvious meaning. A path is a *k-bend path* if it has k bends, i.e., k points that are the endpoints of a horizontal and a vertical segment. Equivalently, k-bend paths are those with precisely $k + 1$ segments.

A *B_k-VPG representation* of a graph G is a set of grid paths (one for each vertex) with at most k bends such that two paths intersect if and only if the corresponding vertices are adjacent in G. For every vertex v we denote the corresponding grid path in a given B_k-VPG representation by $\mathbf{v}$. Consequently a B_k-VPG representation of a graph G is denoted by $\mathbf{G}$. A graph is called *B_k-VPG* if it has a B_k-VPG representation.

3 Planar Graphs are B_2-VPG

In this section we show that every planar graph G has a B_2-VPG representation. We fix any plane embedding of G and assume without loss of generality that G is a maximally planar graph, i.e., all faces are triangular. To achieve this we may put a dummy vertex into each face of G and triangulate it. In a B_2-VPG representation of this graph the paths corresponding to dummy vertices may be removed to obtain a B_2-VPG representation of G.

Our construction of the B_2-VPG representation of the maximally planar graph G relies on two well-known concepts. Using the separation tree T_G of G, we show in Section 3.1 how to divide G into its 4-connected maximally planar subgraphs. Each such subgraph, if we remove one outer edge, has a rectangular dual, i.e., a contact representation with axis-aligned rectangles. In Section 3.2 we show how to construct a B_2-VPG representation from a rectangular dual. In particular we will convert each rectangle to a Z-shaped path by choosing "part" of the top of it, the complementary "part" of the bottom of it and connecting them via a vertical segment. In Section 3.3 we put the obtained representations of all 4-connected maximally planar subgraphs of G together to obtain a B_2-VPG representation of our graph. The same method has been used to prove that every planar graph is a B_4-EPG graph, where EPG stands for emphedge-intersecting paths in the grid [18].

3.1 Separation Tree

A triangle Δ in a graph is a triple of pairwise adjacent vertices. We say that a triangle is *separating* when its removal disconnects the graph. Also, in a maximally planar graph G a triangle Δ is said to be *non-empty* when at least one vertex of G lies inside the bounded region inscribed by Δ. Notice that every separating triangle is non-empty. In fact, each non-empty triangle is either the outer triangle or separating.

We say that a triangle Δ_1 *is contained in a triangle* Δ_2, denoted by $\Delta_1 \sqsubset \Delta_2$, if the bounded region enclosed by Δ_1 is strictly contained in the one enclosed by Δ_2. For example, the outer triangle contains every triangle in the graph (except itself), and no triangle in G is contained in an inner facial triangle.

Definition 1 ([28]) *The* separation tree *of G is the rooted tree T_G whose vertices are the non-empty triangles in G, with Δ being a descendant of Δ' if and only if Δ is contained in Δ'.*

The separation tree has been introduced by Sun and Sarrafzadeh [28]. The root of T_G is the outer triangle. For every non-empty triangle Δ we define H_Δ to be the unique 4-connected maximally planar subgraph of G that contains Δ and at least one vertex of G that lies inside Δ. Equivalently, H_Δ is the union of Δ and all triangles contained in Δ but not contained in any triangle that itself is contained in Δ; i.e., $H_\Delta = \Delta \cup \left(\bigcup_{\Delta' \sqsubset \Delta \text{ and } \nexists \Delta'' : \Delta' \sqsubset \Delta'' \sqsubset \Delta} \Delta'\right)$.

Theorem 1 ([28]) *The separation tree of G and all subgraphs H_Δ can be computed in $\mathcal{O}(n^{3/2})$.*

3.2 Rectangular Duals

Throughout this section let H be a triangulation of the 4-gon, i.e., H is a plane graph with quadrangular outer face and solely triangular inner faces. Such graphs are also known as irreducible triangulations of the 4-gon. We denote the outer vertices by T, R, B, L in this clockwise order around the outer face.

Definition 2 *A* rectangular dual *of H is a set of $|V(H)|$ non-overlapping axis-aligned rectangles in the plane (one for each vertex) such that every edge of H corresponds to a non-trivial overlap of the boundaries of the corresponding rectangles.*

The rectangle corresponding to a vertex v is denoted by $R(v)$. In every rectangular dual the rectangles $R(T)$, $R(B)$, $R(L)$ and $R(R)$ that correspond to the outer vertices of H inscribe a rectangular hole that contains all the remaining rectangles. We assume without loss of generality that $R(T)$, $R(B)$, $R(L)$ and $R(R)$ are laid out as in Fig. 1 a), i.e., the bottom side of $R(T)$ forms the top side of the hole, the left side of $R(R)$ forms the right side of the hole, and so on.

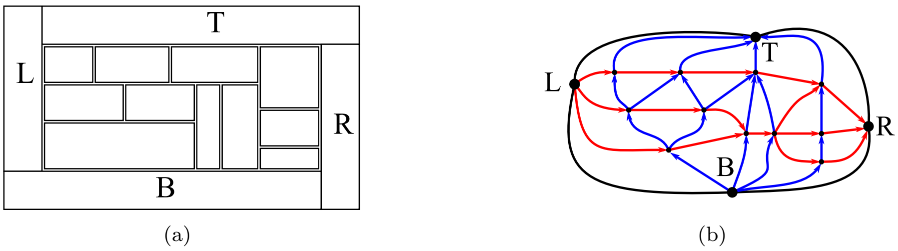

Figure 1: (a) A rectangular dual; and (b) its transversal structure.

Rectangular duals have been considered several times independently in the literature [30, 24, 22, 29, 25]. In particular, the following theorem is well-known.

Theorem 2 *A triangulation of a 4-gon admits a rectangular dual if and only if it is 4-connected, i.e., contains no non-empty triangle.*

We define here *transversal structures* as introduced by Fusy [14], which were independently considered by He [17] under the name *regular edge labelings*. For a nice overview about regular edge labelings and their relations to geometric structures we refer to the introductory article by D. Eppstein [13].

Definition 3 (Fusy [14]) *A* transversal structure *of a triangulation H with outer vertices T, L, B, R is a coloring and orientation of the inner edges of H with colors red and blue such that:*

(i) All edges at T are incoming and blue, all edges at B are outgoing and blue, all edges at R are incoming and red, all edges at L are outgoing and red.

(ii) Around each inner vertex v the edges appear in the following clockwise cyclic order: One or more incoming red edges, one or more outgoing blue edges, one or more outgoing red edges, one or more incoming blue edges.

We denote a transversal structure by (E_r, E_b)*, where* E_r *and* E_b *is the set of red and blue edges, respectively.*

We obtain a transversal structure from any rectangular dual of H as follows. If the right side of a rectangle $R(u)$ has a non-trivial overlap with the left side of some rectangle $R(v)$, then we color the edge $\{u, v\}$ in H red and orient it from u to v. Similarly, if the topside of $R(u)$ overlaps with the bottom side of $R(v)$ then $\{u, v\}$ is colored blue and oriented from u to v. Fig. 1(b) depicts the transversal structure obtained from the rectangular dual in Fig. 1(a). It is known that every transversal structure of H arises from a rectangular dual of H in this way.

Theorem 3 (Kant & He [19]) *Every transversal structure maps to a rectangular dual.*

If we identify *combinatorially equivalent* rectangular duals, i.e., those in which any two rectangles touch with the same sides in both duals, then Theorem 3 actually states that rectangular duals and transversal structures are in bijection. Transversal structures (and hence combinatorially equivalent rectangular duals) can be endowed with a distributive lattice structure [15]. For our purposes, we describe the *minimal transversal structure* of H; i.e., the minimum element in the distributive lattice of all transversal structures of H.

Lemma 1 (Fusy [15]) *Consider four vertices* v, w, x, y *in the minimal transversal structure* (E_r, E_b)*, such that* $v \to w \in E_b$*,* $x \to y \in E_b$*,* $v \to x \in E_r$*,* $w \to y \in E_r$*. Then we have neither* $x \to w \in E_b$ *nor* $v \to y \in E_r$*.*

Moreover, the minimal transversal structure can be computed in linear time.

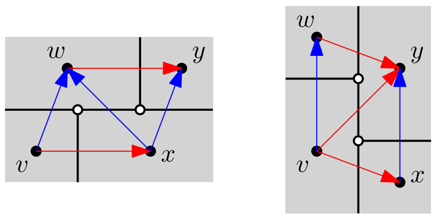

Figure 2: Two configurations that do not appear in the minimal transversal structure.

Fig. 2 shows the two configurations described in Lemma 1 that do not appear in the minimal transversal structure. The rectangular dual that corresponds to the minimal transversal structure is also called the *minimal rectangular dual*. Fig. 3(a) depicts the graph from Fig. 1 together with its minimal rectangular dual and the corresponding transversal structure. We remark that if, besides these two, a third certain configuration is forbidden in the transversal structure, then this already characterizes the minimal transversal structure [15].

Let us call a rectangular dual *non-degenerate* if the top sides of two rectangles lie on the same horizontal line only if there is a rectangle whose bottom side overlaps with both of them. It is not difficult to see that there always exists a non-degenerate minimal rectangular dual.

Given a rectangular dual and any inner vertex v we consider the rightmost rectangle overlapping the top side of $R(v)$. We denote the corresponding vertex of H by $v^\bullet$. In other words, $(v, v^\bullet)$ is the outgoing blue edge at v whose clockwise next edge is red (and outgoing). Similarly, $R(v_\bullet)$ is the bottommost rectangle overlapping the right side of $R(v)$, i.e., $(v, v_\bullet)$ is the outgoing red edge at v whose clockwise next edge is blue (and incoming). Moreover, $R({}_\bullet v)$ $(R({}^\bullet v))$ is the leftmost (topmost) rectangle overlapping the bottom side (left side) of $R(v)$, which means that $({}_\bullet v, v)$ $(({}^\bullet v, v))$ is the incoming blue (red) edge at v whose clockwise next edge is red (blue). Note that if the transversal structure is minimal then every inner edge of H can be written as $(v, v_\bullet)$, $(v, v^\bullet)$, $({}_\bullet v, v)$ or $({}^\bullet v, v)$ for some inner vertex v.

From H and its fixed transversal structure (E_r, E_b) we define a new graph H^*, called the *split graph*, and its transversal structure (E_r^*, E_b^*) as follows.

- The outer vertices of H and H^* are the same.
- For every inner vertex v of H there are two vertices v_1 and v_2 in H^*.
  - There is a red edge $v_1 \to v_2$ in E_r^*.
  - There is a red edge $v_2 \to w_1$ in E_r^* for every edge $v \to w \in E_r$.
  - There are blue edges $v_1 \to w_1$ and $v_1 \to w_2$ in E_b^* for every edge $v \to w \in E_b$.
  - There is a blue edge $v_2 \to (v^\bullet)_2$ in E_b^*.

See Fig. 3(b) for an example of a split graph and its rectangular dual. It is straight-forward to check that (E_r^*, E_b^*) is indeed a transversal structure, namely that for every $v \in V(H)$ incoming and outgoing red and blue edges appear around v_1 and v_2 in accordance with Definition 3. We refer to Fig. 3(b) for an illustration of this fact. Note that defining $R(v) := R(v_1) \cup R(v_2)$ for every vertex v we obtain the transversal structure we started with.

3.3 VPG-representation

We want to construct a B_2-VPG representation for every maximally planar graph G. To this end we split G into its 4-connected maximally planar subgraphs. The outer face Δ of such a subgraph H_Δ is either the outer face of G or an inner face of $H_{\Delta'}$, where Δ' is the father of Δ in the separation tree. We start by representing the outer face of G as depicted in Fig. 4. The highlighted area in the figure is called the frame for H_Δ. Formally, the *frame for H_Δ* is a rectangular area such that either: the paths corresponding to two vertices of Δ pass through it vertically and the path for the third vertex passes through it horizontally, or the paths corresponding to two vertices of Δ pass through it horizontally and third passes through it vertically. When defining the B_2-VPG

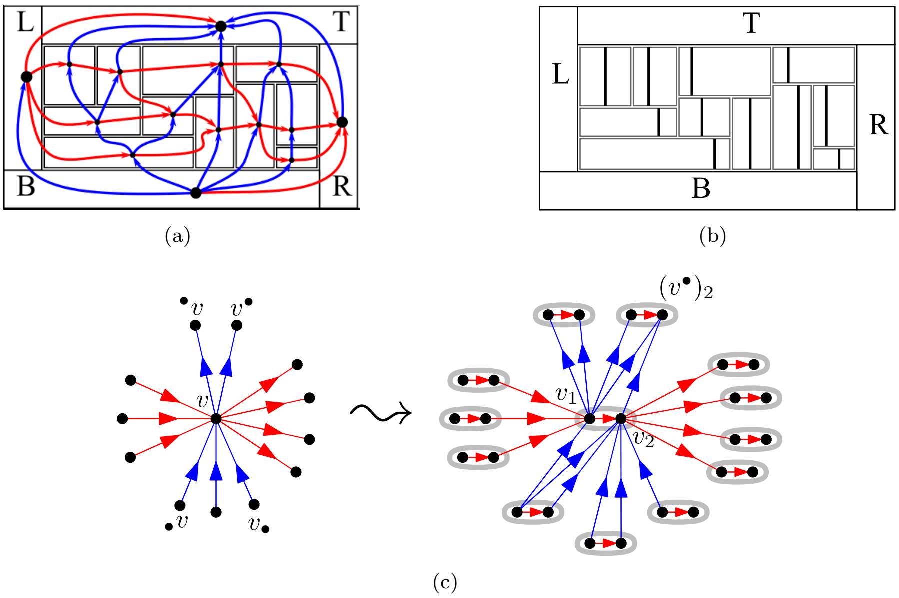

Figure 3: (a) The minimal rectangular dual of the graph in Fig. 1 with its transversal structure overlaid on it. (b) A rectangular dual of the split graph of (a). (c) Splitting a vertex v into v_1 and v_2 and the corresponding transversal structure.

representation of any H_Δ we assume that we have already constructed the paths for the vertices in Δ and that there is a frame for H_Δ.

We now describe how to obtain a B_2-VPG representation of a 4-connected maximally planar graph H_Δ given a frame F for it. Our construction is based on a non-degenerate minimal rectangular dual and its split graph. Let u and w be the two vertices of Δ whose paths do not intersect inside F and denote the third vertex in Δ by v. Then we consider the graph H obtained from H_Δ by removing the edge $\{u, w\}$. Notice that H is a 4-connected triangulation of a 4-gon and we assume without loss of generality that $u = L$, $v = T$, and $w = R$. Consider the minimal transversal structure, a corresponding non-degenerate minimal rectangular dual of H, and its split graph H^* together with the transversal structure (E_r^*, E_b^*). By rotating and stretching it appropriately we place the non-degenerate rectangular dual of H^* inside the frame F, such that the right side of L, the bottom side of T and the left side of R is contained in $\mathbf{u}$, $\mathbf{v}$ and $\mathbf{w}$, respectively.

We define the 2-bend path $\mathbf{B}$ for the vertex B to be a C-shape path that is contained in F and whose horizontal segments intersect $\mathbf{u}$ and $\mathbf{v}$, the upper one being contained in the top side of $R(B)$. See Fig. 4 for an illustration.

We define a 2-bend path $\mathbf{v}$ for every inner vertex v of H as follows. First, let $\mathbf{v}$ be the union of the top side and right side of $R(v_1)$ and the bottom side

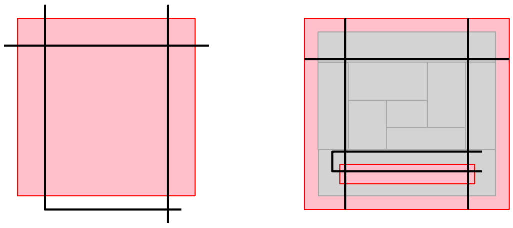

Figure 4: Left: The VPG representation of the outer face of G and its frame. Right: Placing a rectangular dual inside a frame and constructing the path $\mathbf{B}$.

of $R(v_2)$. Now consider the vertex ${}^{\bullet}v$. We extend the left horizontal end of $\mathbf{v}$ to the right side of $R(({}^{\bullet}v)_1)$. In case ${}^{\bullet}v = L$ we do not extend the left end of $\mathbf{v}$. Similarly we extend the right horizontal end of $\mathbf{v}$ horizontally to the right side of $R((v_{\bullet})_1)$, unless $v_{\bullet} = R$. See Fig. 5(a) for an illustration.

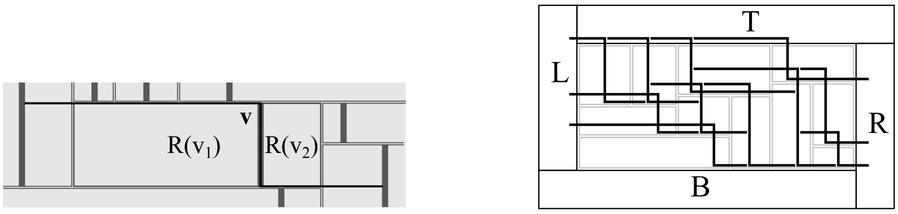

(a) (b)

Figure 5: (a) The path $\mathbf{v}$ based on the rectangles $R(v_1)$ and $R(v_2)$ in the rectangular dual of the split graph. Note: the wide edges indicate the border between split rectangles. (b) The Z-shapes arising from the split graph in Fig. 3(b).

Lemma 2 *The above construction gives a B_2-representation of H.*

Proof: Clearly every path defined above has at most two bends. So it remains to prove that the paths $\mathbf{u}$ and $\mathbf{v}$ intersect if and only if $\{u, v\}$ is an edge in G. Evidently, all outer edges $\{T, L\}$, $\{L, B\}$, $\{B, R\}$, and $\{T, R\}$ are realized, i.e., the corresponding paths intersect. Moreover, $\mathbf{T} \cap \mathbf{B} = \emptyset = \mathbf{L} \cap \mathbf{R}$ which means that no unwanted edge is created.

Now consider a blue edge $u \to v \in E_b$. By definition of the split graph and its transversal structure (E_r^*, E_b^*) we have an edge $u_1 \to v_2$ in E_b^*, i.e., the top side of $R(u_1)$ and the bottom side of $R(v_2)$ overlap. In particular $\mathbf{u} \cap \mathbf{v} \neq \emptyset$, since $\mathbf{u}$ and $\mathbf{v}$ contains the top side of $R(u_1)$ and the bottom side of $R(v_2)$, respectively.

Next consider a red edge of G. Since the underlying rectangular dual is minimal, it does not contain the configuration in the right of Fig. 2. Thus, every red edge can be written as $(v, v_{\bullet})$ or $({}^{\bullet}v, v)$ for some inner vertex v. By definition the right end of $\mathbf{v}$ lies on the right side of $R((v_\bullet)_1)$ (or $\mathbf{R}$ in case $v_\bullet = R$) and the left end of $\mathbf{v}$ lies on the right side of $R((^\bullet v)_1)$ (or $\mathbf{L}$ in case $^\bullet v = L$). Hence both edges are properly represented by intersecting paths.

Finally we need to argue that no two paths that correspond to non-adjacent vertices of G intersect. Therefore consider the parts of $\mathbf{v}$ that lie outside $R(v)$. The left extension of $\mathbf{v}$ passes through $R((^\bullet v)_2)$. This could be along the top side of $R((^\bullet v)_2)$, which is by definition of the split-graph strictly contained in the bottom side of some $R(w_2)$. Similarly, the right extension of $\mathbf{v}$ passes through $R((v_\bullet)_1)$ and this could be along the bottom side of this rectangle, which is strictly contained in some $R(w_1)$. In other words all left extensions are contained in $\bigcup_{v \in V} R(v_2)$ and all right extension are contained in $\bigcup_{v \in V} R(v_1)$. Thus a left extension may intersect a right extension only if these pass through $R(v_2)$ and $R(v_1)$ corresponding to the same vertex v, respectively. Since the underlying rectangular dual is non-degenerate the two extensions lie on distinct y-coordinates and hence are disjoint. □

Slightly changing the paths corresponding to outer vertices we can easily transform them into Z-shapes and make $\mathbf{L}$ and $\mathbf{R}$ intersect. Thus we obtain the following corollary.

Corollary 1 *Every 4-connected planar graph has a B_2-representation where every path has a Z-shape and no two paths cross.* □

We have shown so far how to define a B_2-VPG representation of H_Δ given a frame for H_Δ. It remains to identify a frame for each $\Delta' \sqsubset \Delta$ that is a son of Δ in the separation tree. We modify the representation for this purpose.

Consider a horizontal line ℓ that supports horizontal sides of some rectangles different from $R(T)$. We partition the paths that have a horizontal segment on ℓ into two sets: A contains all paths whose vertical segment lies above ℓ and B all paths whose vertical segment lies below ℓ. Next we extend the vertical segments of all paths in B by some small amount, keeping all lower horizontal segments unchanged. The extension is chosen small enough so that no unwanted intersections are created. See Fig. 6 for an illustration. Since the underlying rectangular dual is minimal, it does not contain the configuration in the left of Fig. 2. It follows that all vertical segments of paths in A lie to the left of the vertical segments of paths in B. Thus, if $\mathbf{v} \in A$ and $\mathbf{w} \in B$ were touching before, then they are crossing after this operation.

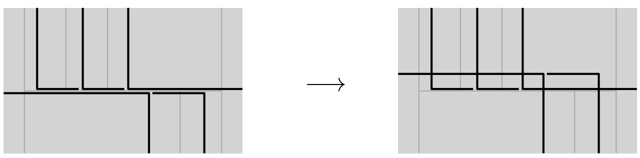

Figure 6: Extending the vertical segments of all paths in B.

Next we identify a frame for every inner face Δ' of H. In case Δ' is a non-empty triangle of G this will be the frame for $H_{\Delta'}$.

Lemma 3 *One can find in* $\mathbf{H}_\Delta$ *a frame for every inner face of* H_Δ*, such that each frame is contained in* F *and all frames are pairwise disjoint.*

Proof: First consider the triangle $\{L, B, R\}$, which is an inner face of H_Δ but not after the removal of the edge $\{L, R\}$. We define the frame for $\{L, B, R\}$ as illustrated in Fig. 4 to partly contain the lower horizontal segment of **B** and the vertical segments of **L** and **R**.

Now consider any inner face f of H_Δ different from $\{L, B, R\}$ and let u, v, w be the vertices of f appearing in this clockwise order. Then f is an inner face of H corresponding to the three mutually touching rectangles $R(u)$, $R(v)$ and $R(w)$ in the rectangular dual. Thus there is a point p_f where those three rectangles meet; two rectangles having a corner at p_f. Without loss of generality let $R(v)$ be the rectangle that does *not* have corner at p_f. We distinguish the four cases according to which side of $R(v)$ contains p_f. See Fig. 7 for an illustration.

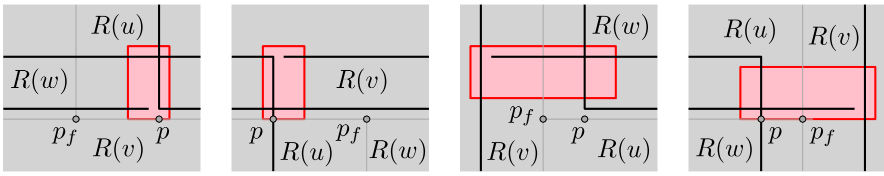

Figure 7: Identifying the frame for an inner face of H_Δ.

If the top side of $R(v)$ contains p_f, then consider the point p where $R(u_1)$, $R(u_2)$ and $R(v_1)$ meet. By definition p is the lower bend of **u** and the right horizontal end of **w**. Moreover, the upper horizontal segment of **v** lies immediately above p, crossing **u**. Now, the frame for f is defined around p as illustrated in Fig. 7 a).

If the bottom side of $R(v)$ contains p_f, then consider the point p where $R(u_1)$, $R(u_2)$ and $R(v_2)$ meet. Now right above p lies the upper bend of **u** and the left horizontal end of **w**, while **v** goes horizontally through p. The frame for f is then defined as illustrated in Fig. 7 b).

If the right side of $R(v)$ contains p_f, let p be the common point of $R(u_1)$, $R(w_1)$ and $R(w_2)$, i.e., p is the lower bend of **u**. The upper horizontal segment of **w** lies right above p and ends on the vertical segment of **v**. The frame for f is then defined as illustrated in Fig. 7 c).

Finally, if the left side of $R(v)$ contains p_f, let p be the common point of $R(u_2)$, $R(w_1)$ and $R(w_2)$, i.e., right above p lies the upper bend of **w**. The lower horizontal segment of **u** runs through p and ends on the vertical segment of **v**. The frame for f is then defined as illustrated in Fig. 7 d).

Clearly, each frame is contained in the frame for H_Δ. Moreover, each frame contains one bend or lies very close to one. Given the bend one can find the corresponding p_f to the left if it is a lower bend, and to the bottom-right if it is an upper bend. It follows that frames and bends are in bijection and hence that all frames are pairwise disjoint. □

We end this section with its main theorem. It is not difficult to see that this theorem follows from Theorem 1, and Lemmas 2 and 3.

Theorem 4 *Every planar graph is B_2-VPG. Moreover, a B_2-VPG representation can be found in $\mathcal{O}(n^{3/2})$, where n denotes the number of vertices in the graph.*

Proof: Given a maximally planar graph G with a fixed embedding, we find the separation tree of G in $\mathcal{O}(n^{3/2})$ and all 4-connected maximally planar subgraphs H_Δ of G (Theorem 1). We define a B_2-VPG representation of the outer triangle Δ of G as explained in Section 3.3 and identify the frame for H_Δ (Fig. 4). Then we traverse the separation tree starting with the root and consider for each non-empty triangle Δ the frame F for the corresponding graph H_Δ. If u and w are the vertices of Δ whose paths $\mathbf{u}$ and $\mathbf{w}$ do not intersect within F, we consider the graph $H = H_\Delta \setminus \{u, w\}$. We find the minimal transversal structure of H in $\mathcal{O}(|V(H)|)$ (Lemma 1) and build the split graph H^* as described in Section 3.2. We then construct a B_2-VPG representation of H within the frame F as described in Section 3.3 and identify frames for each non-empty triangle Δ' that is an inner face of H_Δ. The construction of the split graph and the B_2-VPG representation can be easily done in $\mathcal{O}(|V(H)|)$. Hence the overall running time is dominated by the time needed to find the separation tree, i.e., a B_2-VPG representation can be constructed in $\mathcal{O}(|V(G)|^{3/2})$. □

4 Triangle-Free Planar Graphs are B_1-VPG

In this section we prove that every triangle-free planar graph is B_1-VPG with a very particular B_1-VPG representation. Namely, every vertex is represented by either a 0-bend path or a 1-bend path whose vertical segment is attached to the left end of its horizontal segment. This means that we use only two out of the four possible shapes of a grid path with exactly one bend. Moreover, whenever two paths intersect, it is at an endpoint of exactly one of these paths; i.e., no two paths cross. We call a 1-bend path an L if the left endpoint of the horizontal segment is the lower endpoint of its vertical segment, and a Γ if the left endpoint of the horizontal segment is the upper endpoint of its vertical segment. A VPG representation in which each path that has a bend is an L or a Γ, and in which no two paths cross, is called a *contact-L-Γ representation.*

We say that two contact-L-Γ representations of the same graph G are *equivalent* if the underlying combinatorics is the same. That means that paths corresponding to the same vertex have the same type (either L, Γ, horizontal or vertical segment), the inherited embedding of G is the same, and that the fashion in which two paths touch is the same, e.g., the right endpoint of $\mathbf{u}$ is contained in the vertical segment of $\mathbf{v}$ in both representations. However, it is convenient in our proofs to deal with actual contact-L-Γ representations instead of equivalence classes of contact-L-Γ representations. Therefore we need the following lemma.

Lemma 4 *Let G be a plane graph and v be a vertex of G. Let* $\mathbf{u}$ *and* $\mathbf{w}$ *be two paths in* $\mathbf{G}$ *that touch* $\mathbf{v}$ *at the same segment but from different sides. Then there exists a contact-L-Γ representation of G that is equivalent to* $\mathbf{G}$ *in which the touching points of* $\mathbf{u}$ *and* $\mathbf{w}$ *with* $\mathbf{v}$ *come in the reversed order along* $\mathbf{v}$*.*

Proof: We obtain the required representation from $\mathbf{G}$ with a simple operation, called *slicing*. Assume without loss of generality that the segment s_v of $\mathbf{v}$ that is touched by $\mathbf{u}$ and $\mathbf{w}$ is vertical, i.e., the horizontal segments s_u of $\mathbf{u}$ and s_w of $\mathbf{w}$ touch s_v. Assume further without loss of generality that $s_u \cap s_v$ lies above $s_w \cap s_v$ and that s_u lies to the left and s_w to the right of s_v, respectively. Consider any 2-bend grid path P containing s_u and s_w and extend its left and right endpoints to the left and to the right to infinity, respectively. Then P divides the plane into two unbounded regions. We denote the lower region by A and consider s_u to be contained in A, and the upper region by B and consider s_w to be contained in B. Now we increase the y-coordinates of every point in B by some amount large enough that $s_w \cap s_v$ lies above $s_u \cap s_v$. All vertical segments that cross P, including s_v and maybe the vertical segments of $\mathbf{u}$ and $\mathbf{w}$ are extended so that the corresponding paths are connected again.

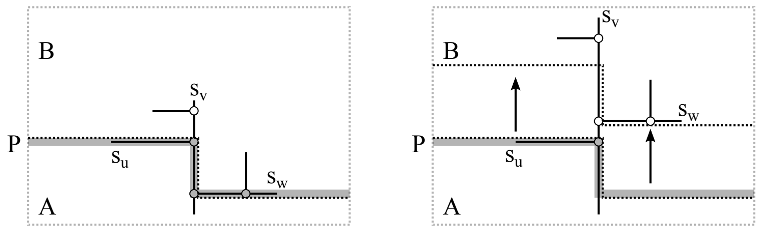

Figure 8: The slicing operation.

The slicing operation is illustrated in Fig. 8. Figuratively speaking, we cut the plane along P and pull the two pieces apart until s_u and s_w change the order along s_v, while paths that cross P are stretched instead of cut. □

The main result of this section is the following.

Theorem 5 *Every triangle-free planar graph has a contact-L-Γ representation.*

Note that if some graph G admits a contact-L-Γ representation then so does every subgraph H of G. Indeed every edge (u, v) in $E(G) \setminus E(H)$ corresponds to a contact point of $\mathbf{u}$ and $\mathbf{v}$ in the representation $\mathbf{G}$. Moreover, this contact point is an endpoint of one of the two paths. If we shorten this path a little bit, and do this for every edge that is in G but not in H, then we obtain a contact-L-Γ representation of H. Thus we assume for the remainder of the section without loss of generality that G is a maximally triangle-free planar graph, i.e., G is 2-connected and every face of G is a quadrangle or a pentagon. Moreover, we can assume by adding one vertex (if necessary) that the outer face of G is a quadrangle.

Consider a contact-L-Γ representation $\mathbf{C}$ of a cycle C on four vertices v_1, v_2, v_3, v_4 and assume without loss of generality that any two paths in $\mathbf{C}$ touch at most once. Then $\mathbf{v_1} \cup \mathbf{v_2} \cup \mathbf{v_3} \cup \mathbf{v_4}$ inscribes a simple rectilinear polygon P. We call the parts of $\mathbf{C}$ that do not lie in the interior of P the *outside of* $\mathbf{C}$. See Fig. 9 for an example.

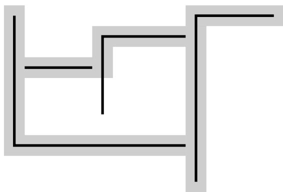

Figure 9: A contact-L-Γ representation of a 4-cycle. Its outside is highlighted.

We prove the following stronger version of Theorem 5.

Theorem 6 *Let G be a maximally triangle-free planar graph with a fixed plane embedding and a quadrangular outer face C_{out}. Let $\mathbf{C_{out}}$ be any contact-L-Γ representation of C_{out}. Then there is a contact-L-Γ representation of G with the same underlying embedding in which the outside of the induced representation of C_{out} is equivalent to that in $\mathbf{C_{out}}$.*

Proof: We do induction on the number of vertices in G, distinguishing the following three cases.

Case 1: G has a separating 4-cycle C. Let V_C be the set of vertices interior to C and G_1 be the graph $G - V_C$. Note that G_1 is maximally triangle-free and with outer face C_{out}. Hence by induction we find a contact-L-Γ representation $\mathbf{G_1}$ of G_1 such that C_{out} is represented with an equivalent outside as in $\mathbf{C_{out}}$. Since the representation $\mathbf{G_1}$ respects the embedding of G_1, the interior of $\mathbf{C}$ is empty. We again apply induction to $G_2 = G[C \cup V_C]$ with respect to the representation $\mathbf{C}$ induced by $\mathbf{G_1}$ and obtain a contact-L-Γ representation $\mathbf{G_2}$. Since the outside of the representation of C in $\mathbf{G_2}$ is equivalent to that in $\mathbf{G_1}$ we can put together $\mathbf{G_1}$ and $\mathbf{G_2}$ and obtain a contact-L-Γ representation $\mathbf{G}$ of G that satisfies our requirements.

Case 2: G has a facial 4-cycle $C = \{v_1, v_2, v_3, v_4\}$. Let v_1 and v_3 be two opposite vertices on C that have distance (counted by the number of edges) at least 4 in $G - \{v_2, v_4\}$. Since G is triangle-free and planar, such vertices exist and we can moreover assume without loss of generality that v_1 is not an outer vertex. Let $\tilde{G}$ be the graph resulting from G by merging v_1 and v_3, and denoting the new vertex by $\tilde{v}$. Note that $\tilde{G}$ is a maximally triangle-free planar graph that inherits a plane embedding from G. Moreover $\tilde{G}$ has outer cycle C_{out} where possibly v_3 is replaced by $\tilde{v}$. By induction we find a contact-L-Γ representation $\tilde{\mathbf{G}}$ of $\tilde{G}$. Next we split the path $\tilde{\mathbf{v}}$ in $\tilde{\mathbf{G}}$ into two, one for v_1 and one for v_3, which will result in a contact-L-Γ representation $\mathbf{G}$ of G. See Fig. 10 for an example.

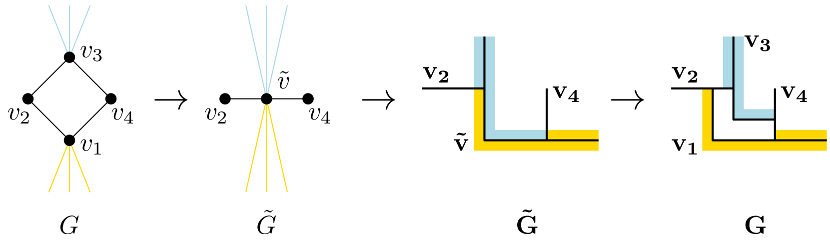

Figure 10: How to split a face in *Case 2*.

Consider the circular ordering of contacts when tracing around $\tilde{\mathbf{v}}$ in $\tilde{\mathbf{G}}$. The paths $\mathbf{v_2}$ and $\mathbf{v_4}$ split the circular ordering into two consecutive blocks, that is, subsets of contacts one corresponding to neighbors of v_1 and one corresponding to neighbors of v_3 in G. (There are no common neighbors of v_1 and v_3 apart from v_2 and v_4, because v_1 and v_3 are at distance at least 4 in $G - \{v_2, v_4\}$.) Now define $\mathbf{v_3}$ to be the sub-path of $\tilde{v}$ defined by the block of neighbors of v_3. Moreover define $\mathbf{v_1}$ in the same way w.r.t. the neighbors of v_1, except that $\mathbf{v_1}$ is translated by some small amount "towards its block". Finally, every path $\mathbf{u}$ corresponding to a neighbor u of v_1 different from v_2 and v_4 is shortened or extended so that it touches $\mathbf{v_1}$. The procedure for *Case 2* is illustrated in Fig. 10.

It is important to note that, even if an outer edge is involved in the above construction, the outsides of $\mathbf{C_{out}}$ in $\mathbf{G}$ is equivalent to that in $\tilde{\mathbf{G}}$.

Case 3: Neither Case 1 nor Case 2 applies and there is an edge (u, v) in G with interior vertices u and v. We contract the edge (u, v) and denote by $\tilde{v}$ the new vertex in the resulting graph $\tilde{G}$. Since neither *Case 1* nor *Case 2* applies, u and v are at distance 4 in $G - (u, v)$ and thus $\tilde{G}$ is maximally triangle-free. Moreover $\tilde{G}$ has outer cycle C_{out} and inherits its plane embedding from G. By induction we find a contact-L-Γ representation $\tilde{\mathbf{G}}$, in which we want to split $\tilde{\mathbf{v}}$ into two paths $\mathbf{v}$ and $\mathbf{u}$, such that the result is a contact-L-Γ representation $\mathbf{G}$ of G.

As in the previous case we trace the contour of $\tilde{\mathbf{v}}$ and see two disjoint blocks, each consisting of those contacts that correspond to neighbors of u and v in G, respectively. We denote the block corresponding to u and v by B_u and B_v, respectively. Without loss of generality assume that $B_u \cup B_v$ is the entire contour of $\tilde{\mathbf{v}}$. We distinguish the following four sub-cases. By symmetry we assume that $\tilde{\mathbf{v}}$ is not a Γ-shape and denote its vertical segment (if existent) by s.

In *Case 3a* either s is completely covered by one block, say B_u, or $\tilde{\mathbf{v}}$ is only a horizontal segment and B_u is the block that contains the left endpoint of it. We define $\mathbf{u}$ and $\mathbf{v}$ to be the sub-paths of $\tilde{\mathbf{v}}$ that are covered by B_u and B_v, respectively. We shift $\mathbf{v}$ a little bit up or down and attach a short vertical segment to its left endpoint so as to touch $\mathbf{u}$. The construction is illustrated in Fig. 11.

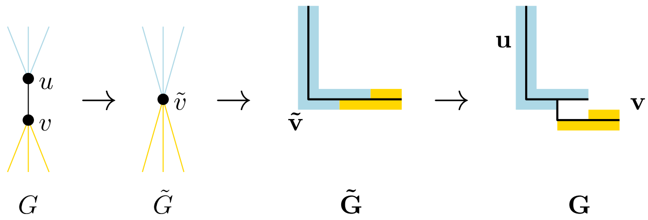

Figure 11: How to split an edge in *Case 3a*.

In *Case 3b* the left side of s is completely covered by one block, say B_u. We define $\mathbf{u}$ to be the sub-path of $\tilde{\mathbf{v}}$ that is covered by B_u. If B_v is contained in s, we define $\mathbf{v}$ to be a very short horizontal segment touching the right side of s immediately below the B_v. Otherwise we define $\mathbf{v}$ to be the sub-path of the horizontal segment of $\tilde{\mathbf{v}}$ that is covered by B_v and shift $\mathbf{v}$ a little bit up. Note that each path that touches the right side of s is only a horizontal segment. We shorten the left endpoint of each such path that corresponds to a neighbor of v a little bit and attach a vertical segment to it that touches $\mathbf{v}$ from above. This can be done so that no two such paths intersect. Moreover, every vertical segment touching $\tilde{\mathbf{v}}$ and corresponding to B_v is shortened or extended a bit so as to touch $\mathbf{v}$. See the left of Fig. 12 for an illustration.

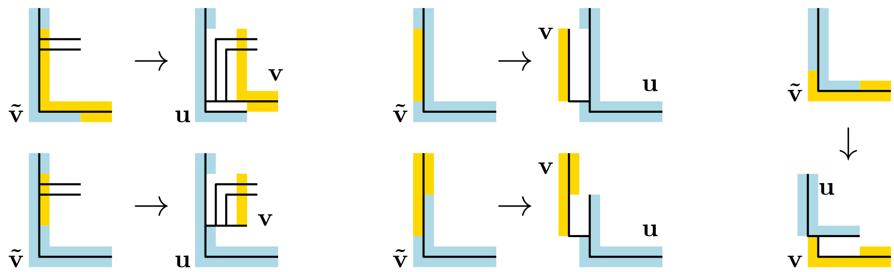

Figure 12: How to split an edge in *Case 3b*, *Case 3c*, and *Case 3d*.

In *Case 3c* either the horizontal segment of $\tilde{\mathbf{v}}$ is completely covered by one block, say again B_u, or $\tilde{\mathbf{v}}$ is only a vertical segment and B_u is the block that contains the lower endpoint of it. Note that since *Case 3b* does not apply, B_v partially covers the left side of s. By Lemma 4 we can assume that no point of s is covered on the left by B_u and on the right by B_v. We define $\mathbf{u}$ and $\mathbf{v}$ to be the sub-paths of $\tilde{\mathbf{v}}$ that are covered by B_u and B_v, respectively, and shift $\mathbf{v}$ a little bit to the left. Again we shorten or extend each path that corresponds to a neighbor of v so that it touches $\mathbf{v}$. See the middle of Fig. 12 for an illustration.

In the remaining case, *Case 3d*, both blocks B_u and B_v appear on both sides of the vertical and horizontal segment of $\tilde{\mathbf{v}}$. Let B_u be the block that contains the upper end of $\tilde{\mathbf{v}}$. Consider paths that touch the horizontal segment of $\tilde{\mathbf{v}}$ on the upper side and within the block B_u. By Lemma 4 may can assume that the horizontal segment of each such path lies above the block B_v. We define $\mathbf{u}$ and $\mathbf{v}$ to be the sub-paths of $\tilde{\mathbf{v}}$ that are covered by B_u and B_v, respectively. We shift the horizontal segment of $\mathbf{u}$ up to the upper endpoint of $\mathbf{v}$ and move $\mathbf{u}$ a little bit to the left so that $\mathbf{v}$ touches $\mathbf{u}$ from below. Moreover, we shorten or extend every path corresponding to a neighbor of u so that it touches $\mathbf{u}$. This completes Case 3.

Finally, if neither of *Case 1*, *Case 2* and *Case 3* applies, then G consists only of the outer cycle C_{out}, for which a Contact-L-Γ representation $\mathbf{C_{out}}$ is given by assumption. This concludes the proof. □

Theorem 5 can be easily transferred into a linear-time algorithm to find a contact-L-Γ representation of a triangle-free planar graph. Note that such an algorithm should first construct the combinatorics of the representation, since slicing operation would have to be done in $\mathcal{O}(1)$. The computation of the actual coordinates of each path can be easily carried out afterwards in linear time. Moreover the constructed representation can be placed into the $n \times n$ grid, since every path requires only one horizontal and one vertical grid line. Here n denotes the number of vertices in G.

5 Future Work and Open Problems

We have disproved the conjecture of Asinowski et al. [2] that B_3-VPG is the simplest B_k-VPG graph class containing planar graphs. Specifically, we have demonstrated that every planar graph is B_2-VPG and that 4-connected planar graphs are the intersection graphs of Z-shapes (i.e., a special subclass of B_2-VPG). We have also shown that these representations can be produced from a planar graph in $\mathcal{O}(n^{3/2})$ time. We have additionally shown that every triangle-free planar graph is a contact graph of: L-shapes, Γ-shapes, vertical segments, and horizontal segments (i.e., it is a specialized contact B_1-VPG graph). Furthermore, we demonstrated how to construct such a contact representation in linear time. As an further consequence, we obtain a new proof that planar bipartite graphs are 2-DIR.

Interestingly, there is no known planar graph which does not have an intersection representation of L-shapes; i.e., even this very restricted form of B_1-VPG is still a good candidate to contain all planar graphs. Further to this, a colleague of ours has observed (via computer search) that all planar graphs on at most ten vertices are intersection graphs of L-shapes [16]. Similarly, all small triangle-free planar graphs seem to be contact graphs of L-shapes. These observations lead to the following two conjectures.

Conjecture 1 *Every planar graph is the intersection graph of L-shapes.*

Conjecture 2 *Every triangle-free planar graph is the contact graph of L-shapes.*

References

[1] M. J. Alam, T. Biedl, S. Felsner, M. Kaufmann, and S. G. Kobourov. Proportional contact representations of planar graphs. In *19th Symposium on Graph Drawing*, GD 2011, pages 26–38, 2011. doi:10.1007/978-3-642-25878-7_4.

[2] A. Asinowski, E. Cohen, M. C. Golumbic, V. Limouzy, M. Lipshteyn, and M. Stern. String graphs of k-bend paths on a grid. *Electronic Notes in Discrete Mathematics*, 37(0):141 – 146, 2011. doi:10.1016/j.endm.2011.05.025.

[3] A. Asinowski, E. Cohen, M. C. Golumbic, V. Limouzy, M. Lipshteyn, and M. Stern. Vertex intersection graphs of paths on a grid. *J. Graph Algorithms Appl.*, 16(2):129–150, 2012. doi:10.7155/jgaa.00253.

[4] I. Ben-Arroyo Hartman, I. Newman, and R. Ziv. On grid intersection graphs. *Discrete Math.*, 87(1):41 – 52, 1991. doi:10.1016/0012-365X(91)90069-E.

[5] J. Chalopin and D. Gonçalves. Every planar graph is the intersection graph of segments in the plane: extended abstract. In *41st annual ACM Symposium on Theory of Computing*, STOC '09, pages 631–638, 2009. doi:10.1145/1536414.1536500.

[6] J. Chalopin, D. Gonçalves, and P. Ochem. Planar graphs are in 1-STRING. In *18th annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, pages 609–617, 2007. URL: http://dl.acm.org/citation.cfm?id=1283383.1283449.

[7] S. Chaplick and T. Ueckerdt. Planar graphs as VPG-graphs. In W. Didimo and M. Patrignani, editors, *Graph Drawing*, volume 7704 of *Lecture Notes in Computer Science*, pages 174–186. Springer, 2012. doi:10.1007/978-3-642-36763-2_16.

[8] N. de Castro, F. J. Cobos, J. C. Dana, A. Márquez, and M. Noy. Triangle-free planar graphs and segment intersection graphs. *J. Graph Algorithms Appl.*, 6(1):7–26, 2002. doi:10.7155/jgaa.00043.

[9] H. de Fraysseix and P. Ossona de Mendez. Representations by contact and intersection of segments. *Algorithmica*, 47(4):453–463, 2007. doi:10.1007/s00453-006-0157-x.

[10] H. de Fraysseix, P. Ossona de Mendez, and J. Pach. Representation of planar graphs by segments. *Intuitive Geometry*, 63:109–117, 1991.

[11] H. de Fraysseix, P. Ossona de Mendez, and P. Rosenstiehl. On triangle contact graphs. *Combinatorics, Probability & Computing*, 3:233–246, 1994. doi:10.1017/S0963548300001139.

[12] G. Ehrlich, S. Even, and R. Tarjan. Intersection graphs of curves in the plane. *J. Comb. Theory, Ser. B*, 21(1):8 – 20, 1976. doi:10.1016/0095-8956(76)90022-8.

[13] D. Eppstein. Regular labelings and geometric structures. In *22nd Canadian Conference on Computational Geometry*, CCCG 2010, pages 125–130, 2010.

[14] E. Fusy. *Combinatoire des cartes planaires et applications algorithmiques.* PhD Thesis, 2007.

[15] E. Fusy. Transversal structures on triangulations: A combinatorial study and straight-line drawings. *Discrete Math.*, 309(7):1870–1894, 2009. doi:10.1016/j.disc.2007.12.093.

[16] T. Gavenčiak. Private communication: Small planar graphs are L-graphs., 2012.

[17] X. He. On finding the rectangular duals of planar triangular graphs. *SIAM J. Comput*, 22:1218–1226, 1993. doi:10.1137/0222072.

[18] D. Heldt, K. Knauer, and T. Ueckerdt. On the bend-number of planar and outerplanar graphs. In *10th Latin American Symposium on Theoretical Informatics*, LATIN 2012, pages 458–469, 2012. doi:10.1007/978-3-642-29344-3_39.

[19] G. Kant and X. He. Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. *Theor. Comput. Sci.*, 172(1-2):175–193, 1997. doi:10.1016/S0304-3975(95)00257-X.

[20] S. Kobourov, T. Ueckerdt, and K. Verbeek. Combinatorial and geometric properties of Laman graphs. In *24th annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, 2013.

[21] P. Koebe. Kontaktprobleme der konformen Abbildung. *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig. Math.-Phys. Klasse*, 88:141–164, 1936.

[22] K. Koźmiński and E. Kinnen. Rectangular dual of planar graphs. *Networks*, 15(2):145–157, 1985.

[23] J. Kratochvíl and J. Matoušek. Intersection graphs of segments. *J. Comb. Theory, Ser. B*, 62(2):289–315, 1994. doi:10.1006/jctb.1994.1071.

[24] Y.-T. Lai and S. M. Leinwand. An algorithm for building rectangular floorplans. In *21st Design Automation Conference*, DAC '84, pages 663–664, 1984.

[25] P. Rosenstiehl and R. Tarjan. Rectilinear planar layouts and bipolar orientations of planar graphs. *Discrete & Computational Geometry*, 1:343–353, 1986. doi:10.1007/BF02187706.

[26] E. R. Scheinerman. *Intersection Classes and Multiple Intersection Parameters of Graphs.* PhD thesis, Princeton University, 1984.

[27] F. Sinden. Topology of thin film circuits. *Bell System Tech. J.*, 45:1639–1662, 1966.

[28] Y. Sun and M. Sarrafzadeh. Floorplanning by graph dualization: L-shaped modules. *Algorithmica*, 10:429–456, 1993. `doi:10.1007/BF01891831`.

[29] C. Thomassen. Interval representations of planar graphs. *J. Comb. Theory, Ser. B*, 40(1):9–20, 1986. `doi:10.1016/0095-8956(86)90061-4`.

[30] P. Ungar. On diagrams representing graphs. *J. London Math. Soc.*, 28:336–342, 1953.

[31] D. West. Open problems. *SIAM Activity Group Newsletter in Discrete Mathematics*, 2:1012, 1991.

Graphs admitting d-realizers: spanning-tree-decompositions and box-representations

William Evans[1,a] Stefan Felsner[2,b] Stephen G. Kobourov[3,c] Torsten Ueckerdt[4]

Abstract

A d-realizer is a collection $R = \{\pi_1, \ldots, \pi_d\}$ of d permutations of a set V representing an antichain in $\mathbb{R}^d$. We use R to define a graph G_R on the suspended set $V^+ = V \cup \{s_1, \ldots, s_d\}$. It turns out that G_R has $dn + \binom{d}{2}$ edges ($n = |V|$), among them the edges of the outer clique on $\{s_1, \ldots, s_d\}$. The inner edges of G_R can be partitioned into d trees such that T_i spans $V + s_i$. In the case $d = 3$ the graph G_R is a planar triangulation and T_1, T_2, T_3 is a Schnyder wood on G_R. The following two results show that d-realizers resemble Schnyder woods in several aspects:

- Complete point-face contact systems of homothetic simplices in $\mathbb{R}^{d-1}$ induce a d-realizer.
- Any spanning subgraph of a graph G with a d-realizer has a d-dimensional proper touching box representation.

We expect that d-realizers will prove to be valuable generalization of Schnyder woods to higher dimensions.

1 Introduction

We consider $\mathbb{R}^d$ equipped with the *dominance order*, i.e., for $x, y \in \mathbb{R}^d$ we have $x \leq_{\mathsf{dom}} y$ if and only if $x_i \leq y_i$ for $i = 1, \ldots, d$. A set $P \subset \mathbb{R}^d$ is in *general position* if no two points of P share a coordinate. If no two points of a set P are in the dominance relation $\leq_{\mathsf{dom}}$, then we call P an *antichain*. If P is in general position, then the projection to the ith coordinate yields a permutation π_i of P. In compliance with the previous definition, we call a family of permutations $\pi_1, \ldots, \pi_d$ of V an *antichain* if for all $x, y \in V$ there are indices i and j such that x precedes y in π_i and y precedes x in π_j. We use the notation $x \prec_i y$ to denote that x precedes y in π_i.

An antichain V in $\mathbb{R}^d$ is *suspended* if V contains a suspension vertex for each i, i.e., a vertex $s_i = (0, \ldots, 0, M_i, 0, \ldots, 0)$ and $0 \leq v_i < M_i$ for all $v \in V \setminus s_i$.

[1]Dept. of Computer Science, Univ. of British Columbia, Vancouver, B.C. V6T 1Z4 Canada
[2]Institut für Mathematik, Technische Universität Berlin, D-10623 Berlin, Germany
[3]Dept. of Computer Science, University of Arizona, Tucson, AZ, USA
[4]Dept. of Mathematics, Karlsruhe Institute of Technology, D-76128 Karlsruhe, Germany

[a]Research supported in part by NSERC of Canada
[b]Partially supported by DFG grant FE-340/7-2 and ESF EuroGIGA project GraDR.
[c]Research funded in part by NSF grants CCF-0545743 and CCF-1115971

Similarly s_i is an i-suspension for $\pi_1, \ldots, \pi_d$ if s_i is the last element of π_i and among the first $d - 1$ elements in π_j for $j \neq i$. The family $\pi_1, \ldots, \pi_d$ is suspended if it has an i-suspension for each $i \in [d]$.

Definition 1 *A* d-realizer *is a suspended antichain $\pi_1, \ldots, \pi_d$ of permutations of V^+ where $V^+ = V \cup S$ and $S = \{s_1, \ldots, s_d\}$ is the set of suspensions.*

Definition 2 *The* graph of a d-realizer $(\pi_1, \ldots, \pi_d)$ *is the graph $G_R = (V^+, E^+)$ with $E^+ = E_R \cup E_S$ where E_S is the set of edges of a clique on S and pairs x, y are edges in E_R if they satisfy two properties:*

(x, y) **is a candidate pair**: *for all $z \neq x, y$ there is an i with $x \prec_i z$ and $y \prec_i z$.*

(x, y) **has the 1-of-d-property**: *there is a unique $i \in [d]$ with $x \prec_i y$, i.e., $y \prec_j x$ for all $j \neq i$.*

The definition of Schnyder woods was originally motivated by the study of the order dimension of incidence posets of graphs. In this line of research the following definition was proposed in [8]:

The dimension *of $G = (V, E)$ is at most k if there are permutations $\pi_1, \ldots, \pi_k$ of V such that each edge $(x, y) \in E$ is a candidate pair.*

If G is two-connected, then it follows that $\pi_1, \ldots, \pi_k$ is a antichain. The following are known:

- $\dim(G) \leq 3$ iff G is planar (Schnyder [12]).
- $\dim(G) \leq 4 \implies G$ has at most $3/8n^2$ edges.
- Exact values of $\dim(K_n)$ are known for $n < 10^{40}$.

The 1-of-d-property naturally leads to a coloring and an orientation of the edges of G_R: The orientation is $x \to y$ if x precedes y only in a single π_i. The color of $x \to y$ is the index i with $x \prec_i y$. Let T_i be the set of edges of color i.

Note that in the case $d = 3$ the 1-of-3-property is fulfilled by all candidate edges; this is where Schnyder's coloring and orientation of edges comes from. Schnyder [12] found that for all i the following two properties hold:

(a) T_i is an in-arborescence with root s_i.

(b) $T_{i-1} + T_{i+1} + T_i^{-1}$ is acyclic.

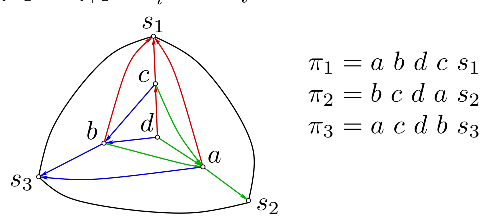

Fig. 1: *An example of a 3-realizer and its graph.*

In the next section we show that this also holds in the case of a d-realizer. In Section 3 we continue to show how d-realizer can be used to construct proper touching box representations; the $d = 3$ case of this result was obtained in [1]. In Section 4 we connect d-realizers to orthogonal surfaces and show how they arise from touching simplices. We conclude with examples and some open problems.

2 Spanning-tree-decompositions

Proposition 1 *Let a graph be defined by a d-realizer $(\pi_1, \ldots, \pi_d)$. If T_i is the set of edges of color i, then T_i is an in-arborescence with root s_i.*

Proof. We first show that each $v \in V$ has a unique out-edge in T_i.

Let $H_i(x)$ be the set of all y with $x \prec_i y$ and $y \prec_j x$ for all $j \neq i$, i.e., the set of all y such that the pair (x, y) has the 1-of-d-property. Since the pair (x, s_i) has the 1-of-d-property $H_i(x) \neq \emptyset$ for all $v \in V$. Let $p_i(x)$ be the first element of $H_i(x)$ with respect to π_i, i.e., $p_i(x)$ is the least element of π_i such that $(x, p_i(x))$ has the 1-of-d-property.

Claim 1. $(x, p_i(x))$ is a candidate.

Consider $z \neq x, p_i(x)$. Since a d-realizer is an antichain there is some j with $x \prec_j z$. If $j \neq i$, then $p_i(x) \prec_j x$ and by transitivity $p_i(x) \prec_j z$. If the only choice for j is i, then $z \in H_i(x)$ and $p_i(x) \prec_j z$ follows from the choice of $p_i(x)$. △

From Claim 1 it follows that $(x, p_i(x)) \in T_i$.

Claim 2. If (x, y) is a candidate with $y \in H_i(x)$, then $y = p_i(x)$.

Indeed if $y \neq p_i(x)$ then there is no π_j where x and y precede $p_i(x)$. In π_i we have $x \prec_i p_i(x) \prec_i y$ and if $j \neq i$, then $p_i(x) \prec_j x$. △

Hence $(x, p_i(x))$ is the only out-edge of x in T_i. Therefore the number of edges of T_i is $|V|$. Since T_i is spanning $V + s_i$ it only remains to show that T_i is connected. For $x \in V$ define $x_0 = x$ and for $k \geq 0$ let $x_{k+1} = p_i(x_k)$. This defines a path that moves to the right on π_i; hence it must reach s_i. □

Corollary 1 *A graph G_R defined by a d-realizer on a vertex set V^+ with $|V^+| = n + d$ has $dn + \binom{d}{2}$ edges.*

Proposition 2 *If G_R is defined by a d-realizer , then $T_i^{-1} + \sum_{j \neq i} T_j$ is acyclic.*

Proof. From the 1-of-d-property it follows that directed edges from T_j with $j \neq i$ point to the left in the order of vertices given by π_i. The same is true if we revert the direction of the edges of T_i, i.e, for the directed edges of T_i^{-1}. □

3 Box-representations

We consider axis aligned boxes in d-space. Such a *box* is a set $B(a, b) = \{x \in \mathbb{R}^d : a \leq_{\mathsf{dom}} x \leq_{\mathsf{dom}} b\}$ or equivalently $B(a, b) = \prod_{i=1}^{d} [a_i, b_i]$. The interior of $B(a, b)$ is $\{x \in \mathbb{R}^d : a <_{\mathsf{dom}} x <_{\mathsf{dom}} b\} = \prod_{i=1}^{d} (a_i, b_i)$. Two boxes B and B' are *properly touching* iff they have a unique separating hyperplane $H = \{x \in \mathbb{R}^d : n_H^T \cdot x = b_H\}$, i.e., $n_H^T \cdot x \leq b_H$ for all $x \in B$ and $n_H^T \cdot x \geq b_H$ for all $x \in B'$. In other words, B and B' are properly touching if their interiors are disjoint and their intersection is $(d-1)$-dimensional.

Definition 3 *A* proper touching box representation *of a graph $G = (V, E)$ in d dimensions consists of a map $v \to B_v$ from the vertices to d-dimensional boxes with pairwise disjoint interiors, such that boxes B_u and B_v are properly touching iff $(u, v) \in E$.*

Box representations of graphs have been studied in 2D with different names, e.g as *rectangle contact graphs.* Surveys of the state of the art can be found in [3] and [5].

For 3D, Thomassen [13] shows that any planar graph has a proper touching box representation. Felsner and Francis [6] prove that any planar graph has a touching cube representation, if the graph is a subgraph of a 4-connected triangulation the representation is proper. New proofs of Thomassen's result and additional results on cube representations can be found in [1].

Theorem 1 *Any spanning subgraph H of a graph G with a d-realizer has a d-dimensional proper touching box representation.*

Proof. Let $(\pi_1, \ldots, \pi_d)$ be the d-realizer for G. We assume that the order of the first $d - 1$ elements in π_i (these are suspensions) is $(s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_d)$. This has the advantage that for $i < j$ the pair (s_j, s_i) has the 1-of-d-property. The pair is a candidate so we can treat it as a regular edge in T_i.

With $\mathsf{rank}_i(x)$ we denote the position of x in π_i, i.e., if we think of π_i as a bijective map $\pi_i : [n + d] \to V^+$, then $\mathsf{rank}_i(x) = \pi_i^{-1}(x)$. For each x and i we define $p_i(x)$ as in Prop 1. For a suspension s_i and all j with $i \leq j$ we assume the value $n + d + 1$ for the strictly speaking undefined expression $\mathsf{rank}_j(p_j(s_i))$.

We now show how to represent G. The box for vertex x in G is $B(x) = \prod_{i=1}^{d} [\mathsf{rank}_i(x), \mathsf{rank}_i(p_i(x))]$.

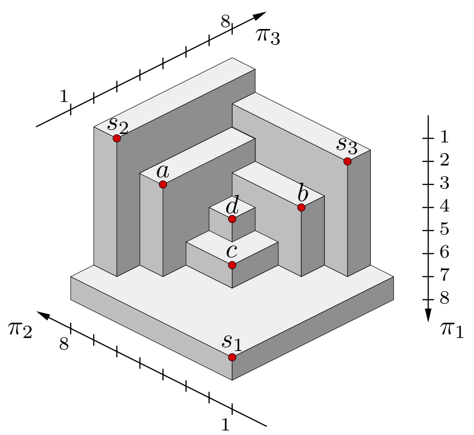

Fig. 2: *The proper touching box representation of the graph from Fig. 1 obtained with our method. The view is from below, i.e., the labeled corners are the minima of the boxes.*

We need to show proper contact between the box $B(x)$ and the box $B(p_i(x))$ for all i. Let $y = p_i(x)$. Since the projection to $B(x)$ and $B(y)$ to dimension i share the point $\mathsf{rank}_i(y)$, it suffices to show that $\mathsf{rank}_j(x) \in (\mathsf{rank}_j(y), \mathsf{rank}_j(p_j(y)))$ for all $j \neq i$. By the 1-of-d-property, $\mathsf{rank}_j(y) < \mathsf{rank}_j(x)$ for all $j \neq i$. So it suffices to check that $\mathsf{rank}_j(x) < \mathsf{rank}_j(p_j(y))$ for all $j \neq i$.

Let $z = p_j(y)$ and suppose $z \prec_j x$. By the 1-of-d-property, $z \prec_k y$ for all $k \neq j$. Since $y \prec_k x$ for all $k \neq i$ transitivity implies that $z \prec_k x$ for all $k \neq i, j$ and by supposition also for $k = j$. Since a d-realizer is an antichain we can conclude that $x \prec_i z$.

It now happens that (x, z) and (x, y) both have the 1-of-d-property and $x \prec_i z \prec_i y$. This however contradicts the choice of $y = p_i(x)$ as the least element of π_i such that (x, y) has the 1-of-d-property. Therefore $x \prec_j z$ as needed for the box contact.

To represent a subgraph of G, remove unneeded boxes and edges from the box representation. To get rid of an edge $(x, p_i(x))$ change the extent of $B(x)$ in dimension i to $[\mathsf{rank}_i(x), \mathsf{rank}_i(p_i(x)) - \varepsilon]$. □

4 Orthogonal surfaces and simplices

In this section we take a more geometric look at the graphs of d-realizers.

With a point $p \in \mathbb{R}^d$ we associate its *cone* $C(p) = \{q \in \mathbb{R}^d : p \leq_{\mathsf{dom}} q\}$. The *filter* $\langle V \rangle$ generated by V is the union of all cones $C(v)$ for $v \in V$. The *orthogonal surface* S_V generated by V is the boundary of $\langle V \rangle$. A point $p \in \mathbb{R}^d$ belongs to S_V if and only if p shares a coordinate with all $v \leq_{\mathsf{dom}} p$, $v \in V$. The generating set V is an antichain if and only if all elements of V appear as minima on S_V.

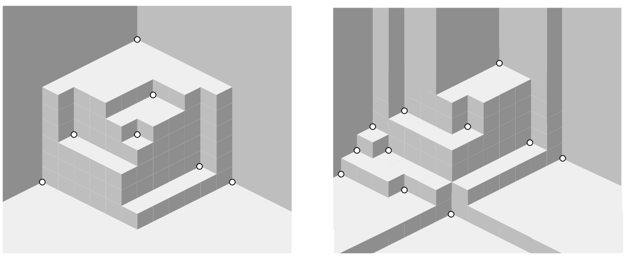

Fig. 3: *Two orthogonal surfaces in $\mathbb{R}^3$: the left one is generated by a suspended antichain in general position; the antichain generating the right one is neither suspended nor in general position. As usual for orthogonal surfaces we take a view from above, the generating points are the minima of the surface.*

Miller [10] observed the connection between Schnyder woods and orthogonal surfaces in $\mathbb{R}^3$. He and subsequently others [4, 9] used orthogonal surfaces to give new proofs for the Brightwell-Trotter theorem about the order dimension of face lattices of 3-polytopes [2]. In fact the dominance order of critical points (maxima, minima, and saddle points) of a 3-dimensional orthogonal surface that is generated by a suspended antichain is the truncated face lattice of a 3-polytope with one facet removed. The converse also holds: every 3-polytope with a facet, selected for removal, has a corresponding orthogonal surface.

The Brightwell-Trotter theorem is an important generalization of Schnyder's dimension theorem. Since orthogonal surfaces can be considered in arbitrary dimensions they provide a direction for generalizing Schnyder structures to higher dimensions. This approach has been taken in [7]. The strongest result in the area is a theorem of Scarf [11] that can be restated as follows: the dominance order of critical points of a d-dimensional orthogonal surface that is generated by a suspended antichain in general position is the truncated face lattice of a simplicial d-polytope with one facet removed. However, the general situation is not nearly as nice as in 3 dimensions. There are simplicial d-polytopes that do not have a corresponding orthogonal surface and if we allow non-general position the dominance order of critical points need not even be a truncated lattice [7].

The orthogonal surface view for graphs given by a d-realizer R is as follows: Embed vertex v at the point p_v whose coordinates are the ranks of v in the realizer. The out-neighbor of v in color i is the vertex w whose cone $C(p_w)$ is first hit by the ray leaving p_v in the ith coordinate direction.

In the 3-dimensional case we can embed every triangulation (graph with a 3-realizer) on an orthogonal surface S_V with a coplanar V, i.e., all $p \in V$ lie in a plane h with normal $\mathbb{1} = (1, 1, 1)$. Identifying h with $\mathbb{R}^2$ we can find the three edges of a vertex v by growing homothetic equilateral triangles with a corner in v until they hit another vertex; Fig. 4 shows an example.

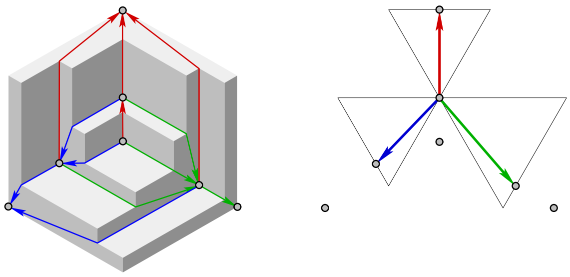

Fig. 4: *The graph from Fig. 1 on a coplanar orthogonal surface and a sketch illustrating how to recover the out-edges of a vertex from the generating set of points in the plane.*

In the same way we may use a set of points in d-space and the homothets of a d-simplex to build a graph from the class defined by $(d+1)$-realizers. The details are as follows: Let Δ be a fixed d-simplex in $\mathbb{R}^d$ and let P be a set of points such that no hyperplane parallel to a facet of Δ contains more than one point (this is the appropriate general position assumption). Let S be the set of corners of a homothet of the dual of Δ that contains P, this is the set of suspensions. Now, for each point $p \in P$ and each corner x of Δ find the unique point q such that there is a homothety that maps Δ to Δ' such that (1) the corner x of Δ' is at p (2) Δ' has no point of P in the interior and (3) q is on the boundary of Δ'. This condition characterizes the edges $p \to q$ of color x in the graph $G_\Delta(P)$.

Problem 1 *Let G be the graph of a d-realizer. Is it always possible to find a point set P in $\mathbb{R}^{d-1}$ such that $G = G_\Delta(P)$?*

There is one class of graphs where we know that the answer to the problem is yes. These are the skeleton graphs of d-dimensional stacked polytopes, also known as simple d-trees. A d-tree is a graph admitting a stacking sequence, i.e., a listing $v_1, v_2, \ldots, v_n$ of the vertices such that $v_1, \ldots, v_{d+1}$ is a clique and for each $j > d+1$ the neighbors of v_j with indices $< j$ induce a clique C_j of size d. A d tree is simple if $C_i \neq C_j$ whenever $i \neq j$, i.e., each d-clique can be used at most once for stacking. If G is a simple d-tree a corresponding point set P can be constructed along the stacking sequence.

In fact, besides this class we know only a few examples of graphs that have a d-realizer with $d > 3$. We know that unlike in the $d = 3$ case we also have non-simple d-trees in the class: Consider a simple d-tree with realizer $(\pi_1, \ldots, \pi_d)$ and let x be a vertex with $\mathsf{deg}(x) = d$, for example the last vertex of the construction sequence has this property. Add a new vertex x' by placing it immediately before x in π_1 and π_2 and immediately after x in all the other π_j. It is easily seen that x and x' have the same neighbors in the same colors, in particular they are stacked over the same clique.

Problem 2 *Characterize the d-trees that have a d-realizer.*

Problem 3 *Find meaningful examples and families of graphs that have a d-realizer.*

Regarding the recognition of graphs that have a d-realizer, we have the criterion that to qualify, a graph G must contain a d-clique of suspensions such that there is an orientation of the edges of G with $\mathsf{out\text{-}deg}(x) = d$ for all non-suspensions x.

Problem 4 *Identify additional obstructions against having a d-realizer.*

Another situation where induced subgraphs of graphs with a $(d+1)$-realizer appear is given by families of interiorly disjoint pairwise homothetic d-simplices in d-space with vertex-facet incidences. To produce a d-realizer for a supergraph add a small d-simplex over each vertex that does not take part in a vertex-facet contact and then use the directions of inward pointing normals of the facets to list the simplices. Figure 5 shows an example in 2 dimensions.

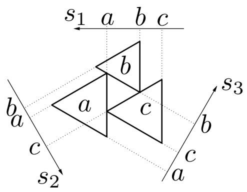

Fig. 5: *A 3-realizer from homothetic triangles.*

Problem 5 *Is it possible to realize every simple d-tree as vertex-facet contact graph of homothetic simplices in $\mathbb{R}^d$?*

Acknowledgments

We started the research reported in this paper during the *Bertinoro Workshop on Graph Drawing 2013*.

References

[1] D. Bremner, W. S. Evans, F. Frati, L. J. Heyer, S. G. Kobourov, W. J. Lenhart, G. Liotta, D. Rappaport, and S. Whitesides, *On representing graphs by touching cuboids*, in Proc. GD'12, LNCS 7704, 2012, 187–198.

[2] G. Brightwell and W. T. Trotter, *The order dimension of convex polytopes*, SIAM J. Discrete Math., 6 (1993), 230–245.

[3] A. L. Buchsbaum, E. R. Gansner, C. M. Procopiuc, and S. Venkatasubramanian, *Rectangular layouts and contact graphs*, ACM Trans. Algorithms, 4 (2008), Art. 8, 28.

[4] S. Felsner, *Geodesic embeddings and planar graphs*, Order, 20 (2003), 135–150.

[5] S. Felsner, *Rectangle and square representations of planar graphs*, in Thirty Essays in Geometric Graph Theory, J. Pach, ed., Springer, 2013, 213–248.

[6] S. Felsner and M. C. Francis, *Contact representations of planar graphs with cubes*, in Proc. ACM Symposium on Computational Geometry, 2011, 315–320.

[7] S. Felsner and S. Kappes, *Orthogonal surfaces and their cp-orders*, Order, 25 (2008), 19–47.

[8] S. Felsner and W. T. Trotter, *Posets and planar graphs*, Journal of Graph Theory, 49 (2005), 273–284.

[9] S. Felsner and F. Zickfeld, *Schnyder woods and orthogonal surfaces*, Discrete & Computational Geometry, 40 (2008), 103–126.

[10] E. Miller, *Planar graphs as minimal resolutions of trivariate monomial ideals*, Documenta Math., 7 (2002), 43–90.

[11] H. Scarf, *The Computation of Economic Equilibria*, vol. 24 of Cowles Foundation Monograph, Yale University Press, 1973.

[12] W. Schnyder, *Planar graphs and poset dimension*, Order, 5 (1989), 323–343.

[13] C. Thomassen, *Plane representations of graphs*, in Progress in graph theory, Bondy and Murty, eds., Acad. Press, 1984, 336–342.

Intersection Graphs of L-Shapes and Segments in the Plane*

Stefan Felsner[1], Kolja Knauer[2], George B. Mertzios[3], and Torsten Ueckerdt[4]

[1] Institut für Mathematik, Technische Universität Berlin, Germany
[2] LIRMM, Université Montpellier 2, France
[3] School of Engineering and Computing Sciences, Durham University, UK
[4] Department of Mathematics, Karlsruhe Institute of Technology, Germany
felsner@math.tu-berlin.de, kolja.knauer@math.univ-montp2.fr,
george.mertzios@durham.ac.uk, torsten.ueckerdt@kit.edu

Abstract. An L-shape is the union of a horizontal and a vertical segment with a common endpoint. These come in four rotations: L, Γ, ⅃ and ⅂. A k-bend path is a simple path in the plane, whose direction changes k times from horizontal to vertical. If a graph admits an intersection representation in which every vertex is represented by an L, an L or Γ, a k-bend path, or a segment, then this graph is called an {L}-graph, {L, Γ}-graph, B_k-VPG-graph or SEG-graph, respectively. Motivated by a theorem of Middendorf and Pfeiffer [Discrete Mathematics, 108(1):365–372, 1992], stating that every {L, Γ}-graph is a SEG-graph, we investigate several known subclasses of SEG-graphs and show that they are {L}-graphs, or B_k-VPG-graphs for some small constant k. We show that all planar 3-trees, all line graphs of planar graphs, and all full subdivisions of planar graphs are {L}-graphs. Furthermore we show that all complements of planar graphs are B_{19}-VPG-graphs and all complements of full subdivisions are B_2-VPG-graphs. Here a full subdivision is a graph in which each edge is subdivided at least once.

Keywords: Intersection graphs, segment graphs, co-planar graphs, k-bend VPG-graphs, planar 3-trees.

1 Introduction and Motivation

A **segment intersection graph**, SEG-graph for short, is a graph that can be represented as follows. Vertices correspond to straight-line segments in the plane and two vertices are adjacent if and only if the corresponding segments intersect. Such representations are called SEG-*representations* and, for convenience, the class of all SEG-graphs is denoted by SEG. SEG-graphs are an important subject of study strongly motivated from an algorithmic point of view. Indeed, having an intersection representation of a graph (in applications graphs often come

* This work was partially supported by (i) the DFG ESF EuroGIGA projects COMPOSE and GraDR, (ii) the EPSRC Grant EP/K022660/1 and (iii) the ANR Project EGOS: ANR-12-JS02-002-01.

E. Csuhaj-Varjú et al. (Eds.): MFCS 2014, Part II, LNCS 8635, pp. 299–310, 2014.

along with such a given representation) may allow for designing better or faster algorithms for optimization problems that are hard for general graphs, such as finding a maximum clique in interval graphs.

More than 20 years ago, Middendorf and Pfeiffer [24], considered intersection graphs of **axis-aligned L-shapes** in the plane, where an axis-aligned L-shape is the union of a horizontal and a vertical segment whose intersection is an endpoint of both. In particular, L-shapes come in four possible rotations: $\llcorner$, $\ulcorner$, $\lrcorner$, and $\urcorner$. For a subset X of these four rotations, e.g., $X = \{\llcorner\}$ or $X = \{\llcorner, \ulcorner\}$, we call a graph an *$X$-graph* if it admits an *X-representation*, i.e., vertices can be represented by L-shapes from X in the plane, each with a rotation from X, such that two vertices are adjacent if and only if the corresponding L-shapes intersect. Similarly to SEG, we denote the class of all X-graphs by X. The question if an intersection representation with polygonal paths or pseudo-segments can be *stretched* into a SEG-representation is a classical topic in combinatorial geometry and Oriented Matroid Theory. Middendorf and Pfeiffer prove the following interesting relation between intersection graphs of segments and L-shapes.

Theorem 1 (Middendorf and Pfeiffer [24]). *Every* $\{\llcorner, \ulcorner\}$*-representation has a combinatorially equivalent* SEG*-representation.*

This theorem is best-possible in the sense that there are examples of $\{\llcorner, \urcorner\}$-graphs which are no SEG-graphs [7, 24], i.e., such $\{\llcorner, \urcorner\}$-representations cannot be stretched. We feel that Theorem 1, which of course implies that $\{\llcorner, \ulcorner\} \subseteq$ SEG, did not receive a lot of attention in the active field of SEG-graphs. In particular, one could use Theorem 1 to prove that a certain graph class $\mathcal{G}$ is contained in SEG by showing that $\mathcal{G}$ is contained in $\{\llcorner, \ulcorner\}$. For example, very recently Pawlik *et al.* [25] discovered a class of triangle-free SEG-graphs with arbitrarily high chromatic number, disproving a famous conjecture of Erdős [18], and it is in fact easier to see that these graphs are $\{\llcorner\}$-graphs than to see that they are SEG-graphs. To the best of our knowledge, the stronger result $\mathcal{G} \subseteq \{\llcorner, \ulcorner\}$ has never been shown for any non-trivial graph class $\mathcal{G}$. In this paper we initiate this research direction. We consider several graph classes which are known to be contained in SEG and show that they are actually contained in $\{\llcorner\}$, which is a proper subclass of $\{\llcorner, \ulcorner\}$ [7].

Whenever a graph is not known (or known not) to be an intersection graph of segments or axis-aligned L-shapes, one often considers natural generalizations of these intersection representations. Asinowski *et al.* [3] introduced **intersection graphs of axis-aligned k-bend paths** in the plane, called B_k-VPG-graphs. An (axis-aligned) k-bend path is a simple path in the plane, whose direction changes k times from horizontal to vertical. Clearly, B_1-VPG-graphs are precisely intersection graphs of all four L-shapes; the union of B_k-VPG-graphs for all $k \geq 0$ is exactly the class STRING of intersection graphs of simple curves in the plane [3]. Now if a graph $G \notin$ SEG is a B_k-VPG-graph for some small k, then one might say that G is "not far from being a SEG-graph".

Our Results and Related Work

Let us denote the class of all planar graphs by PLANAR. A recent celebrated result of Chalopin and Gonçalves [6] states that PLANAR $\subset$ SEG, which was conjectured by Scheinerman [26] in 1984. However, their proof is rather involved and there is not much control over the kind of SEG-representations. Here we give an easy proof for a non-trivial subclass of planar graphs, namely *planar 3-trees*. A *3-tree* is an edge-maximal graph of treewidth 3. Every 3-tree can be built up starting from the clique K_4 and adding new vertices, one at a time, whose neighborhood in the so-far constructed graph is a triangle.

Theorem 2. *Every planar* 3*-tree is an* $\{\mathsf{L}\}$*-graph.*

It remains open to generalize Theorem 2 to planar graphs of treewidth 3 (i.e., subgraphs of planar 3-trees). On the other hand it is easy to see that graphs of treewidth at most 2 are $\{\mathsf{L}\}$-graphs [8]. Chaplick and the last author show in [9] that planar graphs are B_2-VPG-graphs, improving on an earlier result of Asinowski *et al.* [3]. In [9] it is also conjectured that PLANAR $\subset \{\mathsf{L}\}$, which with Theorem 1 would imply the main result of [6], i.e., PLANAR $\subset$ SEG.

Considering line graphs of planar graphs, one easily sees that these graphs are SEG-graphs. Indeed, a straight-line drawing of a planar graph G can be interpreted as a SEG-representation of the line graph $L(G)$ of G, which has the edges of G as its vertices and pairs of incident edges as its edges. We prove the following strengthening result.

Theorem 3. *The line graph of every planar graph is an* $\{\mathsf{L}\}$*-graph.*

Kratochvíl and Kuběna [21] consider the class of all complements of planar graphs (co-planar graphs), CO-PLANAR for short. They show that CO-PLANAR are intersection graphs of convex sets in the plane, and ask whether CO-PLANAR $\subset$ SEG. As the Independent Set Problem in planar graphs is known to be NP-complete [15], Max Clique is NP-complete for any graph class $\mathcal{G} \supseteq$ CO-PLANAR, e.g., intersection graphs of convex sets. Indeed, the longstanding open question whether Max Clique is NP-complete for SEG [22] has recently been answered affirmatively by Cabello, Cardinal and Langerman [4] by showing that every planar graph has an even subdivision whose complement is a SEG-graph. The subdivision is essential in the proof of [4], as it still remains an open problem whether CO-PLANAR $\subset$ SEG [21]. The largest subclass of CO-PLANAR known to be in SEG is the class of complements of partial 2-trees [14]. Here we show that all co-planar graphs are "not far from being SEG-graphs".

Theorem 4. *Every co-planar graph is a* B_{19}*-VPG graph.*

Theorem 4 implies that Max Clique is NP-complete for B_k-VPG-graphs with $k \geq 19$. On the other hand, the Max Clique problem for B_0-VPG-graphs can be solved in polynomial time, while Vertex Colorability remains NP-complete but allows for a 2-approximation [3]. Middendorf and Pfeiffer [24] show that the complement of any *even subdivision* of any graph, i.e., every edge is subdivided with a non-zero even number of vertices, is an {L, ꓶ}-graph. This implies that Max Clique is NP-complete even for {L, ꓶ}-graphs.

We consider *full subdivisions* of graphs, that is, a subdivision H of a graph G where each edge of G is subdivided at least once. It is not hard to see that a full subdivision H of G is in STRING if and only if G is planar, and that if G is planar, then H is actually a SEG-graph. Here we show that this can be further strengthened, namely that H is in an {L}-graph. Moreover, we consider the complement of a full subdivision H of an arbitrary graph G, which is in STRING but not necessarily in SEG. Here, similar to the result of Middendorf and Pfeiffer [24] on even subdivisions we show that such a graph H is "not far from being SEG-graph".

Theorem 5. *Let H be a full subdivision of a graph G.*

(i) If G is planar, then H is an {L}*-graph.*
(ii) If G is any graph, then the complement of H is a B_2-VPG*-graph.*

The graph classes considered in this paper are illustrated in Figure 1. We shall prove Theorems 2, 3, 4 and 5 in Sections 2, 3, 4 and 5, respectively, and conclude with some open questions in Section 6. Due to lack of space, the full proof of Theorem 2 is given in the full version [13].

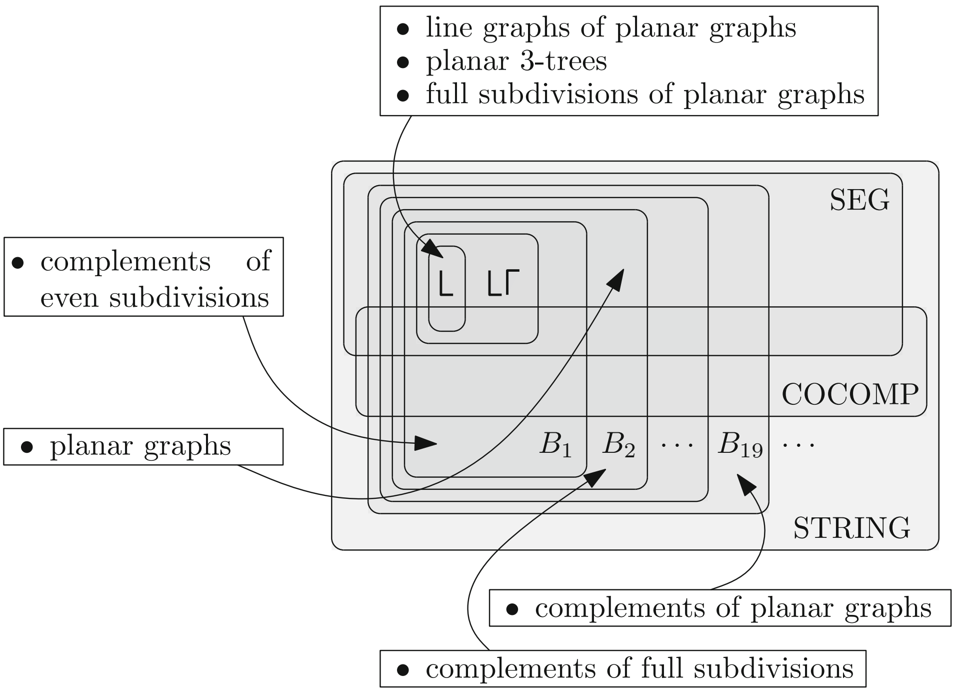

Fig. 1. Graph classes considered in this paper

Related Representations

In the context of *contact representations*, where distinct segments or k-bend paths may not share interior points, it is known that every contact SEG-representation has a combinatorially equivalent contact B_1-VPG-representation, but not vice versa [20]. Contact SEG-graphs are exactly planar Laman graphs and their subgraphs [10], which includes for example all triangle-free planar graphs. Very recently, contact {L}-graphs have been characterized [8]. Necessary and sufficient conditions for stretchability of a contact system of pseudo-segments are known [1, 11].

Let us also mention the closely related concept of *edge*-intersection graphs of paths in a grid (EPG-graphs) introduced by Golumbic *et al.* [16]. There are some notable differences, starting from the fact that *every* graph is an EPG-graph [16]. Nevertheless, analogous questions to the ones posed about VPG-representations of STRING-graphs are posed about EPG-representations of general graphs. In particular, there is a strong interest in finding representations using paths with few bends, see [19] for a recent account.

2 Proof of Theorem 2

Proof (main idea). Let G be a plane 3-tree with a xed plane embedding. We construct an {L}-representation of G satisfying the additional property that for every inner triangular face $\{a, b, c\}$ of G there exists a subset of the plane, called the *private region* of the face, that intersects only the L-paths for a, b and c, and no other L-path. We remark that this technique has also been used by Chalopin *et al.* [5] and refer to Figure 2 for an illustration. □

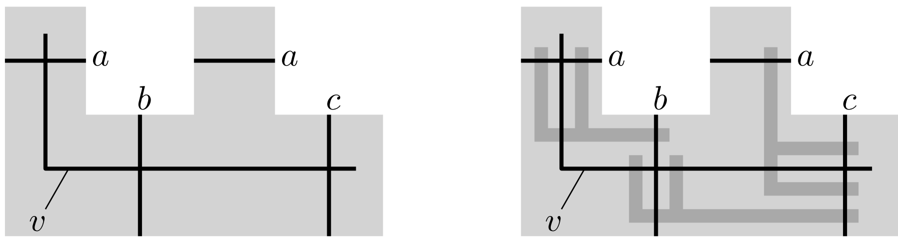

(a) (b)

Fig. 2. (a) Introducing an L-shape for vertex v into the private region for the triangle $\{a, b, c\}$. (b) Identifying a pairwise disjoint private regions for the facial triangles $\{a, b, v\}$, $\{a, c, v\}$ and $\{b, c, v\}$.

3 Proof of Theorem 3

Proof. Without loss of generality let G be a maximally planar graph with a fixed plane embedding. (Line graphs of subgraphs of G are induced subgraphs of $L(G)$.) Then G admits a so-called *canonical ordering* –first defined in [12]–, namely an ordering $v_1, \ldots, v_n$ of the vertices of G such that

– Vertices v_1, v_2, v_n form the outer triangle of G in clockwise order. (We draw G such that v_1, v_2 are the highest vertices.)

– For $i = 3, \dots, n$ vertex v_i lies in the outer face of the induced embedded subgraph $G_{i-1} = G[v_1, \dots, v_{i-1}]$. Moreover, the neighbors of v_i in G_{i-1} form a path on the outer face of G_{i-1} with at least two vertices.

We shall construct an $\{\mathsf{L}\}$-representation of $L(G)$ along a fixed canonical ordering $v_1, \dots, v_n$ of G. For every $i = 2, \dots, n$ we shall construct an $\{\mathsf{L}\}$-representation of $L(G_i)$ with the following additional properties.

For every outer vertex v of G_i we maintain an auxiliary bottomless rectangle $R(v)$, i.e., an axis-aligned rectangle with bottom-edge at $-\infty$, such that:

– $R(v)$ intersects the horizontal segments of precisely those rectilinear paths for edges in G_i incident to v.
– $R(v)$ does not contain any bends or endpoints of any path for an edge in G_i and does not intersect any $R(w)$ for $w \neq v$.
– the left-to-right order of the bottomless rectangles matches the order of vertices on the counterclockwise outer v_1, v_2-path of G_i.

The bottomless rectangles act as placeholders for the upcoming vertices of $L(G)$. Indeed, all upcoming intersections of paths will be realized inside the corresponding bottomless rectangles. For $i = 2$, the graph G_i consist only of the edge v_1v_2. Hence an $\{\mathsf{L}\}$-representation of the one-vertex graph $L(G_2)$ consists of only one L-shape and two disjoint bottomless rectangles $R(v_1)$, $R(v_2)$ intersecting its horizontal segment.

For $i \geq 3$, we shall start with an $\{\mathsf{L}\}$-representation of $L(G_{i-1})$. Let $(w_1, \dots, w_k)$ be the counterclockwise outer path of G_{i-1} that corresponds to the neighbors of v_i in G_{i-1}. The corresponding bottomless rectangles $R(w_1), \dots, R(w_k)$ appear in this left-to-right order. See Figure 3 for an illustration. For every edge v_iw_j, $j = 1, \dots, k$ we define an L-shape $P(v_iw_j)$ whose vertical segment is contained in the interior of $R(w_j)$ and whose horizontal segment ends in the interior of $R(w_k)$. Moreover, the upper end and lower end of the vertical segment of $P(v_iw_j)$ lies on the top side of $R(w_j)$ and below all L-shapes for edges in G_{i-1}, respectively. Finally, the bend and right end of $P(v_iw_j)$ is placed above the bend of $P(v_iw_{j+1})$ and to the right of the right end of $P(v_iw_{j+1})$ for $j = 1, \dots, k-1$, see Figure 3.

It is straightforward to check that this way we obtain an $\{\mathsf{L}\}$-representation of $L(G_i)$. So it remains to find a set of bottomless rectangles, one for each outer vertex of G_i, satisfying our additional property. We set $R'(v) = R(v)$ for every $v \in V(G_i) \setminus \{v_i, w_1, \dots, w_k\}$ since these are kept unchanged. Since $R(w_1)$ and $R(w_k)$ are not valid anymore, we define a new bottomless rectangle $R'(w_1) \subset R(w_1)$ such that $R'(w_1)$ is crossed by all horizontal segments that cross $R(w_1)$ and additionally the horizontal segment of $P(v_iw_1)$. Similarly, we define $R'(w_k) \subset R(w_k)$. And finally, we define a new bottomless rectangle $R'(v_i) \subset R(w_k)$ in such a way that it is crossed by the horizontal segments of exactly $P(v_iw_1), \dots, P(v_iw_k)$. Note that for $1 < j < k$ the outer vertex w_j of G_{i-1} is not an outer vertex of G_i. Then $\{R'(v) \mid v \in v(G_i)\}$ has the desired property. See again Figure 3. □

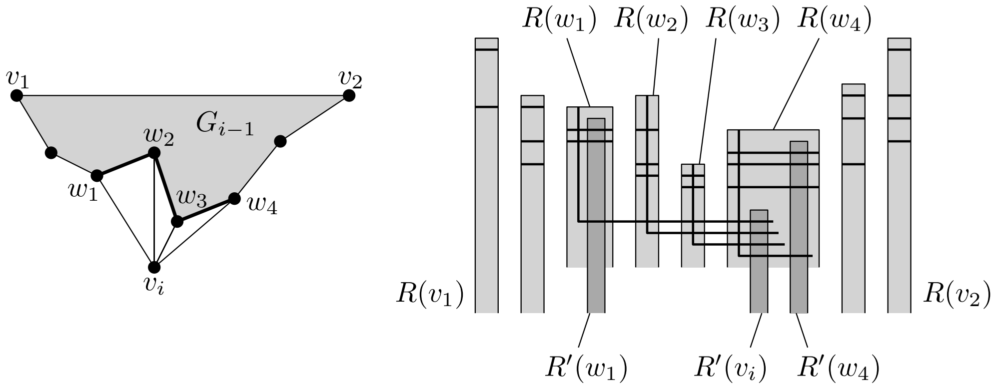

Fig. 3. Along a canonical ordering a vertex v_i is added to G_{i-1}. For each edge between v_i and a vertex in G_{i-1} an L-shape is introduced with its vertical segment in the corresponding bottomless rectangle. The three new bottomless rectangles $R'(w_1), R'(v_i), R'(w_k)$ are highlighted.

4 Proof of Theorem 4

Proof. Let $G = (V, E)$ be any planar graph. We shall construct a B_k-VPG representation of the complement $\bar{G}$ of G for some constant k that is independent of G. Indeed, $k = 19$ is enough. To find the VPG representation we make use of two crucial properties of G: A) G is 4-colorable and B) G is 5-degenerate. Indeed, our construction gives a B_{2d+9}-VPG representation for the complement of any 4-colorable d-degenerate graph. Here a graph is called *d-degenerate* if it admits a vertex ordering such that every vertex has at most d neighbors with smaller index.

Consider any 4-coloring of G with color classes V_1, V_2, V_3, V_4. Further let $\sigma = (v_1, \ldots, v_n)$ be an order of the vertices of V witnessing the degeneracy of G, i.e., for each v_i there are at most 5 neighbors v_j of v_i with $j < i$. We call these neighbors the *back neighbors of* v_i. Consider any ordered pair of color classes, say (V_1, V_2), and denote $W = V_1 \cup V_2$, together with the vertex orders inherited from the order of vertices in V, i.e., $\sigma|_{V_1} = \sigma_1 = (v_1, \ldots, v_{|V_1|})$ and $\sigma|_{V_2} = \sigma_2 = (w_1, \ldots, w_{|V_2|})$. Further consider the axis-aligned rectangle $R = [0, A] \times [0, A]$, where $A = 2(|W| + 2)$. For illustration we divide R into four quarters $[0, A/2] \times [0, A/2]$, $[0, A/2] \times [A/2, A]$, $[A/2, A] \times [0, A/2]$ and $[A/2, A] \times [A/2, A]$. We define a monotone increasing path $Q(v)$ for each $v \in W$ as follows. See Figure 4 for an illustration.

- For $v \in V_1$ let $\{\sigma_2(i_1), \ldots, \sigma_2(i_k)\}$, $i_1 < \cdots < i_k$, be the back neighbors of v in V_2 and $i^* = \max\{0\} \cup \{\sigma_2^{-1}(w) \mid w \in V_2, \sigma^{-1}(w) < \sigma^{-1}(v)\}$ be the largest index with respect to σ_2 of a vertex in V_2 that comes before v in σ or $i^* = 0$ if there is no such vertex. Then we define the path $Q(v)$ so that it starts at $(1, 0)$, uses the horizontal lines at $y = 2i_j - 1$ for $j = 1, \ldots, k$, $y = 2i^* + 1$ and $y = A - 2\sigma_1(v)$ in that order, uses the vertical lines at $x = 1$, $x = 2i_j + 1$ for $j = 1, \ldots, k$ and $x = A - 2\sigma_1(v)$ in that order, and finally ends at $(A, A - 2\sigma_1(v))$.

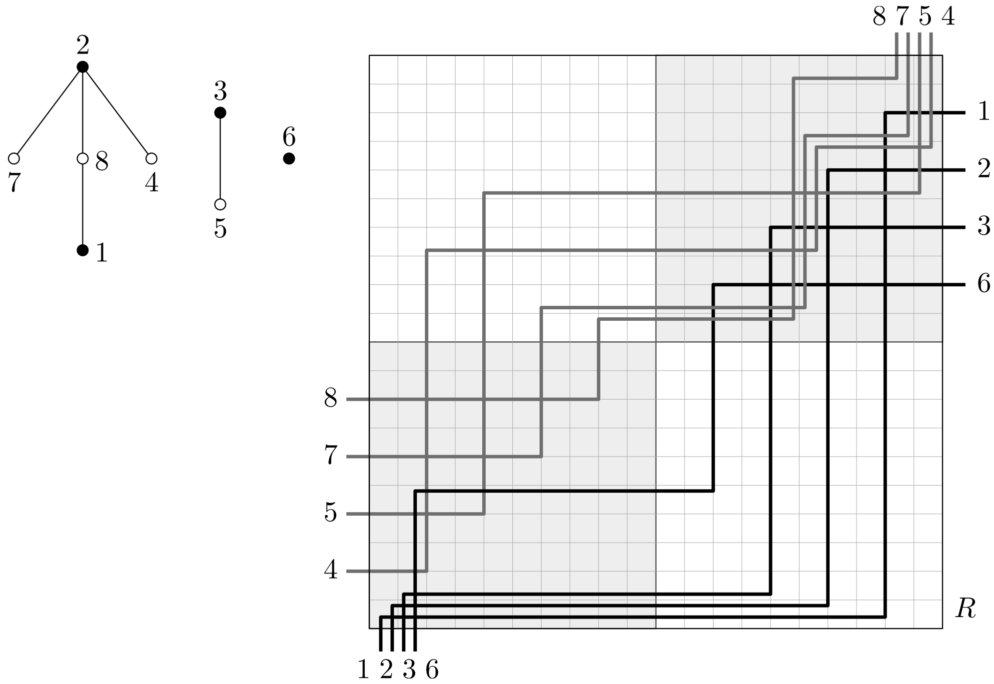

Fig. 4. The induced subgraph $G[W]$ for two color classes $W = V_1 \cup V_2$ of a planar graph G and a VPG representation of its complement $\bar{G}[W]$ in the rectangle $[0, 2(|W|+2)] \times [0, 2(|W|+2)]$

Note that $Q(v)$ avoids the top-left quarter of R, has exactly one bend at $(A - 2\sigma_1(v), A - 2\sigma_1(v))$ in the top-right quarter, and goes above the point $(2i, 2i)$ in the bottom-left quarter if and only if $i \neq i_1, \ldots, i_k$ and $i \leq i^*$.

– For $w_i \in V_2$ the path $P(w_i)$ is defined analogous after rotating the rectangle R by 180 degrees and swapping the roles of V_1 and V_2.

It is straightforward to check that $\{Q(v) \mid v \in W\}$ is a VPG representation of $\bar{G}[W]$ completely contained in R, where each $Q(v)$ starts and ends at the boundary of R and has at most $3 + 2k$ bends, where k is the number of back neighbors of v in W.

Now we have defined for each pair of color classes $V_i \cup V_j$ a VPG-representation of $\bar{G}[V_i \cup V_j]$. For every vertex $v \in V$ we have defined three Q-paths, one for each colors class that v is not in. In total the three Q-paths for the same vertex v have at most $9 + 2k \leq 19$ bends, where $k \leq 5$ is the back degree of v. It remains to place the six representations of $\bar{G}[V_i \cup V_j]$ non-overlapping and to "connect" the three Q-paths for each vertex in such a way that connections for vertices of different color do not intersect. This can easily be done with two extra bends per paths, basically because K_4 is planar (we refer to Figure 5 for one way to do this). Finally, note that the first and last segment of every path in the representation can be omitted, yielding the claimed bound. □

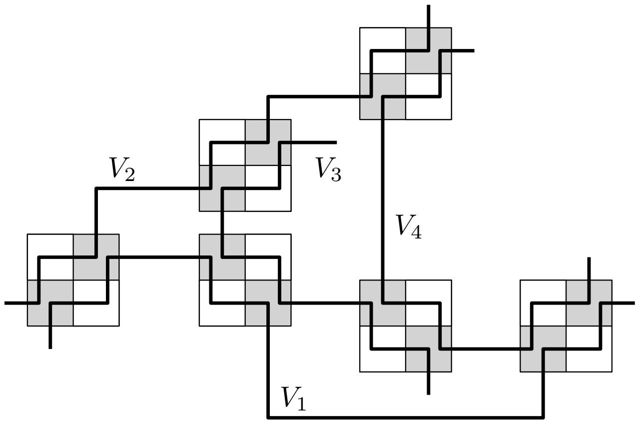

Fig. 5. Interconnecting the VPG representations of $\bar{G}[V_i \cup V_j]$ by adding at most two bends for each vertex. The set of paths corresponding to color class V_i is indicated by a single path labeled V_i, $i = 1, 2, 3, 4$.

5 Proof of Theorem 5

Proof. Let G be any graph and H arise from G by subdividing each edge at least once. Without loss of generality we may assume that every edge of G is subdivided exactly once or twice. Indeed, if an edge e of G is subdivided three times or more, then H can be seen as a full subdivision of the graph G' that arises from G by subdividing e once.

(i) Assuming that G is planar, we shall find an $\{\mathsf{L}\}$-representation of H as follows. Without loss of generality G is maximally planar. We consider a bar visibility representation of G, i.e., vertices of G are disjoint horizontal segments in the plane and edges are disjoint vertical segments in the plane whose endpoints are contained in the two corresponding vertex segments and which are disjoint from all other vertex segments. Such a representation for a planar triangulation exists e.g. by [23]. See Figure 6 for an illustration.

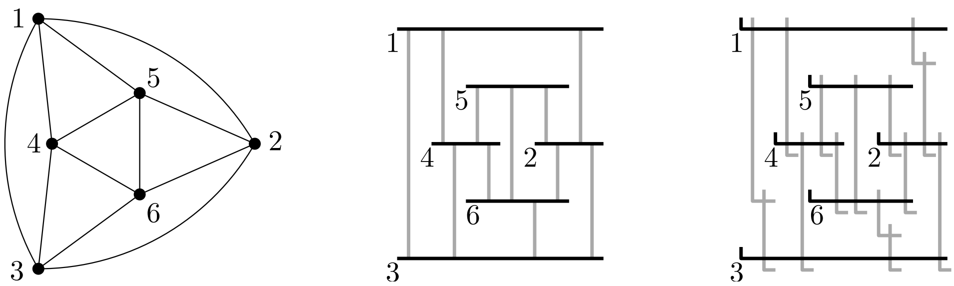

Fig. 6. A planar graph G on the left, a bar visibility representation of G in the center, and an $\{\mathsf{L}\}$-representation of a full division of G on the right. Here, the edges $\{1, 2\}$, $\{1, 3\}$ and $\{3, 6\}$ are subdivided twice.

It is now easy to interpret every segment as an L, and replace an segment corresponding to edge that is subdivided twice by two L-shapes. Let us simply refer to Figure 6 again.

(ii) Now assume that $G = (V, E)$ is any graph. We shall construct a B_2-VPG representation of the complement $\bar{H}$ of $H = (V \cup W, E')$ with monotone increasing paths only. First, we represent the clique $\bar{H}[V]$. Let $V = \{v_1, \ldots, v_n\}$ and define for $i = 1, \ldots, n$ the 2-bend path $P(v_i)$ for vertex v_i to start at $(i, 0)$, have bends at (i, i) and $(i + n, i)$, and end at $(i + n, n + 1)$. See Figure 7 for an illustration. For convenience, let us call these paths *v-paths*.

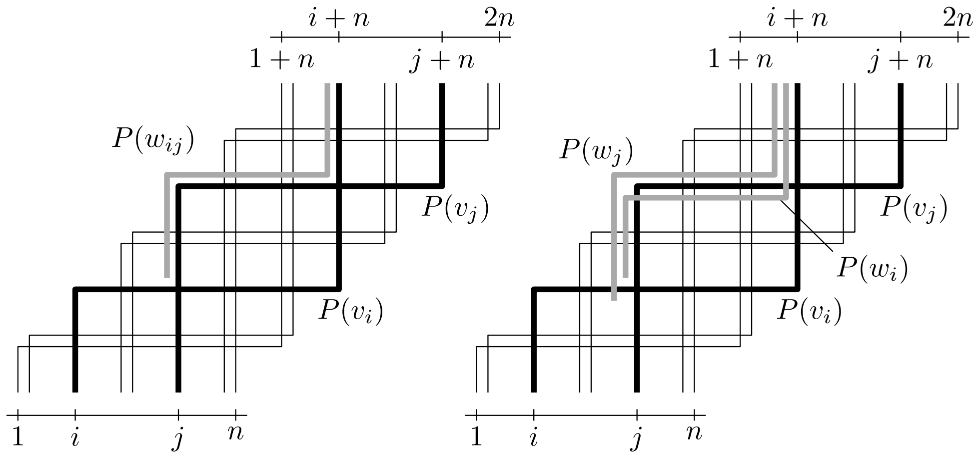

Fig. 7. Left: Inserting the path $P(w_{ij})$ for a single vertex w_{ij} subdividing the edge v_iv_j in G. Right: Inserting the paths $P(w_i)$ and $P(w_j)$ for two vertices w_i, w_j subdividing the edge v_iv_j in G.

Next, we define for every edge of G the 2-bend paths for the one or two corresponding subdivision vertices in $\bar{H}$. We call these paths *w-paths*. So let v_iv_j be any edge of G with $i < j$. We distinguish two cases.

Case 1. The edge v_iv_j is subdivided by only one vertex w_{ij} in H. We define the w-path $P(w_{ij})$ to start at $(j - \frac{1}{4}, i + \frac{1}{4})$, have bends at $(j - \frac{1}{4}, j + \frac{1}{4})$ and $(i + n - \frac{1}{4}, j + \frac{1}{4})$, and end at $(i + n - \frac{1}{4}, n + 1)$, see the left of Figure 7.

Case 2. The edge v_iv_j is subdivided by two vertices w_i, w_j with $v_iw_i, v_jw_j \in E(H)$. We define the start, bends and end of the w-path $P(w_i)$ to be $(j - \frac{1}{4}, i + \frac{1}{4})$, $(j - \frac{1}{4}, j - \frac{1}{4})$, $(i + n - \frac{1}{4}, j - \frac{1}{4})$ and $(i + n - \frac{1}{4}, n + 1)$, respectively. The start, bends and end of the w-path $P(w_j)$ are $(j - \frac{1}{2}, i - \frac{1}{4})$, $(j - \frac{1}{2}, j + \frac{1}{4})$, $(i + n - \frac{1}{2}, j + \frac{1}{4})$ and $(i + n - \frac{1}{2}, n + 1)$, respectively. See the right of Figure 7.

It is easy to see that every w-path $P(w)$ intersects every v-path, except for the one or two v-paths corresponding to the neighbors of w in H. Moreover, the two w-paths in Case 2 are disjoint. It remains to check that the w-paths for distinct edges of G mutually intersect. To this end, note that every w-path for edge v_iv_j starts near (j, i), bends near (j, j) and $(i + n, j)$ and ends near $(i+n,n)$. Consider two w-paths P and P' that start at (j,i) and (j',i'), respectively, and bend near (j,j) and (j',j'), respectively. If $j=j'$ then it is easy to check that P and P' intersect near (j,j). Otherwise, let $j'>j$. Now if $j>i'$, then P and P' intersect near (j',i), and if $j \leq i'$, then P and P' intersect near $(i+n,j')$.

Hence we have found a B_2-VPG-representation of $\bar{H}$, as desired. Let us remark, that in this representation some w-paths intersect non-trivially along some horizontal or vertical lines, i.e., share more than a finite set of points. However, this can be omitted by a slight and appropriate perturbation of endpoints and bends of w-paths. □

6 Conclusions and Open Problems

Motivated by Middendorf and Pfeiffer's theorem (Theorem 1 in [24]) that every $\{\mathsf{L},\Gamma\}$-representation can be stretched into a SEG-representation, we considered the question which subclasses of SEG-graphs are actually $\{\mathsf{L},\Gamma\}$-graphs, or even $\{\mathsf{L}\}$-graphs. We proved that this is indeed the case for several graph classes related to planar graphs. We feel that the question whether PLANAR $\subset \{\mathsf{L},\Gamma\}$, as already conjectured [9], is of particular importance. After all, this, together with Theorem 1, would give a new proof for the fact that PLANAR $\subset$ SEG.

Open Problem 1. *Each of the following is open.*

(i) *When can a* B_1-VPG-*representation be stretched into a combinatorially equivalent* SEG-*representation?*
(ii) *Is* $\{\mathsf{L},\Gamma\} = \text{SEG} \cap B_1\text{-VPG}$*?*
(iii) *Is every planar graph an* $\{\mathsf{L}\}$*-graph, or* B_1-VPG-*graph?*
(iv) *Does every planar graph admit an even subdivision whose complement is an* $\{\mathsf{L}\}$*-graph, or* B_1-VPG-*graph?*
(v) *Recognizing* B_k*-VPG graphs is known to be NP-complete for each* $k \geq 0$ *[7]. What is the complexity of recognizing* $\{\mathsf{L}\}$*-graphs, or* $\{\mathsf{L},\Gamma\}$*-graphs?*

References

1. Aerts, N., Felsner, S.: Straight line triangle representations. In: Wismath, S., Wolff, A. (eds.) GD 2013. LNCS, vol. 8242, pp. 119–130. Springer, Heidelberg (2013)
2. Alon, N., Scheinerman, E.: Degrees of freedom versus dimension for containment orders. Order 5, 11–16 (1988)
3. Asinowski, A., Cohen, E., Golumbic, M.C., Limouzy, V., Lipshteyn, M., Stern, M.: Vertex intersection graphs of paths on a grid. J. Graph Algorithms Appl. 16(2), 129–150 (2012)
4. Cabello, S., Cardinal, J., Langerman, S.: The clique problem in ray intersection graphs. Discrete & Computational Geometry 50(3), 771–783 (2013)
5. Chalopin, J., Gonçalves, D., Ochem, P.: Planar graphs have 1-string representations. Discrete & Computational Geometry 43(3), 626–647 (2010)

6. Chalopin, J., Gonçalves, D.: Every planar graph is the intersection graph of segments in the plane: extended abstract. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, pp. 631–638 (2009)
7. Chaplick, S., Jelínek, V., Kratochvíl, J., Vyskočil, T.: Bend-bounded path intersection graphs: Sausages, noodles, and waffles on a grill. In: Golumbic, M.C., Stern, M., Levy, A., Morgenstern, G. (eds.) WG 2012. LNCS, vol. 7551, pp. 274–285. Springer, Heidelberg (2012)
8. Chaplick, S., Kobourov, S.G., Ueckerdt, T.: Equilateral L-contact graphs. In: Brandstädt, A., Jansen, K., Reischuk, R. (eds.) WG 2013. LNCS, vol. 8165, pp. 139–151. Springer, Heidelberg (2013)
9. Chaplick, S., Ueckerdt, T.: Planar graphs as VPG-graphs. Journal of Graph Algorithms and Applications 17(4), 475–494 (2013)
10. de Fraysseix, H., Ossona de Mendez, P.O.: Representations by contact and intersection of segments. Algorithmica 47(4), 453–463 (2007)
11. de Fraysseix, H., de Mendez, P.O.: Stretching of Jordan arc contact systems. Discrete Applied Mathematics 155(9), 1079–1095 (2007)
12. De Fraysseix, H., Pach, J., Pollack, R.: How to draw a planar graph on a grid. Combinatorica 10(1), 41–51 (1990)
13. Felsner, S., Knauer, K., Mertzios, G.B., Ueckerdt, T.: Intersection graphs of L-shapes and segments in the plane. arXiv preprint arXiv:1405.1476 (2014)
14. Francis, M.C., Kratochvíl, J., Vyskočil, T.: Segment representation of a subclass of co-planar graphs. Discrete Mathematics 312(10), 1815–1818 (2012)
15. Garey, M.R., Johnson, D.S.: Computers and Intractability: A guide to the theory of NP-completeness. W.H. Freeman (1979)
16. Golumbic, M.C., Lipshteyn, M., Stern, M.: Edge intersection graphs of single bend paths on a grid. Networks 54(3), 130–138 (2009)
17. Golumbic, M.C., Rotem, D., Urrutia, J.: Comparability graphs and intersection graphs. Discrete Mathematics 43(1), 37–46 (1983)
18. Gyárfás, A.: Problems from the world surrounding perfect graphs. Zastos. Mat. 19(3-4), 413–441 (1987)
19. Heldt, D., Knauer, K., Ueckerdt, T.: Edge-intersection graphs of grid paths: The bend-number. Discrete Appl. Math. 167, 144–162 (2014)
20. Kobourov, S.G., Ueckerdt, T., Verbeek, K.: Combinatorial and geometric properties of planar Laman graphs. In: SODA, pp. 1668–1678. SIAM (2013)
21. Kratochvíl, J., Kuběna, A.: On intersection representations of co-planar graphs. Discrete Mathematics 178(1-3), 251–255 (1998)
22. Kratochvíl, J., Matousek, J.: Intersection graphs of segments. Journal of Combinatorial Theory, Series B 62(2), 289–315 (1994)
23. Luccio, F., Mazzone, S., Wong, C.K.: A note on visibility graphs. Discrete Mathematics 64(2-3), 209–219 (1987)
24. Middendorf, M., Pfeiffer, F.: The max clique problem in classes of string-graphs. Discrete Mathematics 108(1), 365–372 (1992)
25. Pawlik, A., Kozik, J., Krawczyk, T., Lasoń, M., Micek, P., Trotter, W.T., Walczak, B.: Triangle-free intersection graphs of line segments with large chromatic number. Journal of Combinatorial Theory, Series B (2013)
26. Scheinerman, E.R.: Intersection classes and multiple intersection parameters of graphs. PhD thesis, Princeton University (1984)
27. Warren, H.E.: Lower bounds for approximation by nonlinear manifolds. Trans. Amer. Math. Soc. 133, 167–178 (1968)

Combinatorial Properties of Triangle-Free Rectangle Arrangements and the Squarability Problem

Jonathan Klawitter[1,2], Martin Nöllenburg[3]( ), and Torsten Ueckerdt[2]

[1] Institut für Theoretische Informatik, Karlsruhe Institute of Technology, Karlsruhe, Germany
[2] Institut für Algebra und Geometrie, Karlsruhe Institute of Technology, Karlsruhe, Germany
torsten.ueckerdt@kit.edu
[3] Algorithms and Complexity Group, TU Wien, Vienna, Austria
noellenburg@ac.tuwien.ac.at

Abstract. We consider arrangements of axis-aligned rectangles in the plane. A geometric arrangement specifies the coordinates of all rectangles, while a combinatorial arrangement specifies only the respective intersection type in which each pair of rectangles intersects. First, we investigate combinatorial contact arrangements, i.e., arrangements of interior-disjoint rectangles, with a triangle-free intersection graph. We show that such rectangle arrangements are in bijection with the 4-orientations of an underlying planar multigraph and prove that there is a corresponding geometric rectangle contact arrangement. Using this, we give a new proof that every triangle-free planar graph is the contact graph of such an arrangement. Secondly, we introduce the question whether a given rectangle arrangement has a combinatorially equivalent square arrangement. In addition to some necessary conditions and counterexamples, we show that rectangle arrangements pierced by a horizontal line are squarable under certain sufficient conditions.

1 Introduction

We consider arrangements of axis-aligned rectangles and squares in the plane. Besides *geometric rectangle arrangements*, in which all rectangles are given with coordinates, we are also interested in *combinatorial rectangle arrangements*, i.e., equivalence classes of combinatorially equivalent arrangements. Our contribution is two-fold.

First we consider maximal (with a maximal number of contacts) combinatorial rectangle contact arrangements, in which no three rectangles share a point. For rectangle arrangements this is equivalent to the contact graph being *triangle-free*, unlike, e.g., for triangle contact arrangements. We prove a series of analogues to the well-known maximal combinatorial triangle contact arrangements and to Schnyder realizers. The contact graph G of a maximal triangle contact arrangement is a maximal planar graph. A 3-*orientation* is an orientation of the edges

E. Di Giacomo and A. Lubiw (Eds.): GD 2015, LNCS 9411, pp. 231–244, 2015.
DOI: 10.1007/978-3-319-27261-0_20

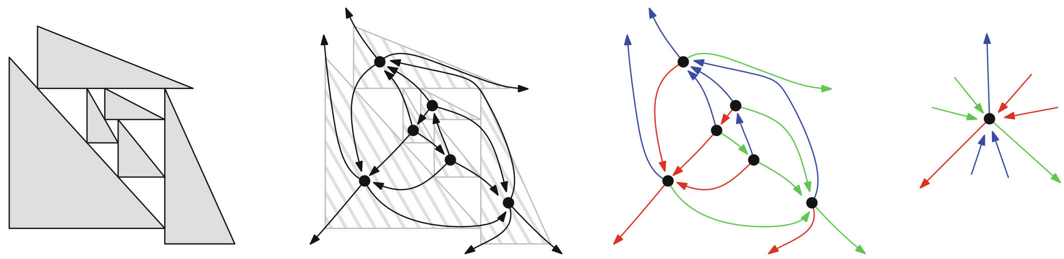

Fig. 1. Left to right: maximal combinatorial contact arrangement with axis-aligned triangles, no three sharing a point. 3-orientation of G'. Schnyder realizer of G'. Local coloring rules for Schnyder realizer (Color figure online).

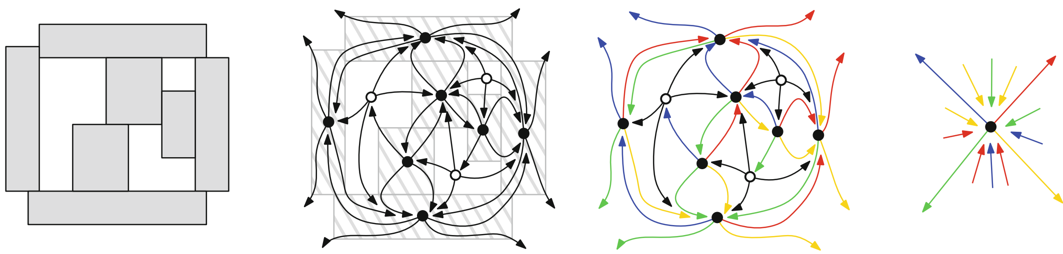

Fig. 2. Left to right: maximal combinatorial contact arrangement with axis-aligned rectangles, no three sharing a point. 4-orientation of underlying graph. Corner-edge-labeling of underlying graph. Local coloring rules for corner-edge-labeling (Color figure online).

of a graph G', obtained from G by adding six edges (two at each outer vertex), in which every vertex has exactly three outgoing edges. Each outer vertex has two outgoing edges that end in the outer face without having an endpoint there. A *Schnyder realizer* [10,11] is a 3-orientation of G' together with a coloring of its edges with colors $0, 1, 2$ such that every vertex has exactly one outgoing edge in each color and incoming edges are colored in the color of the "opposite" outgoing edge. The three outgoing edges represent the three corners of a triangle and the color specifies the corner, see Fig. 1. De Fraysseix *et al.* [3] proved that the maximal combinatorial triangle contact arrangements of G are in bijection with the 3-orientations of G' and the Schnyder realizers of G'. Schnyder proved that for every maximal planar graph G, G' admits a Schnyder realizer and hence G is a triangle contact graph.

In this paper we prove an analogous result, which, roughly speaking, is the following. We consider maximal triangle-free combinatorial rectangle contact arrangements. The corresponding contact graph G is planar with all faces of length 4 or 5. We define an underlying plane multigraph $\bar{G}$, whose vertex set also includes a vertex for each inner face of the contact graph, and define 4-*orientations* of $\bar{G}$. Here, every vertex has exactly four outgoing edges, where each outer vertex has two edges ending in the outer face. For a 4-orientation we introduce *corner-edge-labelings* of $\bar{G}$, which are, similar to Schnyder realizers, colorings of the outgoing edges at vertices of $\bar{G}$ corresponding to rectangles with colors $0, 1, 2, 3$ satisfying certain local rules. Each outgoing edge represents a corner of a rectangle and the color specifies which corner it is, see Fig. 2. We then prove that the combinatorial contact arrangements of G are in bijection with the 4-orientations of $\bar{G}$ and the corner-edge-labelings of $\bar{G}$.

Thomassen [12] proved that rectangle contact graphs are precisely the graphs admitting a planar embedding in which no triangle contains a vertex in its interior. We also prove here that for every maximal triangle-free planar graph G, $\bar{G}$ admits a 4-orientation, obtaining a new proof that G is a rectangle contact graph.

Our second result is concerned with the question whether a given geometric rectangle arrangement can be transformed into a combinatorially equivalent square arrangement. The similar question whether a pseudocircle arrangement can be transformed into a combinatorially equivalent circle arrangement has recently been studied by Kang and Müller [6], who showed that the problem is NP-hard. We say that a rectangle arrangement can be *squared* (or is *squarable*) if an equivalent square arrangement exists. Obviously, squares are a very restricted class of rectangles and not every rectangle arrangement can be squared. The natural open question is to characterize the squarable rectangle arrangements and to answer the complexity status of the corresponding decision problem. As a first step towards solving these questions, we show, on the one hand, some general necessary conditions and, on the other hand, sufficient conditions implying that certain subclasses of rectangle arrangements are always squarable.

Related Work. Intersection graphs and contact graphs of axis-aligned rectangles or squares in the plane are a popular, almost classic, topic in discrete mathematics and theoretical computer science with lots of applications in computational geometry, graph drawing and VLSI chip design. Most of the research for rectangle intersection graphs concerns their recognition [14], colorability [1] or the design of efficient algorithms such as for finding maximum cliques [5]. On the other hand, rectangle contact graphs are mainly investigated for their combinatorial and structural properties. Almost all the research here concerns edge-maximal 3-connected rectangle contact graphs, so called *rectangular duals*. These can be characterized by the absence of separating triangles [9,13] and the corresponding representations by touching rectangles can be seen as dissections of a rectangle into rectangles. Combinatorially equivalent dissections are in bijection with regular edge labelings [7] and transversal structures [4]. The question whether a rectangular dual has a rectangle dissection in which all rectangles are squares has been investigated by Felsner [2].

2 Preliminaries

In this paper a *rectangle* is an axis-aligned rectangle in the plane, i.e., the cross product $[x_1, x_2] \times [y_1, y_2]$ of two bounded closed intervals. A *geometric rectangle arrangement* is a finite set $\mathcal{R}$ of rectangles; it is a *contact arrangement* if any two rectangles have disjoint interiors. In a contact arrangement, any two non-disjoint rectangles R_1, R_2 have one of the two contact types *side contact* and *corner contact*, see Fig. 3 (left); we exclude the degenerate case of two rectangles

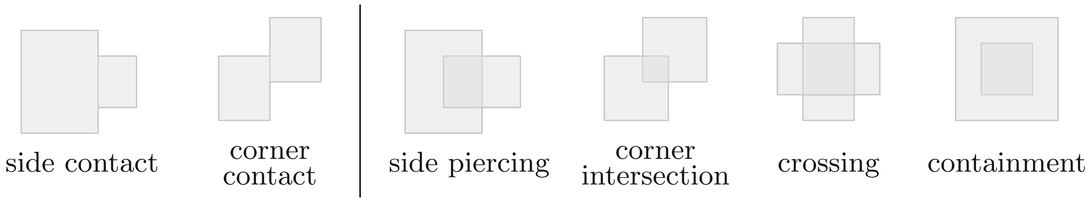

Fig. 3. Contact types (left) and intersection types (right) of rectangles.

sharing only one point. If $\mathcal{R}$ is not a contact arrangement, four intersection types are possible: *side piercing*, *corner intersection*, *crossing*, and *containment*, see Fig. 3 (right). Note that side contact and corner contact are degenerate cases of side piercing and corner intersection, whereas crossing and containment have no analogues in contact arrangements. If no two rectangles form a crossing, we say that $\mathcal{R}$ is *cross-free*. Moreover, in each type (except containment) it is further distinguished which sides of the rectangles touch or intersect.

Two rectangle arrangements $\mathcal{R}_1$ and $\mathcal{R}_2$ are *combinatorially equivalent* if $\mathcal{R}_1$ can be continuously deformed into $\mathcal{R}_2$ such that every intermediate state is a rectangle arrangement with the same intersection or contact type for every pair of rectangles. An equivalence class of combinatorially equivalent arrangements is called a *combinatorial rectangle arrangement*. So while a geometric arrangement specifies the coordinates of all rectangles, think of a combinatorial arrangement as specifying only the way in which any two rectangles touch or intersect. In particular, a combinatorial rectangle arrangement is defined by **(1)** for each rectangle R and each side of R the counterclockwise order of all intersecting (touching) rectangle edges, labeled by their rectangle R' and the respective side of R' (top, bottom, left, right), **(2)** for containments the respective component of the arrangement, in which a rectangle is contained.

In the *intersection graph* of a rectangle arrangement there is one vertex for each rectangle and two vertices are adjacent if and only if the corresponding rectangles intersect. As combinatorially equivalent arrangements have the same intersection graph, combinatorial arrangements themselves have a well-defined intersection graph. For rectangle contact arrangements (combinatorial or geometric) the intersection graph is also called the *contact graph*. Note that such contact graphs are planar, as we excluded the case of four rectangles meeting in a corner.

3 Statement of Results

3.1 Maximal Triangle-Free Planar Graphs and Rectangle Contact Arrangements

We consider so-called *MTP-graphs*, that is, (M)aximal (T)riangle-free (P)lane graphs with a quadrangular outer face. Note that each face in such an MTP-graph is a 4-cycle or 5-cycle, and that every plane triangle-free graph is an induced subgraph of some MTP-graph. Given an MTP-graph G a rectangle contact arrangement of G is one whose contact graph is G, where the embedding

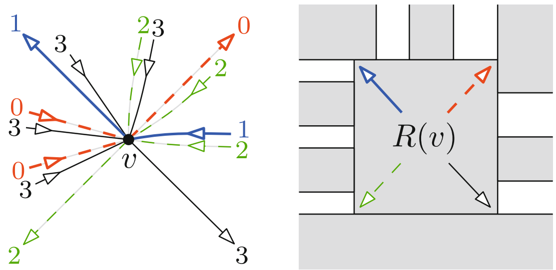

Fig. 4. Local color patterns in corner-edge-labelings of an MTP-graph at a vertex v, together with the corresponding part in a rectangle contact arrangement (Color figure online).

inherited from the arrangement is the given embedding of G, and where each outer rectangle has two corners in the unbounded region[1]. We define the closure, 4-orientations and corner-edge-labelings:

The *closure* $\bar{G}$ of G is derived from G by replacing each edge of G with a pair of parallel edges, called an *edge pair*, and adding into each inner face f of G a new vertex, also denoted by f, connected by an edge, called a *loose edge*, to each vertex incident to that face. At each outer vertex we add two loose edges pointing into the outer face, although we do not add a vertex for the outer face. Note that $\bar{G}$ inherits a unique plane embedding with each inner face being a triangle or a 2-gon.

A 4-orientation of $\bar{G}$ is an orientation of the edges and half-edges of $\bar{G}$ such that every vertex has outdegree exactly 4. An edge pair is called *uni-directed* if it is oriented consistently and *bi-directed* otherwise.

A corner-edge-labeling of $\bar{G}$ is a 4-orientation of $\bar{G}$ together with a coloring of the outgoing edges of $\bar{G}$ at each vertex of G with colors $0, 1, 2, 3$ (see Fig. 4) such that

(i) around each vertex v of G we have outgoing edges in color $0, 1, 2, 3$ in this counterclockwise order and

(ii) in the wedge, called *incoming wedge*, at v counterclockwise between the outgoing edges of color i and $i+1$ there are some (possibly none) incoming edges colored $i+2$ or $i+3$, $i = 0, 1, 2, 3$, all indices modulo 4.

In a corner-edge-labeling the four outgoing edges at a vertex of $\bar{G}$ corresponding to a face of G are not colored. Further we remark that (i) implies that uni-directed pairs are colored i and $i-1$, while (ii) implies that bi-directed pairs are colored i and $i+2$, for some $i \in \{0, 1, 2, 3\}$, where all indices are considered modulo 4. The following theorem is proved in Sect. 4.

Theorem 1. *Let G be an MTP-graph, then each of the following are in bijection:*

- *the combinatorial rectangle contact arrangements of G*
- *the corner-edge-labelings of $\bar{G}$*
- *the 4-orientations of $\bar{G}$.*

[1] Other configurations of the outer four rectangles can be easily derived from this.

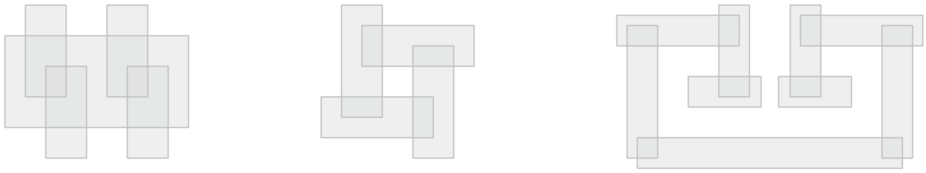

Fig. 5. Three cross-free unsquarable rectangle arrangements.

Using the bijection between 4-orientations of $\bar{G}$ and combinatorial rectangle contact arrangements of G given in Theorem 1, we can give a new proof that every MTP-graph G is a rectangle contact graph, which is the statement of the next theorem; its proof is given in the full paper [8] and sketched in Sect. 5.

Theorem 2. *For every MTP-graph G, $\bar{G}$ has a 4-orientation and it can be computed in linear time. In particular, G has a rectangle contact arrangement.*

We remark that our technique in the proof of Theorem 1 constructs from a given 4-orientation of $\bar{G}$ in linear time a geometric rectangle contact arrangement of G in the $2n \times 2n$ square grid, where n is the number of vertices in G. Thus also the rectangle contact arrangement in Theorem 2 can be computed in linear time and uses only a linear-size grid.

3.2 Squarability and Line-Pierced Rectangle Arrangements

In the squarability problem, we are given a rectangle arrangement $\mathcal{R}$ and want to decide whether $\mathcal{R}$ can be squared. The first observation is that there are obvious obstructions to the squarability of a rectangle arrangement. If any two rectangles in $\mathcal{R}$ are crossing (see Fig. 3) then there are obviously no two combinatorially equivalent squares.

But even if we restrict ourselves to cross-free rectangle arrangements, we can find unsquarable configurations. One such arrangement is depicted in Fig. 5 (left). To get an unsquarable arrangement with a triangle-free intersection graph, we can use the fact that two side-piercing rectangles translate immediately into a smaller-than relation for the corresponding squares: the side length of the square to pierce into the side of another square needs to be strictly smaller. Hence any rectangle arrangement that contains a cycle of side-piercing rectangles cannot be squarable, see Fig. 5 (middle). Moreover, we may even create a counterexample of a rectangle arrangement whose intersection graph is a path and that causes a geometrically infeasible configuration for squares, see Fig. 5 (right).

Proposition 1. *Some cross-free rectangle arrangements are unsquarable, even if the intersection graph is a path.*

Therefore we focus on a non-trivial subclass of rectangle arrangements that we call line-pierced. A rectangle arrangement $\mathcal{R}$ is *line-pierced* if there exists a horizontal line ℓ such that $\ell \cap R \neq \emptyset$ for all $R \in \mathcal{R}$. The line-piercing strongly restricts the possible vertical positions of the rectangles in $\mathcal{R}$, which lets us prove two sufficient conditions for squarability in the following theorem.

Theorem 3. *Let $\mathcal{R}$ be a cross-free, line-pierced rectangle arrangement.*

- *If $\mathcal{R}$ is triangle-free, then $\mathcal{R}$ is squarable.*
- *If $\mathcal{R}$ has only corner intersections, then $\mathcal{R}$ is squarable, even using line-pierced unit squares.*

On the other hand, cross-free, line-pierced rectangle arrangements in general may have forbidden cycles or other geometric obstructions to squarability. We give two examples in Sect. 6, together with a sketch of the proof of Theorem 3.

4 Bijections Between 4-Orientations, Corner-Edge-Labelings and Rectangle Contact Arrangements – Proof of Theorem 1

Throughout this section let $G = (V, E)$ be a fixed MTP-graph and $\bar{G}$ be its closure. By definition, every corner-edge-labeling of $\bar{G}$ induces a 4-orientation of $\bar{G}$. We prove Theorem 1, i.e., that combinatorial rectangle contact arrangements of G, 4-orientations of $\bar{G}$ and corner-edge-labelings of $\bar{G}$ are in bijection, in three steps:

- Every rectangle contact arrangement of G induces a 4-orientation of $\bar{G}$. (Lemma 1)
- Every 4-orientation of $\bar{G}$ induces a corner-edge-labeling of $\bar{G}$. (Lemma 3)
- Every corner-edge-labeling of $\bar{G}$ induces a rectangle contact arrangement of G. (Lemma 4)

Omitted proofs are provided in the full version of this paper [8].

4.1 From Rectangle Arrangements to 4-Orientations

Lemma 1. *Every rectangle contact arrangement of G induces a 4-orientation of $\bar{G}$.*

The proof idea is already given in Fig. 2: For every rectangle draw an outgoing edge through each of the four corners and for every inner face draw an outgoing edge through each of the four extremal sides.

We continue with a crucial property of 4-orientations. For a simple cycle C of G, consider the corresponding cycle $\bar{C}$ of edge pairs in $\bar{G}$. The *interior* of $\bar{C}$ is the bounded component of $\mathbb{R}^2$ incident to all vertices in C after the removal of all vertices and edges of $\bar{C}$. In a fixed 4-orientation of $\bar{G}$ a directed edge $e = (u, v)$ *points inside* C if $u \in V(C)$ and e lies in the interior of $\bar{C}$, i.e., either v lies in the interior of C, or e is a chord of $\bar{C}$ in the interior of $\bar{C}$.

Lemma 2. *For every 4-orientation of $\bar{G}$ and every simple cycle C of G the number of edges pointing inside C is exactly $|V(C)| - 4$.*

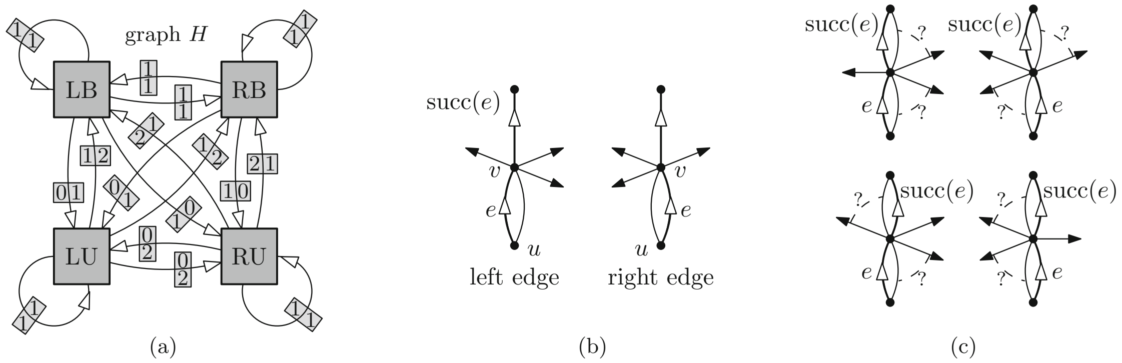

Fig. 6. (a) The graph H. **L**, **R**, **U**, **B** stands for left edge, right edge, uni-directed and bi-directed edge pair, respectively. The number of outgoing edges in the left and right wedge are shown on the left and right of the corresponding arrow. (b) Illustration of the definition of succ(e). (c) Summarizing the 16 possible cases for e and succ(e). Edges connected by a dashed arc may or may not coincide.

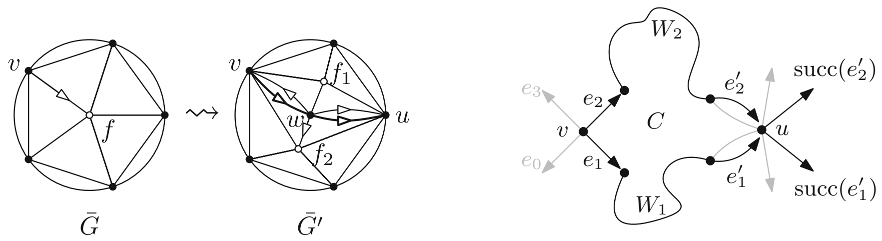

Fig. 7. Left: Stacking a new vertex w into a 5-face f of G. The orientation of edges on the boundary of f, as well as outgoing edges at f, f_1, f_2 is omitted. The directed edge (v, w) and its successor (w, u) are highlighted. Right: Illustration of the proof of the Claim in the proof of Lemma 3.

4.2 From 4-Orientations to Corner-Edge-Labelings

Next we shall show how a 4-orientation of $\bar{G}$ can be augmented (by choosing colors for the edges) into a corner-edge-labeling. Fix a 4-orientation. If e is a directed edge in an edge pair, then e is called a *left edge*, respectively *right edge*, when the 2-gon enclosed by the edge pair lies on the right, respectively on the left, when going along e in its direction. Thus, a uni-directed edge pair consists of one left edge and one right edge, while a bi-directed edge pair either consists of two left edges (clockwise oriented 2-gon) or two right edges (counterclockwise oriented 2-gon).

If $e = (u, v)$ is an edge in an edge pair, let e_2 and e_3 be the second and third outgoing edge at v when going counterclockwise around v starting with e. We define the *successor* of e as $\text{succ}(e) = e_2$ if e is a right edge, and $\text{succ}(e) = e_3$ if e is a left edge, see Fig. 6 (b,c). Note that in a corner-edge-labeling succ(e) is exactly the outgoing edge at v that has the same color as e, see Fig. 4.

Note that $e' = \text{succ}(e)$ may be a loose edge in $\bar{G}$ at the concave vertex for some 5-face in G. For the sake of shorter proofs below, we shall avoid the treatment of this case. To do so, we augment G to a supergraph G' such that starting with any edge in any edge pair and repeatedly taking the successor, we never run into a loose edge pointing to an inner face.

The graph G' is formally obtained from G by stacking a new vertex w into each 5-face f, with an edge to the incoming neighbor v of f in $\bar{G}$ and the vertex u at f that comes second after v in the clockwise order around f in $\bar{G}$. (Indeed, the second vertex in counterclockwise order would be equally good for our purposes.) Let f_1 and f_2 be the resulting 4-face and 5-face incident to w, respectively. We obtain a 4-orientation of the closure $\bar{G}'$ of G' by orienting all edges at f_1 as outgoing, both edges between v and w as right edges (counterclockwise), the remaining three edges at w as outgoing, and the remaining four edges at f_2 as outgoing. See Fig. 7 (left) for an illustration.

Before we augment the 4-orientation of $\bar{G}'$ into a corner-edge-labeling, we need one last observation. Let e and $\text{succ}(e)$ be two edges in edge pairs of $\bar{G}'$ with common vertex v. Consider the wedges at v between e and $\text{succ}(e)$ when going clockwise (left wedge) and counterclockwise (right wedge) around v. Each of e, $\text{succ}(e)$ can be a left edge or right edge, and in a uni-directed pair or a bi-directed pair. This gives us four types of edges and 16 possibilities for the types of e and $\text{succ}(e)$. The graph H in Fig. 6(a) shows for each of these 16 possibilities the number of outgoing edges at v in the left and right wedge at v.

Observation 4. *For every directed closed walk on k edges in the graph H in Fig. 6(a) we have*

$$\#\text{edges in left wedges} = \#\text{edges in right wedges} = k.$$

Proof. It suffices to check each directed cycle on k edges, $k = 1, 2, 3, 4$. □

Lemma 3. *Every 4-orientation of $\bar{G}$ induces a corner-edge-labeling of $\bar{G}$.*

A detailed proof of Lemma 3 is given in the full version of this paper [8].

Proof (Sketch). Consider the augmented graph G', its closure $\bar{G}'$ and 4-orientation as defined above. For any edge e in an edge pair in $\bar{G}'$ (and hence every edge of $\bar{G}$ outgoing at some vertex of G) consider the directed walk W_e in $\bar{G}'$ starting with e by repeatedly taking the successor as long as it exists (namely the current edge is in an edge pair).

First we show that W_e is a simple path ending at one of the eight loose edges in the outer face. Indeed, otherwise W_e would contain a simple cycle C where every edge on C, except the first, is the successor of its preceding edge on C. From the graph H of Fig. 6(a) we see that every wedge of C contains at most two outgoing edges. With Observation 4 the number of edges pointing inside C is at least $|V(C)| - 2$ and at most $|V(C)| + 2$, which is a contradiction to Lemma 2.

Now let v_0, v_1, v_2, v_3 be the outer vertices in this counterclockwise order. Define the color of e to be i if W_e ends with the right loose edge at v_i or the left loose edge at v_{i-1}, indices modulo 4. By definition every edge has the same color as its successor in $\bar{G}'$ (if it exists). Thus this coloring is a corner-edge-labeling of $\bar{G}'$ if at every vertex v of G the four outgoing edges are colored 0, 1, 2, and 3, in this counterclockwise order around v.

Claim. Let e_1, e_2 be two outgoing edges at v for which $W_{e_1} \cap W_{e_2}$ consists of more than just v. Then e_1 and e_2 appear consecutively among the outgoing edges around v, say e_1 clockwise after e_2.

Moreover, if $u \neq v$ is a vertex in $W_{e_1} \cap W_{e_2}$ for which the subpaths W_1 of W_{e_1} and W_2 of W_{e_2} between v and u do not share inner vertices, then the last edge e_1' of W_1 is a right edge and the last edge e_2' of W_2 is a left edge, e_1' and e_2' are part of (possibly the same) uni-directed pairs and these pairs sit in the same incoming wedge at u.

To prove this claim, we consider the cycle $C = W_1 \cup W_2$, count the edges pointing inside with the graph H and conclude that neither u nor v may have edges pointing inside C. See Fig. 7 (right) for an illustration.

The claim implies that the two walks W_{e_1} and W_{e_2} can neither cross, nor have an edge in common. Considering the four walks starting in a given vertex, we can argue (with the second part of the claim) that our coloring is a corner-edge-labeling of $\bar{G}'$. Finally, we inherit a corner-edge-labeling of $\bar{G}$ by reverting the stacking of artificial vertices in 5-faces. □

4.3 From Corner-Edge-Labelings to Rectangle Contact Arrangements

It remains to compute a rectangle arrangement of G based on a given corner-edge-labeling of $\bar{G}$. That is, we shall prove the following lemma.

Lemma 4. *Every corner-edge-labeling of* $\bar{G}$ *induces a rectangle contact arrangement of* G.

A detailed proof of Lemma 4 is given in the full version of this paper [8].

Proof (Sketch). Fix a corner-edge-labeling of $\bar{G}$. For every vertex v of G we introduce two pairs of variables $x_1(v), x_2(v)$ and $y_1(v), y_2(v)$ and set up a system of inequalities and equalities such that any solution defines a rectangle contact arrangement $\{R(v) \mid v \in V\}$ of G with $R(v) = [x_1(v), x_2(v)] \times [y_1(v), y_2(v)]$, which is compatible with the given corner-edge-labeling.

For every edge vw of G the way in which $R(v)$ and $R(w)$ are supposed to touch is encoded in the given corner-edge-labeling and this can be described by the inequalities and equalities in Table 1. Here we list the constraint and the conditions (color and orientation) of a single directed edge between v and w or a uni-directed edge pair outgoing at v and incoming at w in $\bar{G}$ under which we have this constraint.

Instead of showing that the system in Table 1 has a solution, we define another set of constraints implying all constraints in Table 1, for which it is easier to prove feasibility.

It suffices to define a system $\mathcal{I}_x$ for x-coordinates and treat the y-coordinates analogously. In $\mathcal{I}_x$ we have $x_1(v) < x_2(v)$ for every vertex v together with all equalities in the left of Table 1, but only those inequalities in the left of Table 1 that arise from edges in bi-directed edge pairs. The inequalities arising from uni-directed edge pairs are implied by the following set of inequalities. For a vertex

Table 1. Constraints encoding the type of contact between $R(v)$ and $R(w)$, defined based on the orientation and color(s) of the edge pair between v and w in $\bar{G}$.

constraint	edge	color	out
$x_1(w) < x_1(v) < x_2(w)$	right	2	v
	left	1	v
$x_1(w) < x_2(v) < x_2(w)$	right	0	v
	left	3	v
$x_1(w) = x_2(v)$	right	1	w
	left	2	w
	uni	0, 3	v

constraint	edge	color	out
$y_1(w) < y_1(v) < y_2(w)$	right	3	v
	left	2	v
$y_1(w) < y_2(v) < y_2(w)$	right	1	v
	left	0	v
$y_1(w) = y_2(v)$	right	2	w
	left	3	w
	uni	1, 0	v

v in G let $S_1(v) = a_1, \ldots, a_k$ and $S_2(v) = b_1, \ldots, b_\ell$ be the counterclockwise sequences of neighbors of v in the incoming wedges at v bounded by its outgoing edges of color 0 and 1, and color 2 and 3, respectively. See the left of Fig. 8. Then we have in $\mathcal{I}_x$ the inequalities

$$x_1(a_i) > x_2(a_{i+1}) \text{ for } i = 1, \ldots, k-1 \text{ and } x_2(b_i) < x_1(b_{i+1}) \text{ for } i = 1, \ldots, \ell - 1. \quad (1)$$

If $k = 1$ we have no constraint for $S_1(v)$ and if $\ell = 1$ we have no constraint for $S_2(v)$.

We associate the system $\mathcal{I}_x$ with a partially oriented graph I_x whose vertex set is $\{x_1(v), x_2(v) \mid v \in V\}$. For each inequality $a > b$ we have an oriented edge (a, b) in I_x, while for each equality $a = b$ we have an undirected edge ab in I_x, see Fig. 8.

We observe that I_x is planar and prove that I_x has no cycle C in which all directed edges are oriented consistently, which clearly implies that $\mathcal{I}_x$ has a solution. This is done by showing that no inner face is such a cycle, and that for every inner vertex u, vertex $x_1(u)$ has an incident undirected edge or incident outgoing edge and vertex $x_2(u)$ has an incident undirected edge or incident incoming edge. □

5 MTP Graphs Are Rectangle Contact Graphs – Proofsketch of Theorem 2

Theorem 2 is formally proven in the full version of this paper [8]. The idea is to prove by induction on the number of vertices that for an MTP-graph G we find a 4-orientation of $\bar{G}$. In the inductive step we either have (Case 1) that G has an inner 4-face, or (Case 2) that one can contract an inner edge e, keeping it an MTP-graph. Figures 9 and 10 illustrate how to find a 4-orientation in Cases 1 and 2, respectively.

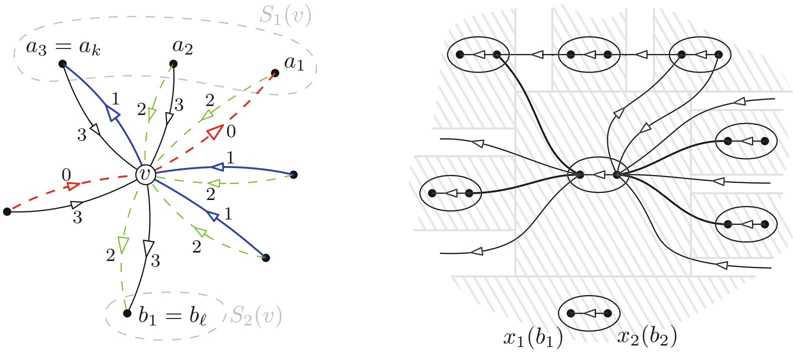

Fig. 8. Illustrating the definition of I_x around a vertex v. On the right a hypothetical rectangle contact arrangement is indicated (Color figure online).

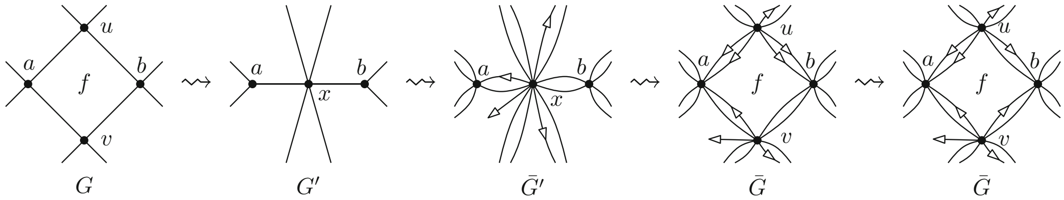

Fig. 9. Collapsing an inner 4-face and inheriting a 4-orientation when uncollapsing.

6 Line-Pierced Rectangle Arrangements and Squarability – Proofsketch of Theorem 3

Recall that a rectangle arrangement $\mathcal{R}$ is line-pierced if there is a horizontal line ℓ that intersects every rectangle in $\mathcal{R}$. Note that by the line-piercing property of $\mathcal{R}$ the intersection graph remains the same if we project each rectangle $R = [a, b] \times [c, d] \in \mathcal{R}$ onto the interval $[a, b] \subseteq \mathbb{R}$. In particular, the intersection graph $G_{\mathcal{R}}$ of a line-pierced rectangle arrangement $\mathcal{R}$ is an *interval graph*, i.e., intersection graph of intervals on the real line.

Line-pierced rectangle arrangements, however, carry more information than one-dimensional interval graphs since the vertical positions of intersection points between rectangles do influence the combinatorial properties of the arrangement. We obtain two squarability results for line-pierced arrangements in Propositions 2 and 3, which yield Theorem 3.

Proposition 2. *Every line-pierced, triangle-free, and cross-free rectangle arrangement $\mathcal{R}$ is squarable.*

There are instances, however, that satisfy the conditions of Proposition 2 and thus have a squaring, but not a line-pierced one. An example is given in Fig. 12.

Proposition 3. *Every line-pierced rectangle arrangement $\mathcal{R}$ restricted to corner intersections is squarable. There even exists a corresponding squaring with unit squares that remains line-pierced.*

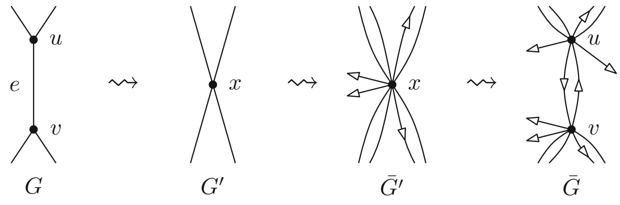

Fig. 10. Contracting an edge and keeping a 4-orientation when uncontracting.

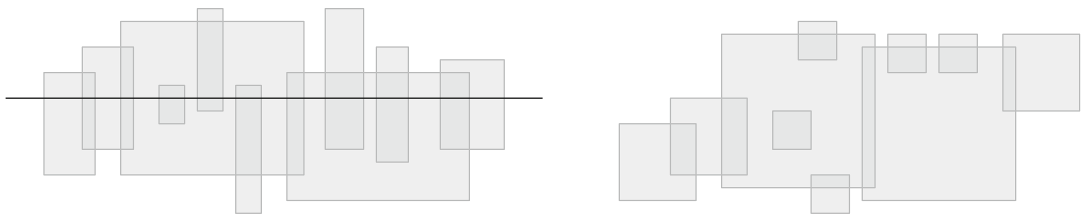

Fig. 11. Constructing a combinatorially equivalent squaring from a line-pierce, triangle-free, and cross-free rectangle arrangement.

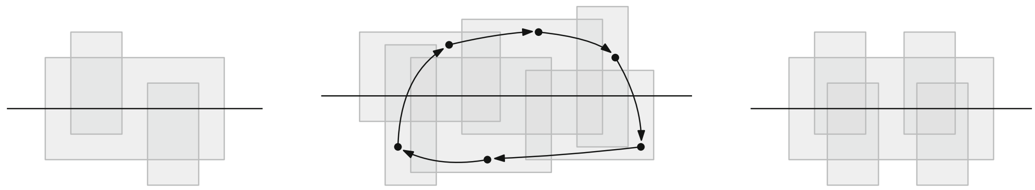

Fig. 12. Left: A line-pierced, triangle-free rectangle arrangement that has no line-pierced squaring. Middle: An unsquarable line-pierced rectangle arrangement due to a forbidden cycle of side-piercing intersections. Right: Squaring the two vertical pairs of rectangles on the right implies that the central square would need to be wider than tall.

Propositions 2 and 3 are proved in the full version of this paper [8]. The crucial observation is that the intersection graph of $\mathcal{R}$ is a caterpillar in the former case (Fig. 11) and a unit-interval graph in the latter case. The results can then be proven by induction on the number of vertices by iteratively removing the "rightmost" rectangle in the representation.

If we drop the restrictions to corner intersections and triangle-free arrangements, we can immediately find unsquarable instances, either by creating cyclic "'smaller than"' relations or by introducing intersection patterns that become geometrically infeasible for squares. Two examples are given in Fig. 12.

7 Conclusions

We have introduced corner-edge-labelings, a new combinatorial structure similar to Schnyder realizers, which captures the combinatorially equivalent maximal rectangle arrangements with no three rectangles sharing a point. Using this, we gave a new proof that every triangle-free planar graph is a rectangle contact graph. We also introduced the squarability problem, which asks for a given rectangle arrangement whether there is a combinatorially equivalent arrangement using only squares. We provide some forbidden configuration for the squarability of an arrangement and show that certain subclasses of line-pierced arrangements are always squarable. It remains open whether the decision problem for general arrangements is NP-complete.

Surprisingly, every unsquarable arrangement that we know has a crossing or a side-piercing. Hence we would like to ask whether every rectangle arrangement with only corner intersections is squarable. Another natural question is whether every triangle-free planar graph is a square contact graph.

References

1. Asplund, E., Grünbaum, B.: On a coloring problem. Mathematica Scandinavica **8**, 181–188 (1960)
2. Felsner, S.: Rectangle and square representations of planar graphs. In: Pach, J. (ed.) Thirty Essays in Geometric Graph Theory, pp. 213–248. Springer, New York (2012)
3. de Fraysseix, H., de Mendez, P.O., Rosenstiehl, P.: On triangle contact graphs. Comb. Probab. Comput. **3**, 233–246 (1994)
4. Fusy, E.: Transversal structures on triangulations: a combinatorial study and straight-line drawings. Discrete Math. **309**(7), 1870–1894 (2009)
5. Imai, H., Asano, T.: Finding the connected components and a maximum clique of an intersection graph of rectangles in the plane. J. Algorithms **4**(4), 310–323 (1983)
6. Kang, R.J., Müller, T.: Arrangements of pseudocircles and circles. Discrete Comput. Geom. **51**, 896–925 (2014)
7. Kant, G., He, X.: Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. Theor. Comput. Sci. **172**(1–2), 175–193 (1997)
8. Klawitter, J., Nöllenburg, M., Ueckerdt, T.: Combinatorial properties of triangle-free rectangle arrangements and the squarability problem. CoRR, arXiv:1509.00835, September 2015
9. Koźmiński, K., Kinnen, E.: Rectangular duals of planar graphs. Networks **15**, 145–157 (1985)
10. Schnyder, W.: Planar graphs and poset dimension. Order **5**(4), 323–343 (1989)
11. Schnyder, W.: Embedding planar graphs on the grid. In: 1st ACM-SIAM Symposium on Discrete Algorithms, SODA 1990, pp. 138–148 (1990)
12. Thomassen, C.: Interval representations of planar graphs. J. Comb. Theor. Ser. B **40**(1), 9–20 (1986)
13. Ungar, P.: On diagrams representing graphs. J. London Math. Soc. **28**, 336–342 (1953)
14. Yannakakis, M.: The complexity of the partial order dimension problem. SIAM J. Algebraic Discrete Methods **3**(3), 351–358 (1982)

Combinatorial and Geometric Properties of Planar Laman Graphs

Stephen Kobourov*
Department of Computer Science
University of Arizona
kobourov@cs.arizona.edu

Torsten Ueckerdt†
Department of Applied Mathematics
Charles University
torsten@kam.mff.cuni.cz

Kevin Verbeek‡
Department of Mathematics and Computer Science
TU Eindhoven
k.a.b.verbeek@tue.nl

Abstract

Laman graphs naturally arise in structural mechanics and rigidity theory. Specifically, they characterize minimally rigid planar bar-and-joint systems which are frequently needed in robotics, as well as in molecular chemistry and polymer physics. We introduce three new combinatorial structures for planar Laman graphs: angular structures, angle labelings, and edge labelings. The latter two structures are related to Schnyder realizers for maximally planar graphs. We prove that planar Laman graphs are exactly the class of graphs that have an angular structure that is a tree, called *angular tree*, and that every angular tree has a corresponding angle labeling and edge labeling.

Using a combination of these powerful combinatorial structures, we show that every planar Laman graph has an L-contact representation, that is, planar Laman graphs are contact graphs of axis-aligned L-shapes. Moreover, we show that planar Laman graphs and their subgraphs are the only graphs that can be represented this way.

We present efficient algorithms that compute, for every planar Laman graph G, an angular tree, angle labeling, edge labeling, and finally an L-contact representation of G. The overall running time is $\mathcal{O}(n^2)$, where n is the number of vertices of G, and the L-contact representation is realized on the $n \times n$ grid.

*Research supported by NSF grants CCF-0545743 and CCF-1115971.

†Research supported by GraDR EUROGIGA project No. GIG/11/E023.

‡Research supported by the Netherlands Organisation for Scientific Research (NWO) under project no. 639.022.707.

1 Introduction

A *contact graph* is a graph whose vertices are represented by geometric objects (like curves, line segments, or polygons), and edges correspond to two objects touching in some specified fashion. There is a large body of work about representing planar graphs as contact graphs. An early result is Koebe's 1936 theorem [18] that all planar graphs can be represented by touching disks.

In the late 1990's Schnyder showed that maximally planar graphs contain rich combinatorial structure [22]. With an angle labeling and a corresponding edge labeling, Schnyder shows that maximally planar graphs can be decomposed into three edge disjoint spanning trees. This combinatorial structure can be transformed into a geometric structure to produce a straight-line crossing-free planar drawing of the graph with vertex coordinates on the integer grid. Later, de Fraysseix *et al.* [10] show how to use the combinatorial structure to produce a representation of planar graphs as T-contact graphs (vertices are axis-aligned T's and edges correspond to point contact between T's) and triangle contact graphs.

We study the class of *planar Laman graphs* and show that we can find similarly powerful combinatorial structures. In particular, we show that every planar Laman graph G contains an *angular structure*—a graph on the vertices and faces of G with certain degree restrictions—that is also a tree and hence called an *angular tree*. We also show that every angular tree has a corresponding *angle labeling* and *edge labeling*, which can be thought of as a special Schnyder realizer [22]. Using a combination of these combinatorial structures we show that planar Laman graphs are *L-contact graphs*, graphs that can be represented as the contacts of axis-aligned non-degenerate L's (where the vertices correspond to the L's and the edges correspond

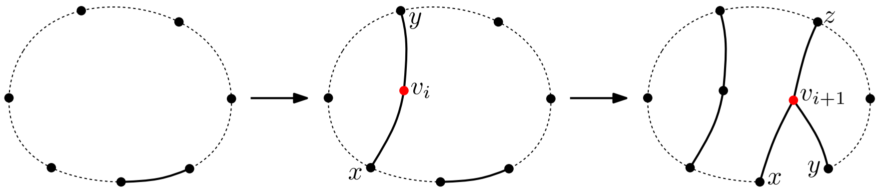

Figure 1: Two operations of a planar Henneberg construction: an $\mathbf{H_1}$-operation followed by an $\mathbf{H_2}$-operation.

to non-degenerate point contacts between the corresponding L's). As a by-product of our approach we obtain a new characterization of planar Laman graphs: a planar graph is a Laman graph if and only if it admits an angular tree. The L-contact representation can be computed in $\mathcal{O}(n^2)$ time and realized on the $n \times n$ grid, where n is the number of vertices of G.

Related Work. Koebe's theorem [18] is an early example of point-contact representation and shows that a planar graph can be represented by touching disks. Any planar graph also has a contact representation where all the vertices are represented by triangles in 2D [10], or even cubes in 3D [12].

Planar bipartite graphs can be represented by axis-aligned segment contacts [4, 9, 20]. Triangle-free planar graphs can be represented via contacts of segments with only three slopes [6]. Furthermore, every 4-connected 3-colorable planar graph and every 4-colored planar graph without an induced C_4 using four colors can be represented as the contact graph of segments [8]. More generally, planar Laman graphs can be represented with contacts of segments with arbitrary number of slopes and every contact graph of segments is a subgraph of a planar Laman graph [1].

The class of planar Laman graphs is of interest due to the fact that it contains several large classes of planar graphs (e.g., series-parallel graphs, outer-planar graphs, planar 2-trees). Laman graphs are also of interest in structural mechanics, robotics, chemistry and physics, due to their connection to rigidity theory, which dates back to the 1970's [19]. A system of fixed-length bars and flexible joints connecting them is minimally rigid if it becomes flexible once any bar is removed; planar Laman graphs correspond to rigid planar bar-and-joint systems [15, 16].

While Schnyder realizers were defined for maximally planar graphs [21, 22], the notion generalizes to 3-connected planar graphs [11]. Fusy's transversal structures [14] for irreducible triangulations of the 4-gon also provide combinatorial structure that can be used to obtain geometric results. Both concepts are closely related to certain angle labelings. Angle labelings of quadrangulations and plane Laman graphs have been considered before [13]. However, for planar Laman graphs the labeling does not have the desired Schnyder-like properties. In contrast, the labelings presented in this paper do have these properties.

Results and Organization. In Section 2 we introduce three combinatorial structures for planar Laman graphs. We first show that planar Laman graphs admit an angular tree. Next, we use this angular tree to obtain a corresponding angle labeling and edge labeling. In Section 3 we use a combination of these combinatorial structures to show that planar Laman graphs are L-contact graphs. We then describe an algorithm to compute the L-contact representation of a planar Laman graph G in $\mathcal{O}(n^2)$ time on the $n \times n$ grid. The running time of our algorithm is dominated by the computation of an angular tree of G. Given an angular tree, the algorithm runs in $\mathcal{O}(n)$ time. Proofs omitted due to space restrictions can be found in the full version of the paper [17].

2 Combinatorial Structures for Planar Laman Graphs

Let $G(W)$ be the subgraph of $G = (V, E)$ induced by $W \subseteq V$ and let $E(W)$ be the set of edges of $G(W)$.

Definition 2.1. *A Laman graph is a connected graph* $G = (V, E)$ *with* $|E| = 2|V| - 3$ *and* $|E(W)| \leq 2|W| - 3$ *for all* $W \subset V$.

Laman graphs admit a *Henneberg construction*: an ordering $v_1 \ldots v_n$ of the vertices such that, if G_i is the graph induced by $v_1 \ldots v_i$, then G_3 is a triangle and G_i is obtained from G_{i-1} by one of these operations:

($\mathbf{H_1}$) Choose two vertices x, y from G_{i-1} and add v_i together with the edges (v_i, x) and (v_i, y).

($\mathbf{H_2}$) Choose an edge (x, y) and a third vertex z from G_{i-1}, remove (x, y) and add v_i together with the three edges (v_i, x), (v_i, y), and (v_i, z).

Planar Laman graphs also admit a *planar Henneberg construction* [16]. That is, the graph can be constructed together with a plane straight-line embedding, with each vertex remaining in the position it is inserted. The

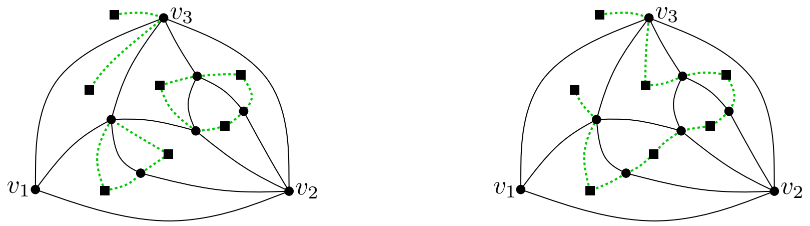

Figure 2: Two angular structures of the same plane Laman graph. The one on the right is an angular tree.

two operations of a (planar) Henneberg construction are illustrated in Figure 1.

Let G be a planar Laman graph. Since Laman graphs have $2|V| - 3$ edges, it easily follows that G contains a facial triangle. We choose an embedding of G in which such a triangle $\{v_1, v_2, v_3\}$ is the outer face. We can assume that the outer face remains intact during a Henneberg construction, i.e., we never perform an $\mathbf{H_2}$-operation on an edge on the outer face. Let v_1, v_2, v_3 appear in this counterclockwise order around the outer triangle. We call v_1, v_2 the *special vertices* and the outer edge $e^* = (v_1, v_2)$ the *special edge* of G.

2.1 Angular Structure The *angular graph* A_G of a plane graph G is a plane bipartite graph defined as follows. The vertices of A_G are the vertices $V(G)$ and faces $F(G)$ of G and there exists an edge (v, f) between $v \in V(G)$ and $f \in F(G)$ if and only if v is incident to f. If G is 2-connected, then A_G is a maximal bipartite planar graph and every face of A_G is a quadrangle.

DEFINITION 2.2. *An* angular structure *of a* 2*-connected plane graph G with special edge $e^* = (v_1, v_2)$ is a set T of edges of A_G with the following two properties:*

Vertex rule: *Every vertex $v \in V(G) \setminus \{v_1, v_2\}$ has exactly* 2 *incident edges in T. Special vertices have no incident edge in T.*

Face rule: *Every face $f \in F(G)$ has exactly* 2 *incident edges not in T.*

Let S be the set of edges of A_G that are not in T. The angular structure T can be represented by orienting the edges of A_G as follows. Every edge (v, f) is oriented from v to f if $(v, f) \in T$, and from f to v if $(v, f) \in S$. This way every vertex of A_G has exactly two outgoing edges (except for the special vertices). Such orientations of a maximal bipartite planar graph are called *2-orientations* and were introduced by de Fraysseix and Ossona de Mendez [7]. From a 2-orientation of A_G we can obtain an angular structure of G.

LEMMA 2.1. ([7]) *Every maximal bipartite planar graph has a 2-orientation. Thus every* 2*-connected plane graph has an angular structure.*

If G is a Laman graph, then $|F(G)| = |V(G)| - 1$ by Euler's formula. Thus every angular structure T consists of $2|V(G)| - 4$ edges and spans $|V(A_G)| - 2 = 2|V(G)| - 3$ vertices. Hence, if T is connected, then T is a spanning tree of $V(A_G) \setminus \{v_1, v_2\}$. An angular structure that is a tree is called an *angular tree*. In Fig. 2 two angular structures of the same plane Laman graph are shown – one is an angular tree.

Next we show that every plane Laman graph admits an angular tree. Our proof is constructive and computes an angular tree along a planar Henneberg sequence of G. Consider a cycle C in A_G such that the edges of C are alternatingly in S and T. We say that C is an *alternating cycle*. We can perform a *flip* on C by removing all edges in $C \cap T$ from T and adding all edges in $C \cap S$ to T. The resulting set of edges satisfies the properties of an angular structure. A flip corresponds to reversing the edges of a directed cycle in the corresponding 2-orientation.

LEMMA 2.2. *Let T consist of two connected components A and B, where A is a tree and B contains a cycle. If we perform a flip on an alternating 4-cycle C that contains an edge of A and an edge of the cycle in B, then the resulting angular structure is a tree.*

Proof. If we remove the edges in $C \cap T$ from T, then B becomes a tree, and we split up A into trees A_1 and A_2. The edges in $C \cap S$ connect A_1 to B and A_2 to B. Thus the resulting angular structure is a tree. ■

THEOREM 2.1. *Every plane Laman graph G admits an angular tree T and it can be computed in $\mathcal{O}(|V(G)|^2)$ time.*

Proof. We build G and T simultaneously along a planar Henneberg construction, which can be found in $\mathcal{O}(|V(G)|^2)$ time using an algorithm of Bereg [2]. T remains a tree during the construction. We begin with the triangle $\{v_1, v_2, v_3\}$ and T containing the two edges incident to v_3 in A_G. Now assume we insert a vertex v into a face f of G, which is split into two faces f_1 and f_2.

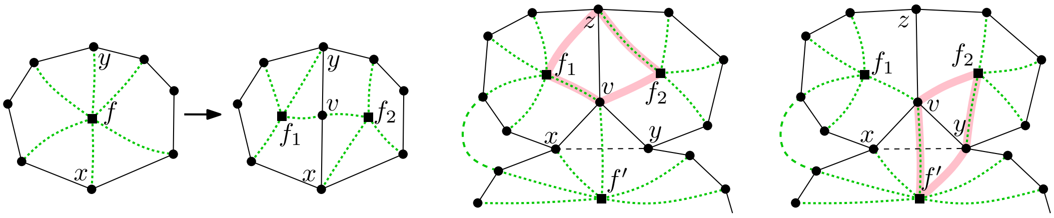

Figure 3: Left: updating T (drawn dotted) for ($\mathbf{H_1}$). Right: alternating cycles after ($\mathbf{H_2}$).

For an $\mathbf{H_1}$-operation, let x and y be the original vertices of the graph. We add an edge (u, f_1) to T if and only if u is incident to f_1 and $(u, f) \in T$ before the operation. We do the same for f_2. Furthermore, we add edges (v, f_1) and (v, f_2) to T. If $(x, f) \in T$ before the operation, then we remove either (x, f_1) or (x, f_2) from T. Similarly, if $(y, f) \in T$ before the operation, then we remove either (y, f_1) or (y, f_2) from T. By choosing these edges correctly, we can ensure that f_1 and f_2 satisfy the degree constraints; see Fig. 3(left). This operation cannot introduce a cycle, so T must remain a tree.

For an $\mathbf{H_2}$-operation, let (x, y) and z be the edge and vertex of the operation. Furthermore, let f' be the face of G that shares the edge (x, y) with f before the operation. We add an edge (u, f_1) to T if and only if u is incident to f_1 and $(u, f) \in T$ before the operation (same for f_2). Furthermore, we add edges (v, f') and either (v, f_1) or (v, f_2) to T. If $(z, f) \in T$ before the operation, then we remove either (z, f_1) or (z, f_2) from T. As above, we can choose the edges to ensure that f_1 and f_2 satisfy the degree constraints. However, this operation can introduce a cycle in T containing the new vertex v (if not, we are done). Assume w.l.o.g. that f_1 is part of this cycle, and hence $(v, f_1) \in T$; see Fig. 3(right).

If $(z, f_2) \in T$, then $(z, f_1) \notin T$, and the cycle formed by (z, f_1), (z, f_2), (v, f_2), and (v, f_1) is alternating and satisfies the requirements of Lemma 2.2. We can flip this cycle to turn T into a tree. If $(z, f_2) \notin T$, then $(y, f_2) \in T$ by the degree constraints on f_2. Also, $(y, f') \notin T$, for otherwise T would contain a cycle before the operation. Thus, the cycle formed by (y, f_2), (y, f'), (v, f'), and (v, f_2) is alternating and satisfies the requirements of Lemma 2.2. As before, we can flip this cycle to turn T into a tree.

At each step in the above procedure one vertex is added to G. The operations carried out to maintain the angular tree can be performed in $\mathcal{O}(1)$ time for an $\mathbf{H_1}$-operation and in $\mathcal{O}(|V(G)|)$ time for an $\mathbf{H_2}$-operation. Indeed, the bottleneck in the latter case is identifying the unique cycle in the intermediate angular structure. Thus the total runtime is $\mathcal{O}(|V(G)|^2)$, which concludes the proof. ∎

LEMMA 2.3. *If T is an angular tree and f is a triangular face of G, then T contains a perfect matching between non-special vertices of G and faces of G different from f.*

Proof. Remove the vertex corresponding to f (leaf in T) from T and let v be the non-special vertex with $(v, f) \in T$. Direct all edges of T towards v. Now every face $f' \neq f$ has exactly one outgoing edge in T and every non-special vertex has exactly one incoming edge in T. The desired matching can be obtained by matching each face different from f to the unique endpoint $v \in V(G)$ of its outgoing edge in T. ∎

2.2 Angle Labeling

Next we define a labeling of the angles of G, using the angular structure above; see Fig. 4. This labeling for 2-connected plane graphs is similar to the Schnyder angle labeling for maximally plane graphs.

DEFINITION 2.3. *An* angle labeling *of a 2-connected plane graph G with special edge $e^* = (v_1, v_2)$ is a labeling of the angles of G by $1, 2, 3, 4$, with the following two properties:*

Vertex rule: *Around every vertex $v \neq v_1, v_2$, in clockwise order, we get the following sequence of angles: exactly one angle labeled 3, zero or more angles labeled 2, exactly one angle labeled 4, zero or more angles labeled 1. All angles at v_1 are labeled 1, all angles at v_2 are labeled 2.*

Face rule: *Around every face, in clockwise order, we get the following sequence of angles: exactly one angle labeled 1, zero or more angles labeled 3, exactly one angle labeled 2, zero or more angles labeled 4.*

THEOREM 2.2. *Every 2-connected plane graph admits an angle labeling.*

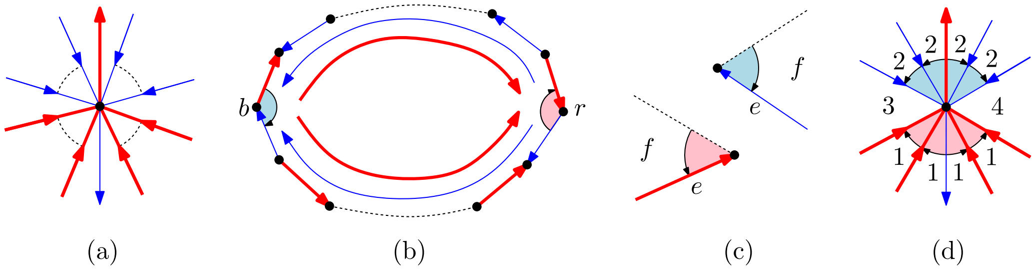

Figure 4: Vertex rule (a), face rule (b), and edge rule (c)-(d). Red edges are drawn thick.

Proof. By Lemma 2.1 every 2-connected plane graph G admits an angular structure, which corresponds to a 2-orientation of the angular graph A_G. The edges of a 2-orientation can be colored in red and blue, such that the edges around each vertex v are ordered as follows: one outgoing red edge, zero or more incoming red edges, one outgoing blue edge, zero or more incoming blue edges (the order is clockwise for $v \in V(G)$ and counterclockwise for $v \in F(G)$). Such an orientation and coloring of the edges of a maximal bipartite planar graph is called a *separating decomposition* [7].

We now label each angle at a vertex v of G based on the color and orientation of the corresponding edge (v, f) in the separating decomposition. If the edge is incoming at v and colored blue, we label the angle 1. If the edge is incoming at v and colored red, we label the angle 2. If the edge is outgoing at v and colored red, we label the angle 3. If the edge is outgoing at v and colored blue, we label the angle 4. It is now straightforward to verify that the vertex rule and face rule are implied by the order in which incident edges appear around each vertex in the separating decomposition. ■

Note that the correspondence derived above between an angular structure T and an angle labeling of G is such that $(v, f) \in T$ if and only if the corresponding angle label is 3 or 4. Moreover, from an angle labeling one can derive the corresponding separating decomposition of A_G and hence the corresponding angular structure. In particular, there is a bijection between angular structures of G and angle labelings of G.

2.3 Edge Labeling

Finally, we define an orientation and coloring of the edges of a 2-connected plane graph G based on an angular tree T of G; see Fig. 4. This edge labeling for 2-connected plane graphs is similar to the Schnyder edge labeling for maximally plane graphs.

DEFINITION 2.4. *An* edge labeling *of a 2-connected plane graph G with special edge $e^* = (v_1, v_2)$ is an orientation and coloring of the non-special edges of G with colors 1 (red) and 2 (blue), such that each of the following holds:*

Vertex rule: *Around every vertex $v \neq v_1, v_2$, in clockwise order, we get the following sequence of edges: exactly one outgoing red edge, zero or more incoming blue edges, zero or more incoming red edges, exactly one outgoing blue edge, zero or more incoming red edges, and zero or more incoming blue edges. All non-special edges at v_1 are incoming and red, all non-special edges at v_2 are incoming and blue.*

Face rule: *For every inner face f there are two distinguished vertices r and b. Every red edge on f is directed from b towards r, and every blue edge is directed from r towards b. The vertices r and b are called the red and blue sink of f, respectively.*

We denote the edge labeling by (E_r, E_b), where E_r and E_b is the set of all red and blue edges, respectively.

In an edge labeling (E_r, E_b) of G every non-special vertex has two outgoing edges. Together with the special edge this makes $2|V(G)| - 3$ edges in total. Thus $|E(G)| = 2|V(G)| - 3$ and $|F(G)| = |V(G)| - 1$. Every inner face has exactly two sinks, which makes $2|F(G)| = |E(G)| - 1$ in total. Indeed, there is a one-to-one correspondence between the non-special edges of G and sinks of inner faces in (E_r, E_b). We associate every directed edge e with the inner face f incident to it as illustrated in Fig. 4(d). This way we have the following for every edge labeling (E_r, E_b) of G.

Edge rule: Every non-special edge e corresponds to one incident inner face f, such that the endpoint of e is a sink of f in the color of e.

THEOREM 2.3. *If a 2-connected plane graph admits an angular tree, then it admits an edge labeling.*

Proof. Let G be a 2-connected plane graph and T be an angular structure of G. By Theorem 2.2, G admits an angle labeling that corresponds to T, i.e., the angle

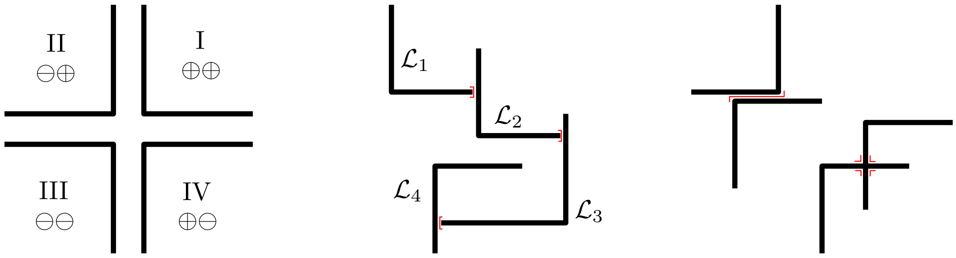

Figure 5: Left: possible L-shapes. Middle: valid contacts. Right: invalid contacts.

of a face f at a vertex v is labeled 3 or 4 if and only if $(v, f) \in T$. We split every vertex v in G, except for v_1 and v_2, into two vertices v^1 and v^2, in such a way that for $i = 1, 2$ all edges incident to an angle labeled i are incident to v^i. We call the resulting graph H. In other words H arises from G by splitting each non-special vertex along its two edges in T. Thus, as T is acyclic, H is connected. Since H consists of $2|V(G)| - 2$ vertices (v_1, v_2 plus $2(|V(G)| - 2)$ split vertices) and $|E(G)| = 2|V(G)| - 3$ edges, H is a tree.

We orient every edge e in H towards the special edge e^* of G. We color e red if it is outgoing at some v^2 and blue if it is outgoing at some v^1. It is now straightforward to check, using the vertex rule and face rule of the angle labeling, that this orientation and coloring of the edges is a valid edge labeling of G. ■

Not every edge labeling corresponds to an angular tree. Also, some but not all angular structures that are not trees correspond to an edge labeling. For example, the angular structure in Fig. 2 (left) has no corresponding edge labeling. Hence, edge labelings of G and angular structures (or angular trees) of G are *not* in bijection.

THEOREM 2.4. *An edge labeling (E_r, E_b) of a 2-connected plane graph G with special edge $e^* = (v_1, v_2)$ has the following two properties:*

(i) The graph $E_r \cup E_b^{-1}$ ($E_b \cup E_r^{-1}$) is acyclic, where E_b^{-1} is E_b with the direction of all edges reversed.

(ii) The graph E_r (E_b) is a spanning tree of $G \setminus \{v_2\}$ ($G \setminus \{v_1\}$) with all edges directed towards v_1 (v_2).

Proof. Consider the graph $E_r \cup E_b^{-1}$. Since every vertex except for v_1 and v_2 has an outgoing red edge and an outgoing blue edge, there is only one source (all edges are outgoing at v_2) and one sink (all edges are incoming at v_1) in $E_r \cup E_b^{-1}$. By the face rule, every face has exactly one source (the blue sink) and one sink (the red sink). The face rule for the inner face of G containing the special edge e^* implies that the outer cycle as well has exactly one source (v_1) and exactly one sink (v_2). Every nesting minimal (the set of faces it circumscribes is inclusion minimal) directed cycle in a plane graph is either a facial cycle or has a source or sink in its interior. This proves (i). Part (ii) follows directly from part (i) and the fact that every non-special vertex has one outgoing edge in E_r (E_b). ■

3 L-Contact Graphs

An *L-shape* $\mathcal{L}$ is a path consisting of exactly one horizontal segment and exactly one vertical segment. There are four different types of L-shapes; see Fig. 5(left). Two L-shapes $\mathcal{L}_1$ and $\mathcal{L}_2$ make *contact* if and only if the endpoint of one of the two L-shapes coincides with an interior point of the other L-shape; see Fig. 5(middle). If the endpoint belongs to $\mathcal{L}_1$, then we say that $\mathcal{L}_1$ makes contact with $\mathcal{L}_2$. Note that we do not allow contact using the bend of an L-shape; see Fig. 5(right).

A graph $G = (V, E)$ is an *L-contact graph* if there exist non-crossing L-shapes $\mathcal{L}(v)$ for each $v \in V$, such that $\mathcal{L}(u)$ and $\mathcal{L}(v)$ make contact if and only if $(u, v) \in E$. We call these L-shapes the *L-contact representation* of G. We can match edges of L-contact graphs to endpoints of L-shapes. However, an endpoint that is bottommost, topmost, leftmost, or rightmost cannot correspond to an edge. We call an L-contact representation *maximal* if every endpoint that is neither bottommost, topmost, leftmost, nor rightmost makes a contact, and there are at most three endpoints that do not make a contact. We assume that the bottommost, topmost, leftmost, and rightmost endpoints are uniquely defined.

In a maximal L-contact representation of a graph G, each inner face of G is bounded by a simple rectilinear polygon, which is contained in the union of all L-shapes. Now each $\mathcal{L}(v)$ has a right angle, which is a convex corner of the polygon corresponding to one incident face at v and a concave corner corresponding to another incident face at v, provided the corresponding face is an inner face.

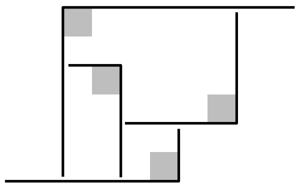

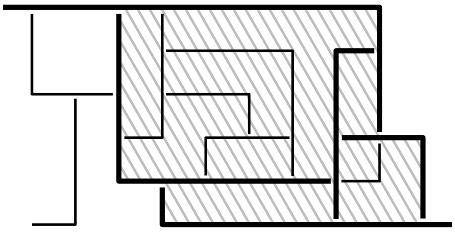

Figure 6: Left: K_4 is an L-contact graph but not a Laman graph. Right: Illustration of the proof of Lemma 3.1.

LEMMA 3.1. *If a graph G has a maximal L-contact representation in which each inner face contains the right angle of exactly one $\mathcal{L}$, then G is a plane Laman graph.*

Proof. Consider a maximal L-contact representation of G in which every inner face contains the right angle of exactly one $\mathcal{L}$. By the definition of maximal L-contact representations, we get that $|E(G)| \geq 2|V(G)| - 3$. We need to show that $|E(W)| \leq 2|W| - 3$ for all subsets $W \subseteq V(G)$ of at least two vertices. For the sake of contradiction, let W be a set ($|W| \geq 2$) with $|E(W)| \geq 2|W| - 2$. It follows that at most two endpoints of L-shapes corresponding to vertices in W do not make contact when restricted to W. Since this holds for one bottommost endpoint and one topmost endpoint, we have $|E(W)| = 2|W| - 2$. Moreover, if we choose W to be inclusion-minimal among all such sets, then $G(W)$ is 2-connected; see thick L-shapes in Fig. 6.

The outer face of $G(W)$ is bounded by a rectilinear polygon $\mathcal{P}$ with two additional ends sticking out. This polygon is highlighted in Fig. 6. Consider the vertex set $W' \supseteq W$ of all vertices whose corresponding L-shapes are contained in $\mathcal{P}$, i.e., $G(W)$ is a subgraph of $G(W')$ and every inner face of $G(W')$ is an inner face of G. Since the representation is maximal we have $|E(W')| = |E(W)| + 2|W' \setminus W| = 2|W'| - 2$. Hence $G(W')$ has too many edges as well. We want to show that one inner face of $G(W')$ has two convex angles, which would then complete the proof.

Let k be the number of outer vertices of $G(W')$. Since $\mathcal{P}$ has only two endpoints sticking out, all but two of its convex corners are due to a single $\mathcal{L}$, so the number of convex corners of $\mathcal{P}$ is at most $k + 2$. Each outer edge of $G(W')$, except for two, corresponds to a contact that is a concave corner of $\mathcal{P}$, so the number of concave corners of $\mathcal{P}$ is at least $k-2$. In every rectilinear polygon the number of concave corners is exactly the number of its convex corners minus four. Thus we conclude that both inequalities above must hold with equality. In particular, every concave corner of $\mathcal{P}$ corresponds to a contact of two L-shapes and no concave corner is due to a single $\mathcal{L}$. Moreover, every L-shape corresponding to an outer vertex in $G(W')$ forms a convex corner of $\mathcal{P}$. Hence for every $w \in W'$ the right angle of $\mathcal{L}(w)$ lies inside $\mathcal{P}$.

By Euler's formula $G(W')$ has precisely $|W'| - 1$ inner faces. Since there are $|W'|$ right angles among those inner faces, one inner face must have two right angles. ∎

DEFINITION 3.1. *A maximal L-contact representation is* proper *if every inner face contains the right angle of exactly one $\mathcal{L}$. An L-contact graph is* proper *if it has a proper L-contact representation.*

Lemma 3.1 states that all proper L-contact graphs are plane Laman graphs. The main result of the remainder of this section is the following.

THEOREM 3.1. *Plane Laman graphs are precisely proper L-contact graphs.*

To obtain an L-contact representation of a plane Laman graph, we require only the existence of an angular tree with the corresponding edge labeling. Thus, if a 2-connected plane graph G admits an angular tree, then it has a corresponding edge labeling by Theorem 2.3, and we can compute a proper L-contact representation of G. We obtain the following characterization of planar Laman graphs as a by-product of our approach.

THEOREM 3.2. *A planar 2-connected graph is a Laman graph if and only if it admits an angular tree.*

3.1 Vertex Types Assume we have an angular tree T with corresponding edge labeling (E_r, E_b) for a plane Laman graph G. Every non-special vertex v in G has two incident edges in T. The other endpoint of such an edge corresponds to a face in G. These are the two faces that contain the bend of $\mathcal{L}(v)$. The matching M of T obtained from Lemma 2.3 (using the outer face of G as the triangular face) determines for every vertex of G the incident inner face f containing the right angle of $\mathcal{L}(v)$. The outgoing red (blue) edge of a vertex v determines the contact made by the horizontal (vertical) leg of $\mathcal{L}(v)$.

We derive from M and (E_r, E_b) the type of the L-shape $\mathcal{L}(v)$ for every vertex v. The *red sign* and *blue*

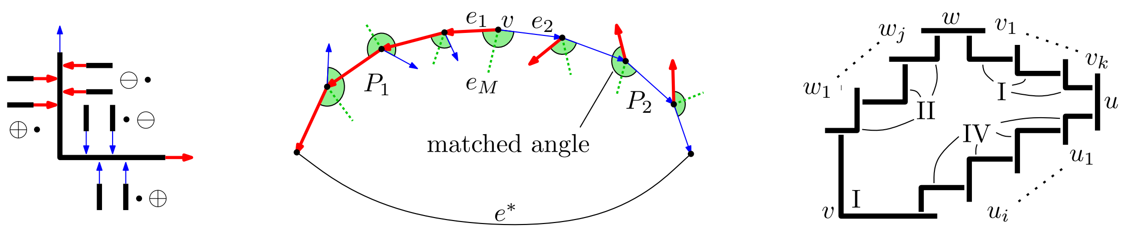

Figure 7: Left/Right: types around a vertex/face ($t(v) = \mathrm{I}$). Middle: proof Lemma 3.2.

sign of a vertex v, denoted by $t_r(v)$ and $t_b(v)$, represent the direction of the horizontal and vertical leg of $\mathcal{L}(v)$, respectively. We write the type of v as $t(v) = t_r(v)t_b(v)$, or as its quadrant number; see Fig. 5(left).

First we set $t_b(v_1) = \oplus$ and $t_r(v_2) = \oplus$ (the red sign of v_1 and the blue sign of v_2 are irrelevant). For every non-special vertex v, let $e_r(v)$ ($e_b(v)$) be its outgoing red (blue) edge, and $e_M(v)$ its incident edge in M. The angle between $e_r(v)$ and $e_b(v)$ that contains $e_M(v)$ is called the *matched angle*. The opposite angle is called the *unmatched angle* (v_1 and v_2 have only an unmatched angle). We use the following rule.

Type rule: Let $e = (u, v)$ be a directed edge from u to v of color c. If e lies in the unmatched angle of v, we set $t_c(u) = t_c(v)$, otherwise $t_c(u) \neq t_c(v)$.

We need to check if this type rule, along with T, M, and (E_r, E_b), results in a correct L-contact representation. Around every vertex v, the neighboring vertices with incoming edges to v must have the correct red or blue sign. For example, if $t(v) = \mathrm{I}$ and the edge $u \to v$ is blue and lies in the matched angle of v, then $t_b(u) = \ominus$. Note that this follows directly from the type rule; see Fig. 7(left).

Secondly, the convex angle of an L-shape $\mathcal{L}(v)$ must belong to the face that contains $e_M(v)$. For example, if $e_b(v), e_M(v), e_r(v)$ appear in clockwise order around v, then $t(v) = \mathrm{I}$ or $t(v) = \mathrm{III}$. We say v is *odd* if $e_b(v), e_M(v), e_r(v)$ appear in clockwise order around v, and *even* otherwise.

Lemma 3.2. *A non-special vertex v is odd if and only if $t_r(v) = t_b(v)$.*

Proof. Consider the directed red path P_1 from v to v_1 and the directed blue path P_2 from v to v_2; see Fig. 7(middle). Since $E_r \cup E_b^{-1}$ is acyclic by Theorem 2.4, $P_1 \cap P_2$ consists only of v. Let C be the cycle formed by P_1, P_2 and the special edge e^*, and G' be the maximal subgraph of G whose outer cycle is C. We define r_1, r_2, r_3, r_4 as follows (we define b_1, b_2, b_3, b_4 analogously w.r.t. P_2):

$r_1 := \#\{e = (u, v) \in P_1 \mid e$ in unmatched angle of v and $e_M(v)$ outside $G'\}$

$r_2 := \#\{e = (u, v) \in P_1 \mid e$ in unmatched angle of v and $e_M(v)$ inside $G'\}$

$r_3 := \#\{e = (u, v) \in P_1 \mid e$ in matched angle of v and $e_M(v)$ outside $G'\}$

$r_4 := \#\{e = (u, v) \in P_1 \mid e$ in matched angle of v and $e_M(v)$ inside $G'\}$

Now let $k = |C|$ be the number of vertices on C and $|V(G')| = k + n'$. Then G' has $2n' + k + r_2 + r_3 + b_2 + b_3$ edges and thus by Euler's formula $n' + b_2 + b_3 + r_2 + r_3 + 1$ inner faces. On the other hand $G' \setminus \{v\}$ contains exactly $n' + b_2 + b_4 + r_2 + r_4$ matching edges. So if v is odd, then e_M lies inside G', too. Since the number of inner faces and matching edges must coincide we have $b_3 + r_3 = b_4 + r_4$. In particular $b_3 + b_4$ and $r_3 + r_4$ have the same parity, which means that the red and blue sign of v coincide. If v is even, then e_M lies outside G' and we get $b_3 + r_3 + 1 = b_4 + r_4$, which implies that $b_3 + b_4$ and $r_3 + r_4$ have different parity. Hence the red sign and blue sign at v are distinct. ∎

Finally we consider the faces in the L-contact representation. Every inner face f of G has three special vertices: the two sinks u and w, as well as the vertex v that f is matched to in M. Let $u, u_1, \ldots, u_i, v, w_1, \ldots, w_j, w, v_1, \ldots, v_k$ be the clockwise order of the vertices around f. The type rule implies the following shape of faces in the L-contact representation; see Fig. 7(right).

Lemma 3.3. *Let v be the vertex that is matched to a face f, and t be the type of v. Then we have the following:*

- *Each of $u_1, \ldots, u_i$ has type $t - 1$.*
- *Each of $v_1, \ldots, v_k$ has type t.*
- *Each of $w_1, \ldots, w_j$ has type $t + 1$.*

	$t(v) = \oplus\oplus$	$t(v) = \ominus\oplus$	$t(v) = \ominus\ominus$	$t(v) = \oplus\ominus$
D_r	$w_j \to f; f \to v_1, u_i$	$v_k, w_1 \to f; f \to u_1$	$v_1, u_i \to f; f \to w_j$	$u_1 \to f; f \to w_1, v_k$
D_b	$u_1 \to f; f \to w_1, v_k$	$w_j \to f; f \to v_1, u_i$	$v_k, w_1 \to f; f \to u_1$	$v_1, u_i \to f; f \to w_j$

Table 1: The three inequality edges of a face f of G in D_r and D_b for each type of f.

3.2 Inequalities Given the type of every vertex v, it suffices to find the point $(x(v), y(v)) \in \mathbb{R}^2$ where the bend of $\mathcal{L}(v)$ is located. Additionally we define for each inner face f an auxiliary point $(x(f), y(f)) \in \mathbb{R}^2$, which in the L-contact representation of G will correspond to some point in the bounded region corresponding to f.

We use two directed (multi-)graphs D_r and D_b on the vertices and inner faces of G to describe inequalities for the x- and y-coordinates, respectively. For every inequality $x(u) < x(v)$ $(y(u) < y(v))$ there is an edge $u \to v$ in D_r (D_b), where $u, v \in V(G) \cup F(G)$. Both graphs D_r and D_b contain all edges of G. The direction of an edge (u, v) can be determined by $t(u)$, $t(v)$, and (E_r, E_b). An edge $u \to v$ is in D_r iff (i) $u \to v \in E_r$ and $t_r(u) = \oplus$, (ii) $v \to u \in E_r$ and $t_r(v) = \ominus$, (iii) $u \to v \in E_b$ and $t_r(v) = \ominus$, or (iv) $v \to u \in E_b$ and $t_r(u) = \oplus$. Similarly, $u \to v$ is in D_b iff (i) $u \to v \in E_b$ and $t_b(u) = \oplus$, (ii) $v \to u \in E_b$ and $t_b(v) = \ominus$, (iii) $u \to v \in E_r$ and $t_b(v) = \ominus$, or (iv) $v \to u \in E_r$ and $t_b(u) = \oplus$.

We need to ensure that the L-contact representation is non-crossing. The inequalities above are not sufficient to achieve this. Therefore we add additional inequalities for each inner face. These inequalities ensure that each inner face does not cross itself in the L-contact representation. The inequalities for each type of face are shown in Table 1; see full version of the paper for more details [17].

Lemma 3.4. *The graphs D_r and D_b are acyclic.*

The lemma above is straightforward yet tedious to prove, and hence the proof is in the full version of the paper [17].

3.3 Construction of L-contact Representation Given a planar Laman graph G, an L-contact representation of G is constructed as follows:

(1) Find a planar Henneberg construction for G.

(2) Compute an angular tree T of G (Theorem 2.1).

(3) Compute the angle and edge labeling of G w.r.t. T (Theorem 2.2 and 2.3).

(4) Compute the type of every vertex of G according to the type rule in Section 3.1. This can be computed using a simple traversal of the trees E_r and E_b.

(5) Define the directed graphs D_r and D_b as described in Section 3.2

(6) Compute a topological order of D_r and D_b (which is possible as they are acyclic by Lemma 3.4) and let, for every vertex v in G, $x(v)$ and $y(v)$ be the number of v in these topological orders, respectively.

(7) For every non-special vertex v with $v \to u$ in E_r and $v \to w$ in E_b define an L-shape $\mathcal{L}(v)$ whose horizontal leg spans from $x(v)$ to $x(u)$ on y-coordinate $y(v)$ and whose vertical leg spans from $y(v)$ to $y(w)$ on x-coordinate $x(v)$.

Let n be the number of vertices of G. By Theorem 2.1 we can compute an angular tree of G in $\mathcal{O}(n^2)$ time. The angle labeling w.r.t. T can be computed in $\mathcal{O}(n)$ time using the linear time algorithm of de Fraysseix and Ossona de Mendez [7]. Similarly, the edge labeling w.r.t. T can be computed by a simple traversal of the tree H described in the proof of Theorem 2.3. It is easy to see that the remaining steps of our algorithm can also be computed in $\mathcal{O}(n)$ time. Finally note that the vertices of D_r and D_b that correspond to inner faces of G do not need to be included in the topological order of D_r and D_b. Hence every coordinate used in the L-contact representation is between 1 and n.

Theorem 3.3. *The algorithm above computes an L-contact representation of G on an $n \times n$ grid in $\mathcal{O}(n^2)$ time, where n is the number of vertices of G. If an angular tree is given, then the algorithm runs in $\mathcal{O}(n)$ time.*

4 Future Work and Open Problems

Using our newly discovered combinatorial structure, we showed that planar Laman graphs are L-contact graphs. A detailed example illustrating the constructive algorithm is shown in Fig. 8. Thus, we showed that axis-aligned L's are as "powerful" as segments with arbitrary slopes when it comes to *contact representation of planar graphs* [1]. The equivalent result is not true for *intersection representation of planar graphs*. Indeed there is no k such that all segment intersection graphs have an intersection representation with axis-aligned paths with no more than k bends each [3].

We think that L-contact representations can be used in various settings. For example, by "fattening" the L's we can get proportional side-contact representations similar to those in [1].

Several natural open problems follow from our results:

1. There are L-contact graphs that are not Laman graphs (e.g. K_4). All L-contact graphs are planar and satisfy $|E(W)| \leq 2|W| - 2$ for all $W \subseteq V$. Are these conditions also sufficient?

2. The L-contact representations resulting from our algorithm use all four types of L-shapes. If we limit ourselves to only type-I L's we can represent planar graphs of tree-width at most 2, which include outerplanar graphs. What happens if we limit ourselves to only type-I L's **and** allow degenerate L's?

3. Not every edge labeling corresponds to an angular tree. What are the necessary conditions for an edge labeling to have a corresponding (not necessarily proper) L-contact representation?

4. Planar Laman graphs can be characterized by the existence of an angular tree, which we can compute in $\mathcal{O}(n^2)$ time. This is slower than the fastest known algorithm for recognizing Laman graphs, which runs in $\mathcal{O}(n^{3/2}\sqrt{\log n})$ time [5]. Can we compute angular trees faster, as to obtain a faster algorithm for recognizing planar Laman graphs?

Acknowledgments. Work on this paper began during the *Bertinoro Workshop on Graph Drawing*. The authors gratefully acknowledge the other participants for useful discussions.

References

[1] M. J. Alam, T. Biedl, S. Felsner, M. Kaufmann, and S. G. Kobourov. Proportional contact representations of planar graphs. In *Proc. 19th Symposium on Graph Drawing*, pages 26–38, 2011.

[2] S. Bereg. Certifying and constructing minimally rigid graphs in the plane. In *Proc. 21st Symposium on Computational Geometry*, pages 73–80. 2005.

[3] S. Chaplick, V. Jelínek, J. Kratochvíl, and T. Vyskčil. Bend-bounded path intersection graphs: Sausages, noodles, and waffles on a grill. In *Proc. 38th International Workshop on Graph Theoretic Concepts in Computer Science*. 2012.

[4] J. Czyzowicz, E. Kranakis, and J. Urrutia. A simple proof of the representation of bipartite planar graphs as the contact graphs of orthogonal straight line segments. *Information Processing Letters*, 66(3):125–126, 1998.

[5] O. Daescu and A. Kurdia. Towards an optimal algorithm for recognizing laman graphs. *Journal of Graph Algorithms and Applications*, 13(2):219–232, 2009.

[6] N. de Castro, F. J. Cobos, J. C. Dana, A. Márquez, and M. Noy. Triangle-free planar graphs and segment intersection graphs. *Journal of Graph Algorithms and Applications*, 6(1):7–26, 2002.

[7] H. de Fraysseix and P. Ossona de Mendez. On topological aspects of orientations. *Discrete Mathematics*, 229:57–72, 2001.

[8] H. de Fraysseix and P. Ossona de Mendez. Representations by contact and intersection of segments. *Algorithmica*, 47:453–463, 2007.

[9] H. de Fraysseix, P. Ossona de Mendez, and J. Pach. Representation of planar graphs by segments. *Intuitive Geometry*, 63:109–117, 1991.

[10] H. de Fraysseix, P. Ossona de Mendez, and P. Rosenstiehl. On triangle contact graphs. *Combinatorics, Probability & Computing*, 3:233–246, 1994.

[11] S. Felsner. Lattice structures from planar graphs. *Electronic Journal of Combinatorics*, 11:R15, 2004.

[12] S. Felsner and M. C. Francis. Contact representations of planar graphs with cubes. In *Proc. 27th Symposium on Computational Geometry*, pages 315–320, 2011.

[13] S. Felsner, C. Huemer, S. Kappes, and D. Orden. Binary labelings for plane quadrangulations and their relatives. *Discrete Mathematics & Theoretical Computer Science*, 12(3):115–138, 2010.

[14] É. Fusy. Transversal structures on triangulations: A combinatorial study and straight-line drawings. *Discrete Mathematics*, 309(7):1870–1894, 2009.

[15] J. Graver, B. Servatius, and H. Servatius. *Combinatorial Rigidity*. Grad. Stud. in Math. 2. American Math. Soc., 1993.

[16] R. Haas, D. Orden, G. Rote, F. Santos, B. Servatius, H. Servatius, D. Souvaine, I. Streinu, and W. Whiteley. Planar minimally rigid graphs and pseudo-triangulations. *Computational Geometry*, 31:31–61, 2005.

[17] S. G. Kobourov, T. Ueckerdt, and K. Verbeek. Combinatorial and geometric properties of planar Laman graphs. *CoRR*, abs/1209.6506, 2012.

[18] P. Koebe. Kontaktprobleme der konformen Abbildung. *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig. Math.-Phys. Klasse*, 88:141–164, 1936.

[19] G. Laman. On graphs and rigidity of plane skeletal structures. *Journal of Engineering Mathematics*, 4:331–340, 1970.

[20] P. Rosenstiehl and R. Tarjan. Rectilinear planar layouts and bipolar orientations of planar graphs. *Discrete & Computational Geometry*, 1:343–353, 1986.

[21] W. Schnyder. Planar graphs and poset dimension. *Order*, 5(4):323–343, 1989.

[22] W. Schnyder. Embedding planar graphs on the grid. In *Proc. 1st ACM-SIAM Symposium on Discrete Algorithms*, pages 138–148, 1990.

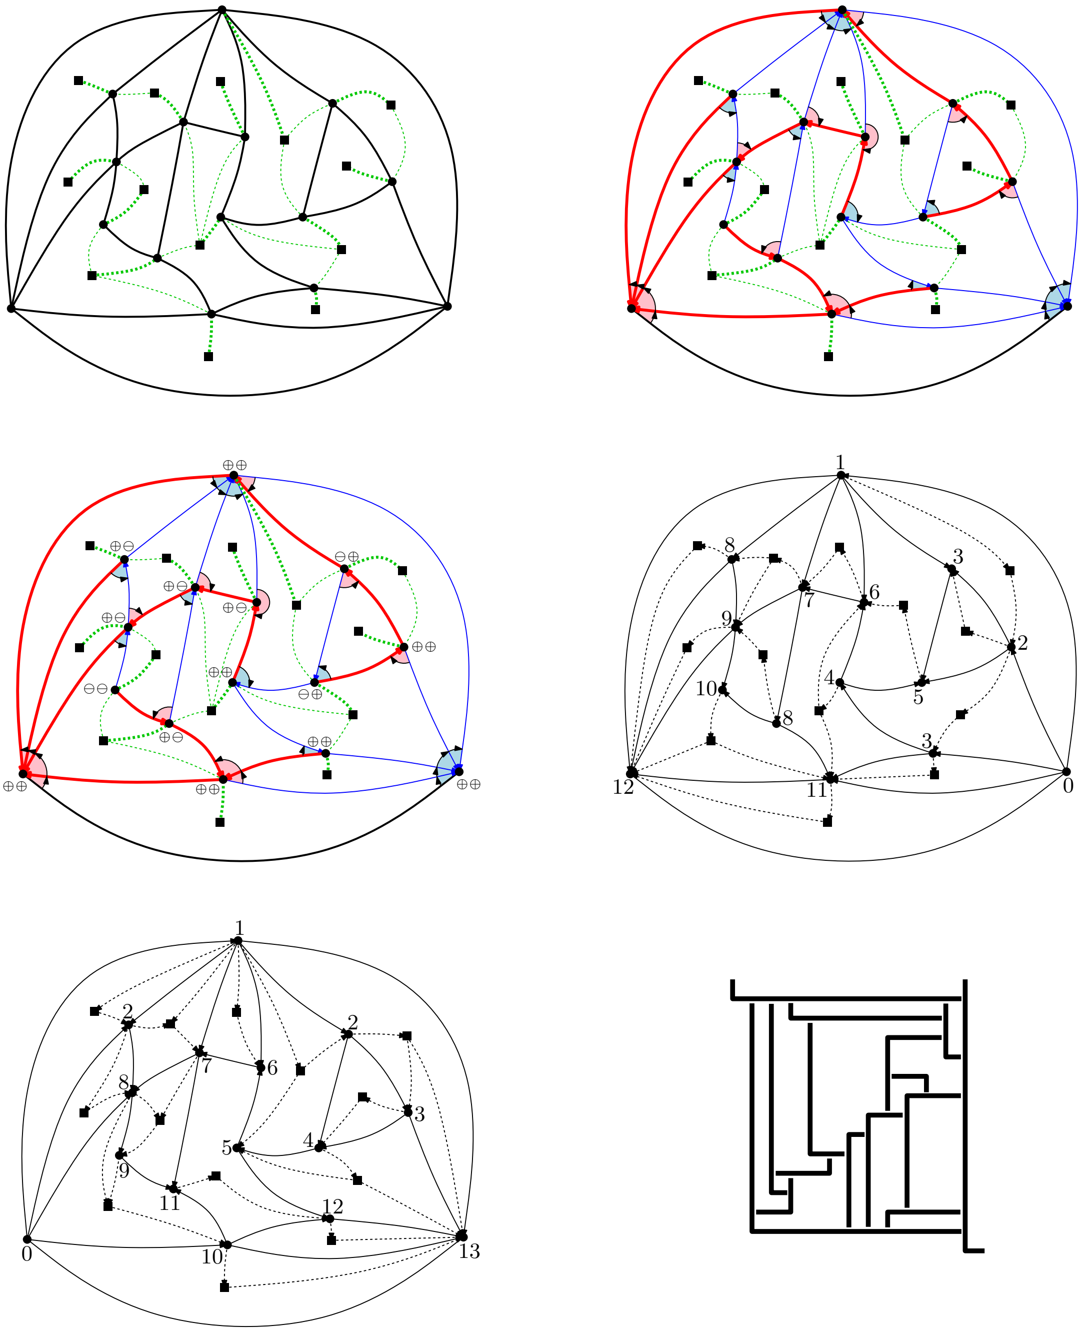

Figure 8: Left top: angular tree, and the corresponding matching (thick). Right top: edge labeling corresponding to angular tree. Middle left: vertex types. Middle right: inequality graph D_r plus x-coordinates. Bottom left: inequality graph D_b plus y-coordinates. Bottom right: L-contact representation.

Articles on Covering Problems

Local and Union Boxicity

Thomas Bläsius, Peter Stumpf and Torsten Ueckerdt

September 30, 2016

Abstract

The boxicity $\text{box}(H)$ of a graph H is the smallest integer d such that H is the intersection of d interval graphs, or equivalently, that H is the intersection graph of axis-aligned boxes in $\mathbb{R}^d$. These intersection representations can be interpreted as covering representations of the complement H^c of H with co-interval graphs, that is, complements of interval graphs. We follow the recent framework of global, local and folded covering numbers (Knauer and Ueckerdt, *Discrete Mathematics* **339** (2016)) to define two new parameters: the local boxicity $\text{box}_\ell(H)$ and the union boxicity $\overline{\text{box}}(H)$ of H. The union boxicity of H is the smallest d such that H^c can be covered with d vertex-disjoint unions of co-interval graphs, while the local boxicity of H is the smallest d such that H^c can be covered with co-interval graphs, at most d at every vertex.

We show that for every graph H we have $\text{box}_\ell(H) \leq \overline{\text{box}}(H) \leq \text{box}(H)$ and that each of these inequalities can be arbitrarily far apart. Moreover, we show that local and union boxicity are also characterized by intersection representations of appropriate axis-aligned boxes in $\mathbb{R}^d$. We demonstrate with a few striking examples, that in a sense, the local boxicity is a better indication for the complexity of a graph, than the classical boxicity.

1 Introduction

An *interval graph* is an intersection graph of intervals on the real line[1]. Such a set $\{I(v) \subseteq \mathbb{R} \mid v \in V(H)\}$ of intervals with $vw \in E(H) \Leftrightarrow I(v) \cap I(w) \neq \varnothing$ is called an *interval representation of* H. A box in $\mathbb{R}^d$, also called a d*-dimensional box*, is the Cartesian product of d intervals. The *boxicity* of a graph H, denoted by $\text{box}(H)$, is the least integer d such that H is the intersection graph of d-dimensional boxes, and a corresponding set $\{B(v) \subseteq \mathbb{R}^d \mid v \in V(H)\}$ is a *box representation of* H. The boxicity was introduced by Roberts [17] in 1969 and has many applications in as diverse areas as ecology and operations research [4].

As two d-dimensional boxes intersect if and only if each of the d corresponding pairs of intervals intersect, we have the following more graph theoretic interpretation of the boxicity of a graph; also see Figure 1(a).

Theorem 1 (Roberts [17])**.** *For a graph H we have $\text{box}(H) \leq d$ if and only if $H = G_1 \cap \cdots \cap G_d$ for some interval graphs $G_1, \ldots, G_d$.*

I.e., the boxicity of a graph H is the least integer d such that H is the intersection of some d interval graphs. For a graph $H = (V, E)$ we denote its

[1] Throughout, we shall just say "intervals" and drop the suffix "on the real line".

complement by $H^c = (V, \binom{V}{2} - E)$. Then by De Morgan's law we have

$$H = G_1 \cap \cdots \cap G_d \iff H^c = G_1^c \cup \cdots \cup G_d^c, \tag{1}$$

i.e., $\text{box}(H)$ is the least integer d such that the complement H^c of H is the union of d co-interval graphs $G_1^c, \ldots, G_d^c$, where a *co-interval graph* is the complement of an interval graph[2]. In other words, $\text{box}(H) \le d$ if H^c can be covered with d co-interval graphs. Strictly speaking, we have to be a little more precise here. In order to use De Morgan's law, we should guarantee that $G_1, \ldots, G_d$ in (1) all have the same vertex set. To this end, if G is a subgraph of H, let $\bar{G} = (V(H), E(G))$ be the graph obtained from G by adding all vertices in $V(H) - V(G)$ as isolated vertices. (Whenever we use $\bar{G}$ it will be clear from the context which supergraph H of G we consider.) Clearly we have

$$H^c = G_1^c \cup \cdots \cup G_d^c \quad \Rightarrow \quad H^c = \bar{G}_1^c \cup \cdots \cup \bar{G}_d^c \quad \Rightarrow \quad H = \bar{G}_1 \cap \cdots \cap \bar{G}_d$$

for any graph H and any subgraphs $G_1, \ldots, G_d$ of H. Now whenever G is a co-interval graph, then so is $\bar{G}$, implying that $\text{box}(H)$ is the least integer d such that H^c can be covered with d co-interval graphs.

Graph covering parameters. In the general graph covering problem one is given an input graph H, a so-called covering class $\mathcal{G}$ and a notion of how to cover H with one or more graphs from $\mathcal{G}$. The most classic notion of covering, which also corresponds to the boxicity as discussed above, is that H shall be the union of $G_1, \ldots, G_t \in \mathcal{G}$, i.e., $V(H) = \bigcup_{i \in [t]} V(G_i)$ and $E(H) = \bigcup_{i \in [t]} E(G_i)$. (Here and throughout the paper, for a positive integer t we denote $[t] = \{1, \ldots, t\}$.) The *global covering number*, denoted by $c_g^{\mathcal{G}}(H)$, is then defined to be the minimum t for which such a cover exists. Many important graph parameters can be interpreted as a global covering number, e.g., the arboricity [15], the track number [9] (this is not the track-number as defined in [5]) and the thickness [1, 14], just to name a few.

Most recently, Knauer and Ueckerdt [11] suggested the following unifying framework for three kinds of covering numbers, differing in the underlying notion of covering. A *graph homomorphism* is a map $\varphi : V(G) \to V(H)$ with the property that if $uv \in E(G)$ then $\varphi(u)\varphi(v) \in E(H)$, i.e., φ maps vertices of G (not necessarily injectively) to vertices of H such that edges are mapped to edges. For abbreviation we shall simply write $\varphi : G \to H$ instead of $\varphi : V(G) \to V(H)$. For an input graph H, a covering class $\mathcal{G}$ and a positive integer t, a *t-global $\mathcal{G}$-cover of H* is an edge-surjective homomorphism $\varphi : G_1 \uplus \cdots \uplus G_t \to H$ such that $G_i \in \mathcal{G}$ for each $i \in [t]$. Here $\uplus$ denotes the vertex-disjoint union of graphs. We say that φ is *injective* if its restriction to G_i is injective for each $i \in [t]$. A $\mathcal{G}$-cover is called *s-local* if $|\varphi^{-1}(v)| \le s$ for every $v \in V(H)$.

Hence, if φ is a $\mathcal{G}$-cover of H, then

φ is t-global if it uses only t graphs from the covering class $\mathcal{G}$,

φ is injective if $\varphi(G_i)$ is a copy of G_i in H for each $i \in [t]$,

φ is s-local if for each $v \in V(H)$ at most s vertices are mapped onto v.

[2]Equivalently, these are the comparability graphs of interval orders.

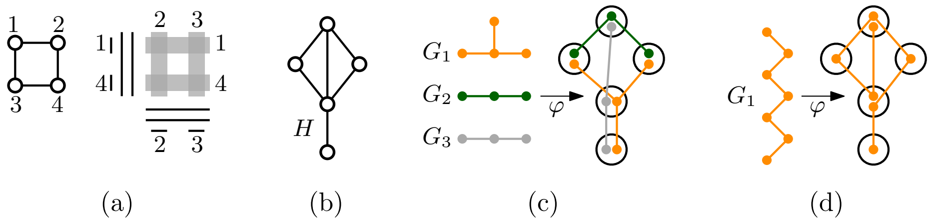

Figure 1: (a) The 4-cycle as the intersection of two interval graphs. (b) Example graph H. (c) An injective covering of H that is 3-global and 2-local. (d) A (non-injective) 1-global 2-local covering of H.

For a covering class $\mathcal{G}$ and an input graph H the *global covering number* $c_g^{\mathcal{G}}(H)$, the *local covering number* $c_\ell^{\mathcal{G}}(H)$, and the *folded covering number* $c_f^{\mathcal{G}}(H)$ are then defined as follows; see also Figure 1(b)–(d):

$$c_g^{\mathcal{G}}(H) = \min\{t : \text{there exists a } t\text{-global injective } \mathcal{G}\text{-cover of } H\}$$

$$c_\ell^{\mathcal{G}}(H) = \min\{s : \text{there exists an } s\text{-local injective } \mathcal{G}\text{-cover of } H\}$$

$$c_f^{\mathcal{G}}(H) = \min\{s : \text{there exists a 1-global } s\text{-local } \mathcal{G}\text{-cover of } H\}$$

Intuitively speaking, for $c_\ell^{\mathcal{G}}(H)$ we want to represent the input graph H as the union of graphs from the covering class $\mathcal{G}$, where the number of graphs we use is not important. Rather we want to "use" each vertex of H in only few of these subgraphs. For $c_f^{\mathcal{G}}(H)$ it is convenient to think of the "inverse" mapping for φ. If $\varphi : G_1 \to H$ is a 1-global $\mathcal{G}$-cover of H, then the preimage under φ of a vertex $v \in V(H)$ is an independent set S_v in G_1. Moreover, for every $u, v \in V(H)$ we have $uv \in E(H)$ if and only if there is at least one edge between S_u and S_v in G_1. So G_1 is obtained from H by a series of *vertex splits*, where splitting a vertex v into an independent set S_v is such that for each edge vw incident to v there is at least one edge between w and S_v after the split. Now $c_f^{\mathcal{G}}(H)$ is the smallest s such that each vertex can be split into at most s vertices so that the resulting graph G_1 lies in the covering class $\mathcal{G}$.

It is known, that if the covering class $\mathcal{G}$ is closed under certain graph operations, we can deduce inequalities between the folded, local and global covering numbers. For a graph class $\mathcal{G}$ we define the following.

- $\mathcal{G}$ is *homomorphism-closed* if for any connected $G \in \mathcal{G}$ and any homomorphism $\varphi : G \to H$ into some graph H we have that $\varphi(G) \in \mathcal{G}$.
- $\mathcal{G}$ is *hereditary* if for any $G \in \mathcal{G}$ and any induced subgraph G' of G we have that $G' \in \mathcal{G}$.
- $\mathcal{G}$ is *union-closed* if for any $G_1, G_2 \in \mathcal{G}$ we have that $G_1 \cup G_2 \in \mathcal{G}$.

Proposition 2 (Knauer-Ueckerdt [11]). *For every input graph H and every covering class $\mathcal{G}$ we have*

(i) $c_\ell^{\mathcal{G}}(H) \le c_g^{\mathcal{G}}(H)$, and if $\mathcal{G}$ is union-closed, then $c_f^{\mathcal{G}}(H) \le c_\ell^{\mathcal{G}}(H)$,

(ii) if G is hereditary and homomorphism-closed, then $c_f^{\mathcal{G}}(H) \ge c_\ell^{\mathcal{G}}(H)$.

Boxicity variants. Let us put the boxicity into the graph covering framework by Knauer and Ueckerdt [11] as described above. To this end, let $\mathcal{C}$ denote the class of all co-interval graphs. Then we have $\mathrm{box}(H) = c_g^{\mathcal{C}}(H^c)$ and we can investigate the new parameters

$$\mathrm{box}_f(H) := c_f^{\mathcal{C}}(H^c) \quad \text{and} \quad \mathrm{box}_\ell(H) := c_\ell^{\mathcal{C}}(H^c).$$

Clearly, if H is an interval graph, i.e., $H^c \in \mathcal{C}$, then $\mathrm{box}_f(H) = \mathrm{box}_\ell(H) = \mathrm{box}(H) = 1$. As it turns out, if H is not an interval graph, then $\mathrm{box}_f(H)$ is not very meaningful.

Theorem 3. *For every graph H we have $\mathrm{box}_f(H) = 1$ if $H^c \in \mathcal{C}$ and $\mathrm{box}_f(H) = \infty$ otherwise.*

Basically, Theorem 3 says that if H^c is not a co-interval graph, there is no way to obtain a co-interval graph from H^c by vertex splits. For example, if H has an induced 4-cycle and hence H^c has two independent edges, then $H^c \notin \mathcal{C}$ and whatever vertex splits are applied, the result will always have two independent edges, i.e., not be a co-interval graph. To overcome this issue, it makes sense to define $\overline{\mathcal{C}}$ to be the class of all vertex-disjoint unions of co-interval graphs and consider the parameters

$$\overline{\mathrm{box}}(H) := c_g^{\overline{\mathcal{C}}}(H^c), \quad \overline{\mathrm{box}}_\ell(H) := c_\ell^{\overline{\mathcal{C}}}(H^c), \quad \overline{\mathrm{box}}_f(H) := c_f^{\overline{\mathcal{C}}}(H^c).$$

We have defined in total six boxicity-related graph parameters, one of which (namely $\mathrm{box}_f(H)$) turned out to be meaningless by Theorem 3. Somehow luckily, three of the remaining five parameters always coincide.

Theorem 4. *For every graph H we have $\mathrm{box}_\ell(H) = \overline{\mathrm{box}}_\ell(H) = \overline{\mathrm{box}}_f(H)$.*

Proposition 2 gives $\overline{\mathrm{box}}_\ell(H) = c_\ell^{\overline{\mathcal{C}}}(H^c) \le c_g^{\overline{\mathcal{C}}}(H^c) = \overline{\mathrm{box}}(H)$ for every input graph H. As $\mathcal{C} \subset \overline{\mathcal{C}}$ we have $\overline{\mathrm{box}}(H) = c_g^{\overline{\mathcal{C}}}(H^c) \le c_g^{\mathcal{C}}(H^c) = \mathrm{box}(H)$ for every input graph H. Hence with Theorem 4 for every graph H the remaining three boxicity-related parameters fulfil:

$$\mathrm{box}_\ell(H) \le \overline{\mathrm{box}}(H) \le \mathrm{box}(H). \tag{2}$$

We refer to $\mathrm{box}_\ell(H)$ as the *local boxicity of H* and to $\overline{\mathrm{box}}(H)$ as the *union boxicity of H*. Indeed, the three parameters boxicity, local boxicity and union boxicity are non-trivial and reflect different aspects of the graph, as will be investigated in more detail in this paper.

Theorem 5. *For every positive integer k there exist graphs H_k, H'_k, H''_k with*

(i) $\mathrm{box}_\ell(H_k) \ge k$,

(ii) $\mathrm{box}_\ell(H'_k) = 2$ *and* $\overline{\mathrm{box}}(H'_k) \ge k$,

(iii) $\overline{\mathrm{box}}(H''_k) = 1$ *and* $\mathrm{box}(H''_k) = k$.

We also give geometric interpretations of the local and union boxicity of a graph H in terms of intersecting high-dimensional boxes. For positive integers k, d with $k \le d$ we call a d-dimensional box $B = I_1 \times \cdots \times I_d$ *k-local* if for at most k indices $i \in \{1, \ldots, d\}$ we have $I_i \neq \mathbb{R}$. Thus a k-local d-dimensional box is the Cartesian product of d intervals, at least $d - k$ of which are equal to the entire real line $\mathbb{R}$.

Theorem 6. *Let H be a graph.*

(i) We have $\overline{\text{box}}(H) \leq k$ if and only if there exist $d_1, \dots, d_k$ such that H is the intersection graph of Cartesian products of k boxes, where the ith box is 1-local d_i-dimensional, $i = 1, \dots, k$.

(ii) We have $\text{box}_\ell(H) \leq k$ if and only if there exists some d such that H is the intersection graph of k-local d-dimensional boxes.

There is a number of results in the literature stating that the boxicity of certain graphs is low, for which we can easily see that the local boxicity is even lower. Indeed, often an intersection representation with d-dimensional boxes is constructed, in order to show that $\text{box}(H) \leq d$, and in many cases these representations consist of s-local d-dimensional boxes for some $s < d$ (or can be turned into such quite easily). Hence, with Theorem 6 we can conclude in such cases that $\text{box}_\ell(H) \leq s$.

Let us restrict here to one such case, which is comparably simple. For a graph H the *acyclic chromatic number*, denoted by $\chi_a(H)$, is the smallest k such that there exists a proper vertex coloring of H with k colors in which any two color classes induce a forest. In other words, an acyclic coloring has no monochromatic edges and no bicolored cycles. Esperet and Joret [6] have recently shown that for any graph H with $\chi_a(H) = k$ we have $\text{box}(H) \leq k(k-1)$. Indeed, their proof (which we include here for completeness) gives an intersection representation of H with $2(k-1)$-local $k(k-1)$-dimensional boxes, implying the following theorem.

Theorem 7. *For every graph H we have $\text{box}_\ell(H) \leq 2(\chi_a(H) - 1)$.*

Proof. Let c be an acyclic coloring of H with k colors. For any pair $\{i, j\}$ of colors consider the subgraph $G_{i,j}$ induced by the vertices of colors i and j. As $G_{i,j}$ is a forest, we have $\text{box}(G_{i,j}) \leq 2$ (this follows from [18] but can also be seen fairly easily). Moreover, since H is the union of all $G_{i,j}$, the complement H^c of H is the intersection of the complements of all $\bar{G}_{i,j}$ (note the use of $\bar{G}_{i,j}$ instead of $G_{i,j}$ here).

Now take an intersection representation of $G_{i,j}$ with 2-dimensional boxes and extend it to one for $\bar{G}_{i,j}$ by putting the box $\mathbb{R}^2$ for each vertex colored neither i nor j. Then the Cartesian product of all these $\binom{k}{2}$ box representations is an intersection representation of H with $2(k-1)$-local $k(k-1)$-dimensional boxes. This proves that $\text{box}(H) \leq k(k-1)$ and $\text{box}_\ell(H) \leq 2(k-1)$, as desired. □

Organization of the paper. In Section 2 we prove Theorem 3, i.e., that $\text{box}_f(H)$ is meaningless, and Theorem 4, i.e., that three of the remaining five boxicity variants coincide. In Section 3 we consider the problem of separation for boxicity and its local and union variants, that is, we give a proof of Theorem 5. In Section 4 we describe and prove the geometric interpretations of local and union boxicity from Theorem 6. Finally, we give some concluding remarks and open problems in Section 5.

2 Local and Union Boxicity

Recall that a graph class $\mathcal{G}$ is homomorphism-closed if for every *connected* graph $G \in \mathcal{G}$ and any homorphism $\varphi : G \to H$ into some graph H we have $\varphi(G) \in \mathcal{G}$.

Since φ is a homomorphism, $\varphi(G)$ arises from G by a series of "inverse vertex splits", i.e., an independent set in G is identified into a single vertex of $\varphi(G)$. If $\mathcal{G}$ is not only homomorphism-closed, but also closed under identifying non-adjacent vertices in *disconnected* graphs, then the folded covering number $c_f^{\mathcal{G}}$ turns out to be somewhat meaningless.

Lemma 8. *If a covering class $\mathcal{G}$ is closed under identifying non-adjacent vertices, then for every non-empty input graph H we have*

$$c_f^{\mathcal{G}}(H) < \infty \iff H \in \mathcal{G} \iff c_f^{\mathcal{G}}(H) = 1.$$

Proof. The right equivalence follows by definition of $c_f^{\mathcal{G}}(H)$.

The implication $H \in \mathcal{G} \Rightarrow c_f^{\mathcal{G}}(H) < \infty$ in the first equivalence is thereby obvious, and it is left to show that $c_f^{\mathcal{G}}(H) = 1$ whenever $c_f^{\mathcal{G}}(H) < \infty$. So let $\varphi : G_1 \to H$ be any 1-global cover of H. We do induction over $|V(G_1)|$, the number of vertices in G_1.

If $|V(G_1)| = |V(H)|$, i.e., no vertices are folded, then φ is injective and therefore $c_f^{\mathcal{G}}(H) = 1$. So assume that $|V(G_1)| > |V(H)|$ and let v, w be distinct vertices in G_1 with $\varphi(v) = \varphi(w)$. Consider the graph G_1' that we obtain by identifying v and w in G_1. Since $\varphi(v) = \varphi(w)$ is only possible if v and w are non-adjacent, and $\mathcal{G}$ is closed under identifying non-adjacent vertices we know that $G_1' \in \mathcal{G}$. Now the 1-global $\mathcal{G}$-cover $\varphi : G_1 \to H$ induces a 1-global $\mathcal{G}$-cover $\varphi' : G_1' \to H$ by $\varphi = \varphi' \circ \psi$, where $\psi : G_1 \to G_1'$ identifies v and w in G_1 and fixes all other vertices. As $|V(G_1')| = |V(G_1)| - 1$, we can apply induction to φ' to conclude that $c_f^{\mathcal{G}}(H) = 1$. □

Lemma 9. *Let $\mathcal{C}$ be the class of all co-interval graphs and $\overline{\mathcal{C}}$ be the class of all vertex-disjoint unions of co-interval graphs. Then*

(i) $\mathcal{C}$ and $\overline{\mathcal{C}}$ are hereditary,

(ii) $\mathcal{C}$ is closed under identifying non-adjacent vertices, and

(iii) $\overline{\mathcal{C}}$ is homomorphism-closed.

Proof. (i) Consider any graph $G \in \overline{\mathcal{C}}$. Then $G = G_1 \cupdot \cdots \cupdot G_t$ for some $G_1, \ldots, G_t \in \mathcal{C}$. If $G \in \mathcal{C}$, then $t = 1$. For $i \in [t]$ consider an intersection representation $\{I_i(v) \mid v \in V(G_i)\}$ of G_i^c with intervals. For any vertex set $S \subseteq V(G)$, consider the induced subgraphs when restricted to vertices in S, i.e., $G' = G[S]$ and $G_i' = G_i[V(G_i) \cap S]$ for $i \in [t]$. Note that $\{I_i(v) \mid v \in V(G_i) \cap S\}$ is an interval representation of $(G_i')^c$, i.e., $G_i' \in \mathcal{C}$. Hence $G' = G_1' \cupdot \cdots \cupdot G_t' \in \overline{\mathcal{C}}$ and $G' \in \overline{\mathcal{C}}$ if $t = 1$. This shows that $\mathcal{C}$ and $\overline{\mathcal{C}}$ are hereditary.

(ii) Let $G \in \mathcal{C}$, x, y be two non-adjacent vertices in G and $\{I(v) \mid v \in V(G)\}$ be an intersection representation of G with intervals. Let G' be the graph obtained from G by identifying x and y into a single vertex z. Since $xy \in E(G^c)$ we have $I(x) \cap I(y) \neq \emptyset$ and hence $I(z) := I(x) \cap I(y)$ is a non-empty interval. As for any interval J we have $J \cap I(z) \neq \emptyset$ if and only if $J \cap I(x) \neq \emptyset$ or $J \cap I(y) \neq \emptyset$ or both, we have that $\{I(v) \mid v \in V(G), v \neq x, y\} \cup \{I(z)\}$ is an intersection representation of $(G')^c$ and thus $G' \in \mathcal{C}$, as desired.

(iii) If $G \in \overline{\mathcal{C}}$ then $G = G_1 \cup \cdots \cup G_t$ for some $G_1, \ldots, G_t \in \mathcal{C}$. If x, y are two non-adjacent vertices in the same connected component, then x, y are in the same G_i, say G_1. By (ii) identifying x and y in G_1 gives a graph $G_1' \in \mathcal{C}$. Moreover, identifying x and y in G gives a graph $G' = G_1' \cup G_2 \cup \cdots \cup G_t$. As $G_1' \in \mathcal{C}$ we have $G' \in \overline{\mathcal{C}}$ and hence $\overline{\mathcal{C}}$ is homomorphism-closed. □

Proof of Theorem 3. This is a direct corollary of Lemma 8 and Lemma 9 (ii). □

Proof of Theorem 4. We have that $\overline{\mathcal{C}}$ is hereditary by Lemma 9 (i), homomorphism-closed by Lemma 9 (iii) and union-closed by definition. Hence by Proposition 2 we have $\overline{\mathrm{box}}_f(H) = c_f^{\overline{\mathcal{C}}}(H^c) = c_\ell^{\overline{\mathcal{C}}}(H^c) = \overline{\mathrm{box}}_\ell(H)$.

As $\mathcal{C} \subset \overline{\mathcal{C}}$ we clearly have $\overline{\mathrm{box}}_\ell(H) = c_\ell^{\overline{\mathcal{C}}}(H^c) \leq c_\ell^{\mathcal{C}}(H^c) = \mathrm{box}_\ell(G)$. Finally, consider any s-local t-global $\overline{\mathcal{C}}$-cover $\varphi : G_1 \cup \cdots \cup G_t \to H^c$. For $i = 1, \ldots, t$ we have $G_i \in \overline{\mathcal{C}}$ and hence G_i is the vertex-disjoint union of some graphs in $\mathcal{C}$. Thus we can interpret φ as an s-local t'-global $\mathcal{C}$-cover of H^c for some $t' \geq t$. This shows that $\mathrm{box}_\ell(H) = c_\ell^{\mathcal{C}}(H^c) \leq c_\ell^{\overline{\mathcal{C}}}(H^c) = \overline{\mathrm{box}}_\ell(H)$ and thus concludes the proof. □

3 Separating the Variants

Proof of Theorem 5.

(i) For a fixed integer $k \geq 1$ we consider any graph F_k that is $2k$-regular and has girth at least 6 (i.e., its shortest cycle has length at least 6). Now let φ be an injective s-local $\mathcal{C}$-cover of F_k, i.e., a cover of $E(F_k)$ with t co-interval graphs $G_1, \ldots, G_t \subseteq F_k$ for some $t \in \mathbb{N}$ such that every vertex of F_k is contained in at most s such G_i. We shall show that $s \geq k$, proving that $c_\ell^{\mathcal{C}}(F_k) \geq k$ and hence $\mathrm{box}_\ell(H_k) \geq k$, where $H_k = F_k^c$ denotes the complement of F_k.

A co-interval graph G does not contain any induced matching on two edges. Hence G does not contain any induced cycle of length at least 6. (Moreover, as G is perfect, it also contains no induced cycles of length 5.) Since F_k has girth at least 6, this implies that every subgraph of F_k that is a co-interval graph is a forest. In particular, every G_i has average degree less than 2, i.e., $\sum_{v \in V(G_i)} \deg_{G_i}(v) < 2|V(G_i)|$. We conclude that

$$2k \cdot |V(F_k)| = \sum_{v \in V(F_k)} \deg_{F_k}(v) \leq \sum_{v \in V(F_k)} \sum_{\substack{i \in [t] \\ v \in V(G_i)}} \deg_{G_i}(v)$$
$$= \sum_{i=1}^{t} \sum_{v \in V(G_i)} \deg_{G_i}(v) < \sum_{i=1}^{t} 2|V(G_i)| \leq 2s \cdot |V(F_k)|,$$

where the first inequality holds since every edge of F_k is covered and the last inequality holds since every vertex is contained in at most s of the G_i, $i = 1, \ldots, t$. From the above it follows that $s \geq k$, as desired.

(ii) Our proof follows the ideas of Milans *et al.* [13], who consider $L(K_n)$, the line graph of K_n, and prove that $c_g^{\mathcal{I}}(L(K_n)) \to \infty$ for $n \to \infty$, while $c_\ell^{\mathcal{I}}(L(K_n)) = 2$ for every $n \in \mathbb{N}$, where $\mathcal{I}$ denotes the class of all interval graphs. However, instead of using the ordered Ramsey numbers (which is also possible in our case) we shall rather use the following hypergraph Ramsey numbers: Let K_n^3, $n \in \mathbb{N}$, denote the complete 3-uniform hypergraph on n vertices, i.e., $K_n^3 = ([n], \binom{[n]}{3})$. For an integer $k \geq 1$, the Ramsey number $R_k(K_6^3)$ is the smallest integer n such that every coloring of the hyperedges of K_n^3 with k colors contains a monochromatic copy of K_6^3. The hypergraph Ramsey theorem implies that $R_k(K_6^3)$ exists for every k [16].

Now for fixed $k \geq 1$, choose an integer $n = n(k) > R_k(K_6^3)$ and consider $L(K_n)$, the line graph of K_n. Let φ be any injective t-global $\overline{\mathcal{C}}$-cover of $L(K_n)$ with co-interval graphs $G_1, \ldots, G_t \subseteq L(K_n)$ for some $t \in \mathbb{N}$. We shall show that $t > k$, proving that $c_g^{\overline{\mathcal{C}}}(L(K_n)) > k$ and hence $\overline{\text{box}}(H'_k) > k$, where $H'_k = (L(K_n))^c$ denotes the complement of $L(K_n)$.

Assume for the sake of contradiction that $t \leq k$. From the $\overline{\mathcal{C}}$-cover φ of $L(K_n)$, we define a coloring c of $E(K_n^3)$ with t colors. Given $x, y, z \in [n]$ with $x < y < z$, let $c(x, y, z) = \min\{i \in [t] \mid \{xy, yz\} \in E(G_i)\}$ be the smallest index of a co-interval graph in $\{G_1, \ldots, G_t\}$ that covers the edge between xy and yz in $L(K_n)$. Since $n > R_k(K_6^3) \geq R_t(K_6^3)$ under c there is a monochromatic copy of K_6^3, say it is in color i and that its vertices are $\{x_1, \ldots, x_6\}$. This means that G_i has a connected component containing $x_1, \ldots, x_6$ and in particular the edges $\{x_1x_2, x_2x_3\}$ and $\{x_4x_5, x_5x_6\}$ of $L(K_n)$. However, these two edges induce a matching in $L(K_n)$ and hence also in that connected component of G_i. This is a contradiction to that component being a co-interval graph, and thus implies that $t > k$, as desired.

Finally, observe that for any $n \in \mathbb{N}$ the following is an injective 2-local $\mathcal{C}$-cover of $L(K_n)$: For each $i \in [n]$ let G_i be the clique in $L(K_n)$ formed by all edges incident to vertex i of K_n. Then $\{G_1, \ldots, G_n\}$ is a set of n co-interval graphs in $L(K_n)$ with the property that every edge of $L(K_n)$ lies in exactly one G_i and every vertex of $L(K_n)$ lies in exactly two G_i. This shows that $c_\ell^{\mathcal{C}}(L(K_n)) = \text{box}_\ell(H'_k) \leq 2$.

(iii) For fixed $k \geq 1$ consider M_k the matching on k edges. We shall show that $c_g^{\overline{\mathcal{C}}}(M_k) = 1$ and $c_g^{\mathcal{C}}(M_k) = k$, proving that $\overline{\text{box}}(H''_k) = 1$ and $\text{box}(H''_k) = k$, where $H''_k = M_k^c$ is the complement of M_k. Indeed, as every co-interval graph has at most one component containing an edge, any $\mathcal{C}$-cover of M_k contains at least k co-interval graphs to cover all k components of M_k. Since K_2 is a co-interval graph, there actually is an injective k-global $\mathcal{C}$-cover of M_k. Thus, we have $c_g^{\mathcal{C}}(M_k) = \text{box}(H''_k) = k$.

On the other hand, the class $\overline{\mathcal{C}}$ is union-closed and, since K_2 is a co-interval graph, $\overline{\mathcal{C}}$ contains all matchings. In particular $M_k \in \overline{\mathcal{C}}$ and therefore we have $c_g^{\overline{\mathcal{C}}}(M_k) = \overline{\text{box}}(H''_k) = 1$. □

4 Geometric Interpretations

Lemma 10. *A graph H is the intersection graph of 1-local d-dimensional boxes if and only if H^c is the vertex-disjoint union of d co-interval graphs.*

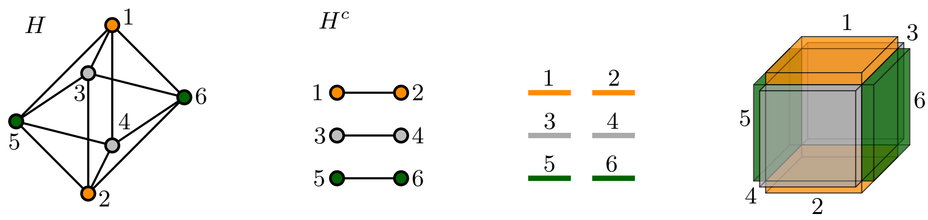

(a) (b) (c) (d)

Figure 2: (a) The octahedron H. (b) Its complement H^c. (c) H^c as the vertex-disjoint union of three co-interval graphs (given in their interval representation). (d) The corresponding intersection representation of H with 1-local 3-dimensional boxes. The two long sides of each box have actually infinite length.

Proof. For an illustration of the proof, see Figure 2. First, if $\{B(v) \mid v \in V(H)\}$ is an intersection representation of H with 1-local boxes in $\mathbb{R}^d$, then for each $v \in V(H)$ let $B(v) = I_1(v) \times \cdots \times I_d(v)$. Without loss of generality assume that for every $v \in V(H)$ there is some coordinate $i \in [d]$ for which $I_i(v) \neq \mathbb{R}$. For each $i \in [d]$ consider the set $V_i = \{v \in V(H) \mid I_i(v) \neq \mathbb{R}\}$ of those vertices v for which $B(v)$ is bounded in the i^{th} coordinate. Then $V_1, \ldots, V_d$ is a partition of $V(H)$ and for each $i \in [d]$ the set $\{I_i(v) \mid v \in V_i\}$ is an intersection representation with intervals of some graph G_i with vertex set V_i. Then we have $H = \bar{G}_1 \cap \cdots \cap \bar{G}_d$ and hence $H^c = \bar{G}_1^c \cup \cdots \cup \bar{G}_d^c = G_1^c \uplus \cdots \uplus G_d^c$. Thus H^c is the vertex-disjoint union of the d co-interval graphs, as desired.

Now let $H^c = G_1^c \uplus \cdots \uplus G_d^c$, where $G_i^c \in \mathcal{C}$ for $i = 1, \ldots, d$. Consider for each i an intersection representation $\{I_i(v) \mid v \in V(G_i)\}$ of the complement G_i of G_i^c with intervals. For $v \in V(H)$ we define

$$I_i'(v) = \begin{cases} I_i(v), & \text{if } v \in V(G_i) \\ \mathbb{R}, & \text{if } v \notin V(G_i). \end{cases}$$

Then $B(v) = I_1'(v) \times \cdots \times I_d'(v)$ is a 1-local d-dimensional box. Moreover, $\{B(v) \mid v \in V(H)\}$ is an intersection representation of H, which concludes the proof. □

From Lemma 10 we easily derive Theorem 6, i.e., the geometric intersection representations characterizing the local and union boxicity, respectively.

Proof of Theorem 6.

(i) This follows easily from Lemma 10. Indeed, if $\overline{\text{box}}(H) = c_g^{\overline{\mathcal{C}}}(H^c) \leq k$, then $H^c = G_1 \cup \cdots \cup G_k$ where for $i = 1, \ldots, k$ the graph $G_i \in \overline{\mathcal{C}}$ is the vertex-disjoint union of some d_i co-interval graphs. By Lemma 10 G_i^c has an intersection representation with 1-local d_i-dimensional boxes. Similarly to the proof of Lemma 10, extending this 1-local box representation of G_i^c to all vertices of H by adding a box $\mathbb{R}^{d_i}$ for each vertex in $H - G_i$, and taking the Cartesian product of these k extended 1-local box representations, we obtain an intersection representation of H of the desired kind.

Similarly, consider any intersection representation $\{B_1(v) \times \cdots \times B_k(v) \mid v \in V(H)\}$ of H, where for every $v \in V(H)$ and every $i \in [k]$ the box $B_i(v)$ is

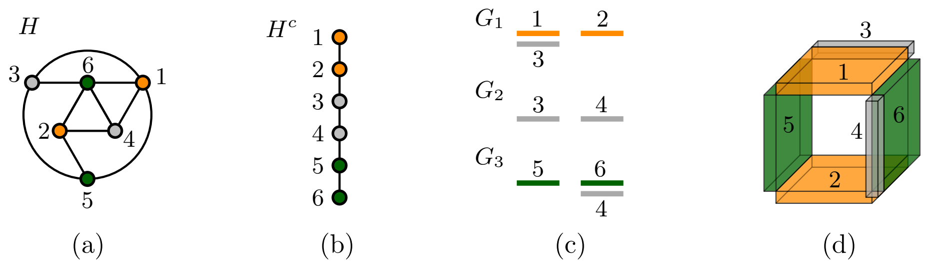

Figure 3: (a, b) A graph H and its complement H^c. (c) H^c can be covered using three co-interval graphs. (d) The resulting intersection representation. Note that the boxes are 3-dimensional as the cover uses three co-interval graphs and the boxes are 1-local and 2-local if the corresponding vertices are covered once $(1,2,5,6)$ and twice $(3,4)$, respectively. The long sides of each box have actually infinite length.

d_i-dimensional and 1-local. Then by Lemma 10 the set $\{B_i(v) \mid v \in V(H)\}$ is an intersection representation of some graph G_i whose complement G_i^c is in $\overline{\mathcal{C}}$. Moreover, H^c is the union of these k graph $G_1^c, \ldots, G_k^c \in \overline{\mathcal{C}}$. This gives $\overline{\text{box}}(H) = c_g^{\overline{\mathcal{C}}}(H^c) \leq k$, as desired.

(ii) For an example illustrating this case, see Figure 3. If $\text{box}_\ell(H) = c_\ell^{\mathcal{C}}(H^c) \leq k$, then there is a set $\{G_1, \ldots, G_t\}$ of t co-interval graphs such that $G_i \subseteq H^c$ for $i = 1, \ldots, t$, $E(H^c) = E(G_1) \cup \cdots \cup E(G_t)$ and every $v \in V(H^c)$ is contained in at most k such G_i, $i = 1, \ldots, t$. For each $i \in [t]$ consider an interval representation $\{I_i(v) \mid v \in V(G_i)\}$ of G_i^c. For $v \in H - G_i$ we set $I_i(v) = \mathbb{R}$. Note that $\{I_i(v) \mid v \in V(H)\}$ is an interval representation of $\bar{G}_i^c$.

Now for $v \in V(G)$ let $B(v) = I_1(v) \times \cdots \times I_t(v)$ be the Cartesian product of the t intervals associated with vertex v. As v is in G_i for at most k indices $i \in [t]$, $I_i(v) \neq \mathbb{R}$ for at most k indices $i \in [t]$. In other words, $B(v)$ is a k-local box. Finally, we claim that $\{B(v) \mid v \in V(H)\}$ is an intersection representation of H. Indeed, if $vw \notin E(H)$, then $vw \in E(H^c)$ and hence $vw \in E(G_i)$ for at least one $i \in [t]$. Then $I_i(v) \cap I_i(w) = \varnothing$ and thus $B(v) \cap B(w) = \varnothing$. And if $vw \in E(H)$, then $vw \notin E(H^c)$ and $vw \notin E(G_i')$ for every $i \in [t]$. Thus $I_i(v) \cap I_i(w) \neq \varnothing$ for every $i \in [t]$ and hence $B(v) \cap B(w) \neq \varnothing$.

This shows that if $\text{box}_\ell(H) \leq k$, then H is the intersection graph of k-local boxes. On the other hand, if H admits an intersection representation with k-local t-dimensional boxes, then for each $i \in [t]$ projecting the boxes to coordinate i and considering the bounded intervals in this projection gives an interval representation of some subgraph G_i of H^c. As before, we can check that $\{G_1, \ldots, G_t\}$ forms an injective k-local $\mathcal{C}$-cover of H^c, showing that $\text{box}_\ell(H) = c_\ell^{\mathcal{C}}(H^c) \leq k$. □

5 Conclusions

In this paper we have introduced the notions of the local boxicity $\mathrm{box}_\ell(H)$ and union boxicity $\overline{\mathrm{box}}(H)$ of a graph H. It holds that $\mathrm{box}_\ell(H) \leq \overline{\mathrm{box}}(H) \leq \mathrm{box}(H)$, where $\mathrm{box}(H)$ denotes the classical boxicity as introduced almost 50 years ago. Indeed, both new parameters are a better measure of the complexity of H. For example, if H is the complement of a matching on n edges, then $\mathrm{box}(H) = n$, simply because the n non-edges each have to be realized in a different dimension. On the other hand, we have $\mathrm{box}_\ell(H) = \overline{\mathrm{box}}(H) = 1$, and as these non-edges are vertex-disjoint, they also should be "counted only once". We have shown this phenomenon in a few more examples in the course of the paper. In fact, in many box representations from the literature many (if not all) dimensions are only used by few vertices. The resulting high boxicity may be misintepreted as the graph being very complex, which could be avoided by using local or union boxicity.

In future research, established boxicity results should be revisited to see whether one can improve the upper bounds using local or union boxicity. For example, it is known that if H is a planar graph, then $\mathrm{box}(H) \leq 3$ [19]. Moreover, the octahedral graph O is planar and has boxicity 3, because its complement O^c is the matching on three edges (c.f. the proof of Theorem 5 (iii) and Figure 2). By (2) we have that $\mathrm{box}_\ell(H) \leq \overline{\mathrm{box}}(H) \leq 3$ whenever H is planar. However, $\mathrm{box}_\ell(O) = \overline{\mathrm{box}}(O) = 1$, because O^c is the vertex-disjoint union of co-interval graphs, i.e., $O^c \in \overline{\mathcal{C}}$. Hence it is natural to ask the following.

Question 11. *Is there a planar graph H with $\mathrm{box}_\ell(H) = 3$?*

For general graphs H we proved that the local boxicity $\mathrm{box}_\ell(H)$ and the union boxicity $\overline{\mathrm{box}}(H)$ can be arbitrarily far from the classical boxicity $\mathrm{box}(H)$. But we do not know whether if $\mathrm{box}(H)$ is large, then $\mathrm{box}_\ell(H)$ and $\overline{\mathrm{box}}(H)$ can be very *close* to $\mathrm{box}(H)$. We construct graphs in the proof of Theorem 5 (i) with large local boxicity, but one can show that these have even larger boxicity.

Question 12. *Is there for every $k \in \mathbb{N}$ a graph H_k such that $\mathrm{box}_\ell(H_k) = \overline{\mathrm{box}}(H_k) = \mathrm{box}(H_k) = k$?*

Another interesting research direction concerns the computational complexity. It is known that for every $k \geq 2$ deciding whether a given graph H satisfies $\mathrm{box}(H) \leq k$ is NP-complete [3, 12]. For $k = 1$ we have $\mathrm{box}(H) \leq k$ if and only if H is an interval graph, and $\overline{\mathrm{box}}(H) \leq k$ (equivalently $\mathrm{box}_\ell(H) \leq k$) if and only if the complement of H is the vertex-disjoint union of co-interval graphs, both of which can be tested in polynomial time via interval graph recognition [2].

Question 13. *For $k \geq 2$, is it NP-complete to decide whether $\mathrm{box}_\ell(H) \leq k$ (or $\overline{\mathrm{box}}(H) \leq k$) for a given graph H?*

Let us remark that for general covering numbers the computational complexity of computing $c_g^{\mathcal{G}}(H)$ tends to be harder than that of $c_\ell^{\mathcal{G}}(H)$, which in turn tends to be harder than for $c_f^{\mathcal{G}}(H)$. For example, for $\mathcal{G}$ being the class of star forests, computing $c_g^{\mathcal{G}}(H)$ is NP-complete [8, 10], while computing $c_\ell^{\mathcal{G}}(H)$ and $c_f^{\mathcal{G}}(H)$ is polynomial-time solvable [11]. The same holds when $\mathcal{G}$ is the class of all matchings as discussed in [11]. And for $\mathcal{G}$ being the class of bipartite graphs, computing $c_g^{\mathcal{G}}(H)$ and $c_\ell^{\mathcal{G}}(H)$ is NP-complete [7], while computing $c_f^{\mathcal{G}}(H)$ is polynomial-time solvable since $c_f^{\mathcal{G}}(H) = 1$ if H is bipartite and $c_f^{\mathcal{G}}(H) = 2$ otherwise.

References

[1] Lowell W. Beineke, *A survey of packings and coverings of graphs*, The Many Facets of Graph Theory (Proc. Conf., Western Mich. Univ., Kalamazoo, Mich., 1968), Springer, Berlin, 1969, pp. 45–53.

[2] Kellogg S. Booth and George S. Lueker, *Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms*, J. Comput. System Sci. **13** (1976), no. 3, 335–379, Working Papers presented at the ACM-SIGACT Symposium on the Theory of Computing (Albuquerque, N. M., 1975).

[3] Margaret B Cozzens, *Higher and multi-dimensional analogues of interval graphs*, Dissertation Abstracts International Part B: Science and Engineering **42** (1981), no. 4.

[4] Margaret B. Cozzens and Fred S. Roberts, *Computing the boxicity of a graph by covering its complement by cointerval graphs*, Discrete Applied Mathematics **6** (1983), no. 3, 217–228.

[5] Vida Dujmovic, Pat Morin, and David R Wood, *Layout of graphs with bounded tree-width*, SIAM Journal on Computing **34** (2005), no. 3, 553–579.

[6] Louis Esperet and Gwenaël Joret, *Boxicity of graphs on surfaces*, Graphs and combinatorics **29** (2013), no. 3, 417–427.

[7] Peter C. Fishburn and Peter L. Hammer, *Bipartite dimensions and bipartite degrees of graphs*, Discrete Math. **160** (1996), no. 1-3, 127–148.

[8] Daniel Gonçalves and Pascal Ochem, *On star and caterpillar arboricity*, Discrete Math. **309** (2009), no. 11, 3694–3702.

[9] András Gyárfás and Douglas B. West, *Multitrack interval graphs*, Proceedings of the Twenty-sixth Southeastern International Conference on Combinatorics, Graph Theory and Computing (Boca Raton, FL, 1995), vol. 109, 1995, pp. 109–116.

[10] S. Louis Hakimi, John. Mitchem, and Edward F. Schmeichel, *Star arboricity of graphs*, Discrete Math. **149** (1996), no. 1-3, 93–98.

[11] Kolja Knauer and Torsten Ueckerdt, *Three ways to cover a graph*, Discrete Mathematics **339** (2016), no. 2, 745–758.

[12] Jan Kratochvíl, *A special planar satisfiability problem and a consequence of its np-completeness*, Discrete Applied Mathematics **52** (1994), no. 3, 233–252.

[13] Kevin G Milans, Derrick Stolee, and Douglas B West, *Ordered Ramsey theory and track representations of graphs*, Journal of Combinatorics **6** (2015), no. 4, 445–456.

[14] Petra Mutzel, Thomas Odenthal, and Mark Scharbrodt, *The thickness of graphs: a survey*, Graphs Combin. **14** (1998), no. 1, 59–73.

[15] Crispin St. J. A. Nash-Williams, *Decomposition of finite graphs into forests*, J. London Math. Soc. **39** (1964), 12.

[16] Frank P. Ramsey, *On a problem of formal logic*, Proceedings of the London Mathematical Society **s2-30** (1930), no. 1, 264–286.

[17] Fred S. Roberts, *On the boxicity and cubicity of a graph*, Recent Progresses in Combinatorics, Academic Press, 1969, pp. 301–310.

[18] Edward R Scheinerman, *Intersection classes and multiple intersection parameters of graphs*, Ph.D. thesis, Princeton University, 1984.

[19] Carsten Thomassen, *Interval representations of planar graphs*, Journal of Combinatorial Theory, Series B **40** (1986), no. 1, 9–20.

Discrete Mathematics 339 (2016) 745–758

Contents lists available at ScienceDirect

Discrete Mathematics

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

Perspective

Three ways to cover a graph

Kolja Knauer [a,*], Torsten Ueckerdt [b]

[a] *Aix-Marseille Université, CNRS, LIF UMR 7279, 13288, Marseille, France*
[b] *Karlsruhe Institute of Technology, Karlsruhe, Germany*

ARTICLE INFO

Article history:
Received 3 January 2013
Received in revised form 2 October 2015
Accepted 5 October 2015
Available online 11 November 2015

Keywords:
Graph covering
Interval number
Track number
Linear arboricity
Caterpillar arboricity
Star arboricity

ABSTRACT

We consider the problem of covering an *input graph* H with graphs from a fixed *covering class* $\mathcal{G}$. The classical covering number of H with respect to $\mathcal{G}$ is the minimum number of graphs from $\mathcal{G}$ needed to cover the edges of H without covering non-edges of H. We introduce a unifying notion of three covering parameters with respect to $\mathcal{G}$, two of which are novel concepts only considered in special cases before: the local and the folded covering number. Each parameter measures "how far" H is from $\mathcal{G}$ in a different way. Whereas the folded covering number has been investigated thoroughly for some covering classes, e.g., interval graphs and planar graphs, the local covering number has received little attention.

We provide new bounds on each covering number with respect to the following covering classes: linear forests, star forests, caterpillar forests, and interval graphs. The classical graph parameters that result this way are interval number, track number, linear arboricity, star arboricity, and caterpillar arboricity. As input graphs we consider graphs of bounded degeneracy, bounded degree, bounded tree-width or bounded simple tree-width, as well as outerplanar, planar bipartite, and planar graphs. For several pairs of an input class and a covering class we determine exactly the maximum ordinary, local, and folded covering number of an input graph with respect to that covering class.

1. Introduction

Graph covering is one of the most classical topics in graph theory. In 1891, in one of the first purely graph-theoretical papers, Petersen [47] showed that any $2r$-regular graph can be covered with r sets of vertex disjoint cycles. A survey on covering problems by Beineke [11] appeared in 1969. Graph covering is a lively field with deep ramifications — over the last decades as well as today [29,30,2,3,25,45]. This is supported through the course of this paper by many references to recent works of different authors.

In every graph covering problem one is given an input graph H, a covering class $\mathcal{G}$, and a notion of how to cover H with one or several graphs from $\mathcal{G}$. One is then interested in $\mathcal{G}$-coverings of H that are in some sense simple, or well structured; the most prevalent measure of simplicity being the number of graphs from $\mathcal{G}$ needed to cover the edges of H.

The main goal of this paper is to introduce the following three parameters, each of which represents how well H can be covered with respect to $\mathcal{G}$ in a different way:

The *global covering number*, or simply *covering number*, is the most classical one. It is the smallest number of graphs from $\mathcal{G}$ needed to cover the edges of H without covering non-edges of H. All kinds of arboricities, e.g. star [4], caterpillar [24], linear [3], pseudo [48], and ordinary [46] arboricity of a graph are global covering numbers, where the covering class is the

* Corresponding author.
E-mail addresses: kolja.knauer@lif.univ-mrs.fr (K. Knauer), torsten.ueckerdt@kit.edu (T. Ueckerdt).

http://dx.doi.org/10.1016/j.disc.2015.10.023

Table 1
Overview of results. (See [29,42,52,23,15,17,6,5,24].)

	Star forests		Caterpillar forests			Interval graphs		
	g	$\ell = f$	g	ℓ	f	g	ℓ	f
Outer-planar	3 [29]	3	3 [42]	3	3	2 [42]	2	2 [52]
Planar bipartite	4 (C18)	3 (C18)	4 [23]	3	3	4	3	3 [52]
Planar	5 [5, 29]	4 (C18)	4 [23]	4	4 [52]	4 [24]	?	3 [52]
stw $\leq k$	$k+1$	$k+1$	$k+1$	$k+1$	$k+1$ (T17)	$k+1$ (T16)	k (T14)	k (T15)
tw $\leq k$	$k+1$ [15, 17]	$k+1$	$k+1$	$k+1$	$k+1$	$k+1$	$k+1$	$k+1$ (T15)
dgn $\leq k$	$2k$ [6]	$k+1$ (C10)	$2k$	$k+1$	$k+1$	$2k$ (T12)	$k+1$	$k+1$

class of star forests, caterpillar forests, linear forests, pseudoforests, and ordinary forests, respectively. Other global covering numbers are the planar and outerplanar thickness [11,45] and the track number [28] of a graph. Here, the covering classes are planar, outerplanar, and interval graphs, respectively.

In the *local covering number* of H with covering class $\mathcal{G}$ one also tries to cover the edges of H with graphs from $\mathcal{G}$ but now minimizes the largest number of graphs in the covering containing a common vertex of H. We are aware of only two local covering numbers in the literature: The bipartite degree introduced by Fishburn and Hammer [21] is the local covering number where the covering class is the class of complete bipartite graphs. It was rediscovered by Dong and Liu [16] as the local biclique cover number, and recently it has been studied in comparison with its global variant by Pinto [49]. The local clique cover number is another local covering number, where the covering class is the class of complete graphs. It was studied by Skums, Suzdal, and Tyshkevich [55] and by Javadi, Maleki, and Omoomi [36].

Finally, the *folded covering number* underlies a different, but related, concept of covering. Here, one looks for a graph in $\mathcal{G}$ which has H as homomorphic image and one minimizes the size of the largest preimage of a vertex of H. Equivalently, one splits every vertex of H into a independent set such that the size of the largest such independent set is minimized, distributing the incident edges to the new vertices, such that the result is a graph from $\mathcal{G}$. The folded covering number has been investigated using interval graphs and planar graphs as covering class. In the former case the folded covering number is known as the interval number [31], in the latter case as the splitting-number [35].

While some covering numbers, like arboricities, are of mainly theoretical interest, others, like thickness, interval number, and track number, have wide applications in VLSI design [1], network design [50], scheduling and resource allocation [10,13], and bioinformatics [39,37]. The three covering numbers presented here not only unify some notions in the literature, they as well seem interesting in their own right and may provide new approaches to attack classical open problems.

In this paper we moreover present new lower and upper bounds for several covering numbers. In the new results, the covering classes are: interval graphs, star forests, linear forests, and caterpillar forests. The input classes are: graphs of bounded degeneracy, bounded tree-width or bounded simple tree-width, as well as outerplanar, planar bipartite, planar, and regular graphs. Not all pairs of these input classes with these covering classes are given new bounds. We provide an overview over some of our new results in Table 1. Each row of the table corresponds to an input class $\mathcal{H}$, each column to a covering class $\mathcal{G}$. Every cell contains the maximum covering number among all graphs $H \in \mathcal{H}$ with respect to the covering class $\mathcal{G}$, where the columns labeled g, ℓ, f stand for the global, local, and folded covering number, respectively. Gray entries follow by Proposition 4 from other stronger results in the table. Letters T and C stand for Theorem and Corollary in the present paper, respectively. Indeed all the entries except the '?' in Table 1 are exact, with matching upper and lower bounds. Note that besides results we prove as new theorems as indicated, many values in the table (written in gray) follow from the point of view offered by our general approach (Proposition 4).

This paper is structured as follows: In order to give a motivating example before the general definition, we start by discussing in Section 2 the linear arboricity and its local and folded variants. In Section 3 the three covering numbers are formally introduced and some general properties are established. In Section 4 we introduce the covering classes star forests, caterpillar forests, and interval graphs, and in Section 5 we present our results claimed in Table 1. In Section 6 we briefly discuss the computational complexity of some covering numbers, giving a polynomial-time algorithm for the local star arboricity. Moreover, we discuss by how much global, local and folded covering numbers can differ.

For the entire paper we assume all graphs to be simple without loops nor multiple edges. Notions used but not introduced can be found in any standard graph theory book; such as [57].

2. Global, local, and folded linear arboricity

We give the general definitions of covers and covering numbers in Section 3. In this section we motivate and illustrate these concepts on the basis of one fixed covering class: the class $\mathcal{L}$ of *linear forests*, which are the disjoint unions of paths.

We want to cover an input graph H by several linear forests $L_1, \dots, L_k \in \mathcal{L}$. That is, every edge $e \in E(H)$ is contained in at least[1] one L_i and no non-edge of H is contained in any L_i. When H is covered by $L_1, \dots, L_k$ we write $H = \bigcup_{i\in[k]} L_i$.

The *linear arboricity of H*, denoted by $\mathrm{la}(H)$, is the minimum k such that $H = \bigcup_{i\in[k]} L_i$ and $L_i \in \mathcal{L}$ for $i \in [k]$. One easily sees that every graph H of maximum degree $\Delta(H)$ has $\mathrm{la}(H) \geq \left\lceil \frac{\Delta(H)}{2} \right\rceil$, and every $\Delta(H)$-regular graph H has $\mathrm{la}(H) \geq \left\lceil \frac{\Delta(H)+1}{2} \right\rceil$. In 1980, Akiyama et al. [3] stated the Linear Arboricity Conjecture (LAC). It says that the linear arboricity of any simple graph H of maximum degree $\Delta(H)$ is either $\left\lceil \frac{\Delta(H)}{2} \right\rceil$ or $\left\lceil \frac{\Delta(H)+1}{2} \right\rceil$. LAC was confirmed for planar graphs by Wu and Wu [59,60] and asymptotically for general graphs by Alon and Spencer [7]. The general conjecture remains open. The best-known general upper bound for $\mathrm{la}(H)$ is $\left\lceil \frac{3\Delta(H)+2}{5} \right\rceil$, due to Guldan [27].

We define the *local linear arboricity of H*, denoted by $\mathrm{la}_\ell(H)$, as the minimum j such that $H = \bigcup_{i\in[k]} L_i$ for some k and every vertex v in H is contained in at most j different L_i. Again, if H has maximum degree $\Delta(H)$, then $\mathrm{la}_\ell(H) \geq \left\lceil \frac{\Delta(H)}{2} \right\rceil$, and if H is $\Delta(H)$-regular, then $\mathrm{la}_\ell(H) \geq \left\lceil \frac{\Delta(H)+1}{2} \right\rceil$. Note that $\mathrm{la}_\ell(H)$ is at most $\mathrm{la}(H)$, and hence the following statement must necessarily hold for LAC to be true.

Conjecture 1. *Local Linear Arboricity Conjecture (LLAC): The local linear arboricity of any simple graph with H maximum degree $\Delta(H)$ is either $\left\lceil \frac{\Delta(H)}{2} \right\rceil$ or $\left\lceil \frac{\Delta(H)+1}{2} \right\rceil$.*

Observation 2. *To prove LAC or LLAC it suffices to consider regular graphs of odd degree: Regularity is obtained by considering a $\Delta(H)$-regular supergraph of H. If $\Delta(H)$ is even, say $\Delta(H) = 2k$, one can find a spanning linear forest L_{k+1} in H* [27], *remove it from the graph, and extend L_{k+1} by a cover $L_1, \dots, L_k$ in the remaining graph of maximum degree $\Delta(H) - 1 = 2k - 1$.*

If H is regular with odd degree, then LLAC states that $H = \bigcup_{i\in[k]} L_i$ with every vertex being an endpoint of exactly one path. LAC additionally requires that the paths can be colored with $\left\lceil \frac{\Delta(H)}{2} \right\rceil$ colors such that no two paths that share a vertex receive the same color. We will see in later sections that sometimes the coloring is the crucial and difficult task.

Next we propose a second way to cover the input graph H with linear forests. A *walk* in H is a sequence of consecutively incident edges of H of the form $\{v_1v_2, v_2v_3, \dots, v_{k-1}v_k\}$ for $v_1, \dots v_k$ being vertices of H. As before, a set $W_1, \dots, W_k$ of walks covers H, denoted by $H = \bigcup_{i\in[k]} W_i$, if the edge-set E of H is the union of the edge-sets of the walks. We are now interested in how often a vertex v in H appears in the walks $W_1, \dots, W_k$ in total. The *folded linear arboricity of H*, denoted by $\mathrm{la}_f(H)$, is the minimum j such that $H = \bigcup_{i\in[k]} W_i$ and every vertex v in H appears at most j times in the walks $W_1, \dots, W_k$. Again if H has maximum degree $\Delta(H)$ then $\mathrm{la}_f(H) \geq \left\lceil \frac{\Delta(H)}{2} \right\rceil$, and if H is $\Delta(H)$-regular then $\mathrm{la}_f(H) \geq \left\lceil \frac{\Delta(H)+1}{2} \right\rceil$. Clearly, $\mathrm{la}_f(H) \leq \mathrm{la}_\ell(H)$. The next theorem follows directly from a short proof of West [56] of a result previously published by Griggs and West [26] (where it is stated in terms of the interval number $i(H)$). It is a weakening of LLAC above.

Theorem 3. *If H has maximum degree $\Delta(H)$ then $\mathrm{la}_f(H) \in \left\{ \left\lceil \frac{\Delta(H)}{2} \right\rceil, \left\lceil \frac{\Delta(H)+1}{2} \right\rceil \right\}$.*

Proof. Add a vertex x to H and connect it to every vertex in H of odd degree. Each component of the resulting graph is Eulerian. Consider any Eulerian tour in $H \cup x$ (or H) and split it into shorter walks by removing x from it. □

3. Covers and covering numbers

In this section we formalize the concepts from Section 2 with respect to general covering and input classes and obtain some general inequalities. The notation we introduce is convenient for making our generalized approach as transparent as possible. When treating concrete covering classes, for which covering numbers already have an established notation in the literature later on in the paper, we will use the latter in order to make results more accessible to readers already familiar with the parameters.

A *homomorphism* from a graph G to a graph H is a map $\varphi : V(G) \to V(H)$ such that $vw \in E(G)$ implies $\varphi(v)\varphi(w) \in E(H)$. We call a homomorphism *edge-surjective* if for all $v'w' \in E(H)$ there exists $vw \in E(G)$ such that $\varphi(v) = v'$ and $\varphi(w) = w'$. For an input graph H and a covering class $\mathcal{G}$, we define a *$\mathcal{G}$-cover of H* as an edge-surjective homomorphism $\varphi : G_1 \dot\cup G_2 \dot\cup \cdots \dot\cup G_k \to H$, where $G_i \in \mathcal{G}$ for $i \in [k]$ and $\dot\cup$ denotes the vertex disjoint union. The *size* of a cover is the number of covering graphs in the disjoint union. A cover φ is called *injective* if $\varphi|G_i$, that is, φ restricted to G_i, is injective for every $i \in [k]$.

Definition 1. For a covering class $\mathcal{G}$ and an input graph H define the *(global) covering number* $c_g^{\mathcal{G}}(H)$, the *local covering number* $c_\ell^{\mathcal{G}}(H)$, and the *folded covering number* $c_f^{\mathcal{G}}(H)$ as follows:

[1] Since linear forests are closed under taking subgraphs, we can indeed assume that $e \in L_i$ for exactly one $i \in [k]$.

$c_g^{\mathcal{G}}(H) = \min \{\text{size of } \varphi : \varphi \text{ is an injective } \mathcal{G}\text{-cover of } H\}$
$c_\ell^{\mathcal{G}}(H) = \min \{\max_{v\in V(H)} |\varphi^{-1}(v)| : \varphi \text{ is an injective } \mathcal{G}\text{-cover of } H\}$
$c_f^{\mathcal{G}}(H) = \min \{\max_{v\in V(H)} |\varphi^{-1}(v)| : \varphi \text{ is a } \mathcal{G}\text{-cover of } H \text{ having size } 1\}.$

Let us rephrase $c_g^{\mathcal{G}}(H)$, $c_\ell^{\mathcal{G}}(H)$, and $c_f^{\mathcal{G}}(H)$. The covering number is the minimum number of graphs in $\mathcal{G}$ needed to cover H exactly, where *covering exactly* means identifying subgraphs in H that are covering graphs, such that every edge of H is contained in some covering graph. In the local covering number the number of covering graphs is not restricted; instead the number of covering graphs at every vertex should be small. We will see later that these two numbers can differ significantly. The folded covering number is the minimum k such that every vertex v of H can be split into at most k vertices, distributing the incident edges at v arbitrarily (even repeatedly) among them, such that the resulting graph belongs to $\mathcal{G}$. The splitting corresponds to representing the vertex by the set of its preimages under the edge-surjective homomorphism φ.

One is often interested in the maximum or minimum value of a graph parameter on a class of input graphs. For $i \in \{g, \ell, f\}$, a covering class $\mathcal{G}$, and an input graph class $\mathcal{H}$, we define $c_i^{\mathcal{G}}(\mathcal{H}) = \sup \left\{c_i^{\mathcal{G}}(H) : H \in \mathcal{H}\right\}$. We close this section with a list of inequalities, most of which are elementary applications of Definition 1 and homomorphisms.

Proposition 4. *For covering classes $\mathcal{G}$, $\mathcal{G}'$, input classes $\mathcal{H}$, $\mathcal{H}'$ and any input graph H we have the following:*

(i) *$c_g^{\mathcal{G}}(H) \geq c_\ell^{\mathcal{G}}(H)$, and if $\mathcal{G}$ is closed under disjoint union, then $c_\ell^{\mathcal{G}}(H) \geq c_f^{\mathcal{G}}(H)$.*
(ii) *If $\mathcal{G}$ is closed under merging non-adjacent vertices within connected components (and afterwards deleting multiple edges) and restriction to maximal connected components, then $c_\ell^{\mathcal{G}}(H) \leq c_f^{\mathcal{G}}(H)$.*
(iii) *If $\mathcal{H} \subseteq \mathcal{H}'$, then $c_i^{\mathcal{G}}(\mathcal{H}) \leq c_i^{\mathcal{G}}(\mathcal{H}')$ for $i \in \{g, \ell, f\}$.*
(iv) *If $H_{\mathcal{G}}$ and $H_{\mathcal{G}'}$ denote the set of subgraphs of H that are homomorphic images of graphs in $\mathcal{G}$ and $\mathcal{G}'$, respectively, then $H_{\mathcal{G}} \subseteq H_{\mathcal{G}'}$ implies $c_i^{\mathcal{G}}(H) \geq c_i^{\mathcal{G}'}(H)$ for $i \in \{g, \ell, f\}$. This holds in particular when $\mathcal{G} \subseteq \mathcal{G}'$.*
(v) *If $\bar{H}$ denotes the set of all subgraphs of H and we have $\mathcal{G} \cap \bar{H} \subseteq \mathcal{G}' \cap \bar{H}$, then $c_i^{\mathcal{G}}(H) \geq c_i^{\mathcal{G}'}(H)$ for $i \in \{g, \ell\}$.*

Proof. The first inequality in (i) follows from the definition, the second one comes by viewing an injective cover $G_1 \uplus G_2 \uplus \cdots \uplus G_k$ as a $\mathcal{G}$-cover of size 1.

To see (ii), let $\varphi : G \to H$ be a $\mathcal{G}$-cover of H of size 1 witnessing $c_f^{\mathcal{G}}(H)$. Now for every $v \in H$ and a component G' of G merge all $\varphi^{-1}(v) \cap V(G')$ into one vertex (and delete multiple edges). Since H has no loops, the merging process creates no loops. Doing this for all components of G yields a new covering graph $\widetilde{G} \in \mathcal{G}$ with homomorphism $\widetilde{\varphi}$ being injective on each component. Clearly, $|\widetilde{\varphi}^{-1}(v)| \leq |\varphi^{-1}(v)|$.

Claims (iii) and (iv) follow immediately from the definition. To see (v) note that it follows similarly as (iv), because $\mathcal{G} \cap \bar{H}$ and $\mathcal{G}' \cap \bar{H}$ are the subgraphs of H that arise as images of *injective* covers. □

Remark 5. Within the scope of this paper we only consider covering classes that are closed under disjoint union even without explicitly saying so. For example, when considering stars or complete graphs as covering graphs, we actually mean star forests and disjoint unions of complete graphs, respectively. If the covering class $\mathcal{G}$ is closed under disjoint union, then the restriction to covers of size 1 in the definition of $c_f^{\mathcal{G}}$ is unnecessary.

It is still interesting to consider covering classes that are not closed under disjoint union. Hajós' Conjecture [43] states that the edges of any n-vertex Eulerian graph H may be partitioned into $\lfloor \frac{n}{2} \rfloor$ cycles. Hajós' Conjecture being widely open, one may consider coverings with cycles. When $\mathcal{C}'$ denotes the class of all simple cycles and H is an n-vertex Eulerian graph, Fan [19] proved $c_g^{\mathcal{C}'}(H) \leq \lfloor \frac{n-1}{2} \rfloor$.

Example

In order to illustrate the notions introduced above, consider the covering class $\mathcal{C}$ of disjoint unions of cycles. As input graph H we take the Petersen graph. See Fig. 1 where we have from left to right: A global cover with three unions of cycles, a local cover of size five with at most three cycles at each vertex, and a folded cover with two preimages per vertex. Note that the local cover does not yield an optimal global cover.

Proposition 6. *For the Petersen graph, we have $3 = c_g^{\mathcal{C}}(H) = c_\ell^{\mathcal{C}}(H) > c_f^{\mathcal{C}}(H) = 2$.*

Proof. All witnesses for the upper bounds are shown in Fig. 1. Clearly, $c_f^{\mathcal{C}}(H) \geq 2$ since otherwise H would have to be a disjoint union of cycles. Now suppose, $c_\ell^{\mathcal{C}}(H) = 2$. Since H is cubic, at each vertex there is exactly one edge contained in two cycles of the covering. Thus, these edges form a perfect matching M of H. Moreover, all cycles involved in the cover are alternating cycles with respect to M. In particular they are all even and of length 6 or 8 (as this graph is not Hamiltonian there is no 10-cycle). Since M is covered twice and the remaining edges of H once, the sum of sizes of cycles in the cover is 20, which can be obtained only as $6 + 6 + 8$. In particular, a 6-cycle C must be involved. Now M restricted to $H \backslash V(C)$ is still a perfect matching, but $H \setminus V(C)$ is a claw. □

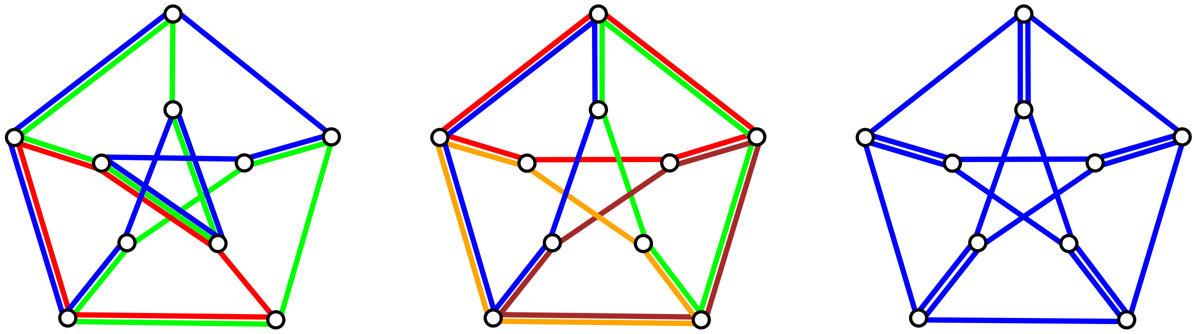

Fig. 1. Coverings of the Petersen graph by disjoint unions of cycles.

4. Covering classes

In this section we introduce the covering classes and covering numbers corresponding to the columns of Table 1. We also include some known results and general observations.

4.1. Forests and pseudoforests

Nash-Williams [46] showed that the minimum number of forests needed to cover the edges of H is $\max_{S\subseteq V(H)} \left\lceil \frac{|E[S]|}{|S|-1} \right\rceil$, where $E[S]$ denotes the set of edges in the subgraph induced by S. This value, denoted by $a(H)$, is now usually called the *arboricity* of H, see Beineke [11] for an early appearance of this name. Clearly, $a(H) = c_g^{\mathcal{G}}(H)$, where $\mathcal{G}$ is the class of forests.

A *pseudoforest* is a graph with at most one cycle per component and the *pseudoarboricity* $p(H)$ is the minimum number of pseudoforests needed to cover the edges of H. Thus, $p(H) = c_g^{\mathcal{G}}(H)$, where $\mathcal{G}$ is the class of pseudoforests. Results of Picard and Queyranne [48] and Frank and Gyárfás [22] yield the following lemma.

Lemma 7 *([22,48]). The pseudoarboricity $p(H)$ of a graph H equals the minimum over all orientations of H of the maximum out-degree of H. Furthermore, $p(H) = \max_{S\subseteq V(H)} \left\lceil \frac{|E[S]|}{|S|} \right\rceil$.*

Using $a(H) = \max_{S\subseteq V(H)} \left\lceil \frac{|E[S]|}{|S|-1} \right\rceil$, one immediate consequence of Lemma 7 is $p(H) \leq a(H) \leq p(H) + 1$.

Theorem 8. *For every graph, the values of global, local, and folded (pseudo)arboricity coincide.*

Proof. Take a folded covering φ of H with a (pseudo)forest, such that for every $v \in H$ we have $|\varphi^{-1}(v)| \leq c$. Since (pseudo)forests are closed under taking induced subgraphs, this in particular yields a covering for every induced subgraph $H[S]$ such that every vertex is covered at most c times. Now, focusing on pseudoforests, we know that the subgraph of the covering graph induced by $\varphi^{-1}(S)$ has at most $c|S|$ edges, and therefore $c|S| \geq |E[S]|$, i.e., $c \geq \left\lceil \frac{|E[S]|}{|S|} \right\rceil$. Now by Lemma 7, we have $p(H) = \max_{S\subseteq V(H)} \left\lceil \frac{|E[S]|}{|S|} \right\rceil$ yielding the result for folded coverings. Now, Proposition 4(i) gives the result for the local covering number.

Along the same lines one obtains $c \geq \left\lceil \frac{|E[S]|+1}{|S|} \right\rceil$ when c is the number of times a vertex is covered in a forest-cover of H. It is then easy to compute $\left\lceil \frac{|E[S]|+1}{|S|} \right\rceil = \left\lceil \frac{|E[S]|}{|S|-1} \right\rceil$, since $|E[S]| \leq \binom{|S|}{2}$. The result follows as in the case of pseudoarboricity. □

4.2. Star forests

The star arboricity $\mathrm{sa}(H)$ of a graph H, introduced by Akiyama and Kano [4], is the minimum number of star forests (forests without paths of length 3) into which the edge-set of H can be partitioned. In particular, if $\mathcal{S}$ denotes the class of star forests, then $\mathrm{sa}(H) = c_g^{\mathcal{S}}(H)$. The star arboricity has been a frequent subject of research. It is known that outerplanar and planar graphs have star arboricity at most 3 and 5, respectively; see Hakimi et al. [42]. That this is best possible was shown by Algor and Alon [5]. Alon et al. [6] showed that $sa(H) \leq 2a(H)$ is a tight upper bound.

Since merging non-adjacent vertices in a star and omitting double edges yields again a star, local and folded star arboricity coincide, by Proposition 4(ii). Here, we show that in contrast to the global star arboricity, the local star arboricity, denoted by $\mathrm{sa}_\ell(H)$, fits nicely into the inequalities relating arboricity and pseudoarboricity from Section 4.1.

Theorem 9. *For any graph H, we have $p(H) \leq a(H) \leq \mathrm{sa}_\ell(H) \leq p(H)+1$, where any inequality can be strict. Moreover, $\mathrm{sa}_\ell(H) = p(H)$ if and only if H has an orientation with maximum out-degree $p(H)$ in which this outdegree occurs only at vertices of degree $p(H)$.*

Proof. Every cover of H with respect to stars can be transformed into an orientation of H by orienting every edge towards the center of the corresponding star. If every vertex is contained in at most $\mathrm{sa}_\ell(H)$ stars, then the orientation has maximum out-degree at most $\mathrm{sa}_\ell(H)$. Lemma 7 then gives $p(H) \leq \mathrm{sa}_\ell(H)$.

In the same way, every orientation can be transferred into a cover with respect to stars by taking at every vertex the star of its incoming edges. If the orientation has maximum out-degree $p(H)$, then each vertex is contained in no more than $p(H)+1$ stars, i.e., $\mathrm{sa}_\ell(H) \leq p(H)+1$. Moreover, the maximum out-degree is $\mathrm{sa}_\ell(H)$ if and only if for every vertex v lying in $\mathrm{sa}_\ell(H)$ stars with centers different from v there is no star with center v. Equivalently, $\mathrm{sa}_\ell(H) = p(H)$ if and only if the maximum out-degree $p(H)$ is attained only at vertices of degree $p(H)$.

If $sa_l(H) = p(H)+1$, then $a(H) \leq sa_l(H)$ follows from $a(H) \leq p(H)-1$. When $sa_l(H) = p(H)$, there is an orientation with maximum out-degree $p(H)$ attained only at vertices with degree $p(H)$. Removing these vertices, we obtain a graph H' with $p(H') \leq p(H)-1$, in particular $a(H') \leq p(H)$. We reinsert the vertices of degree $p(H)$ putting each incident edge into a different one of the $p(H)$ forests that partition H'. We obtain a cover of H with $p(H)$ forests, so $a(H) \leq p(H) = \mathrm{sa}_\ell(H)$.

Finally, we show that each inequality can be strict: First $k = p(H) < a(H)$ holds for every $2k$-regular graph H, due to the number of edges of the covering graphs. Second, we claim that $k = p(H) = \mathrm{sa}_\ell(H)$ holds for the complete bipartite graph $K_{k,n}$ with n large enough. Indeed, $p(K_{k,n}) = \max_{S\subseteq V(K_{k,n})} \left\lceil \frac{|E[S]|}{|S|} \right\rceil = \left\lceil \frac{kn}{k+n} \right\rceil = k$, and taking all maximal stars with centers in the smaller class of the bipartition yields $\mathrm{sa}_\ell(K_{k,n}) \leq k$.

It remains to present a graph H with $k = a(H) < \mathrm{sa}_\ell(H)$. We take H to be the k-dimensional grid of size m. That is, $V(H) = [m]^k$, and there is an edge joining vertices v and w if and only if they differ in exactly one coordinate and differ there by 1. It is straightforward to compute that H has $(m-1)m^{k-1}k$ edges. Observing that H itself is a densest induced subgraph, the formulas for arboricity and pseudoarboricity give $a(H) = p(H) = k$ for large enough m. Also, $a(H) = \mathrm{sa}_\ell(H)$ implies $p(H) = \mathrm{sa}_\ell(H)$. Hence, as proved above, H has an orientation with maximum out-degree k, which furthermore is only attained at vertices of degree k. However, H has only 2^k vertices of degree k. If all other vertices have outdegree at most $k-1$, then H has at most $2^k k + (m^k - 2^k)(k-1)$ edges. Choosing $m > 2^k + k$ yields a contradiction to the number of edges of H calculated above. □

We will derive from Theorem 9 tight upper bounds for the local star arboricity in Section 5, as well as a polynomial-time algorithm to compute the local star arboricity in Section 6.

4.3. Other covering classes

A graph parameter related to the star arboricity is the *caterpillar arboricity* $\mathrm{ca}(H)$ of H. A *caterpillar* is a tree in which all non-leaf vertices form a path, called the *spine*. The caterpillar arboricity is the minimum number of caterpillar forests into which the edge-set of H can be partitioned. It has mainly been considered for outerplanar graphs (Kostochka and West [41]), and for planar graphs (Gonçalves and Ochem [23,24]).

The class $\mathcal{I}$ of *interval graphs* has already been considered in many ways and remains present in today's literature. Interval graphs have been generalized to intersection graphs of systems of intervals by several groups of people: Gyárfás and West [28] proposed the $\mathcal{I}$-covering and introduced the corresponding global covering number called the *track number*, denoted by $t(H)$, i.e., $t(H) = c_g^{\mathcal{I}}(H)$. It has been shown that outerplanar and planar graphs have track number at most 2 [41] and 4 [24], respectively. Already in 1979, Harary and Trotter [31] introduced the folded $\mathcal{I}$-covering number, called the *interval number*, denoted by $i(H)$, i.e., $i(H) = c_f^{\mathcal{I}}(H)$. It is known that trees have interval number at most 2 [31]. Also, outerplanar and planar graphs have interval number at most 2 and 3, respectively, see Scheinerman and West [52]. All these bounds are tight.

The *local track number* $t_\ell(H) := c_\ell^{\mathcal{I}}(H)$ is a natural variation of $i(H)$ and $t(H)$, which to our knowledge has not been considered so far.

5. Results

In this section we present all the new results displayed in Table 1. We proceed input class by input class.

5.1. Bounded degeneracy

The *degeneracy* $\mathrm{dgn}(H)$ of a graph H is the minimum of the maximum out-degree over all acyclic orientations of H. It is a classical measure for the sparsity of H. By Lemma 7 and the definition we have $p(H) \leq a(H) \leq \mathrm{dgn}(H)$. Thus, the next corollary follows directly from Theorem 9.

Corollary 10. *For every H we have $\mathrm{sa}_\ell(H) \leq \mathrm{dgn}(H)+1$.*

Let $\mathcal{I}$ be the class of interval graphs and $\mathcal{C}a$ be the class of caterpillar forests, i.e., the class of bipartite interval graphs. Since homomorphisms the image of a homomorphism has chromatic number at least as large as its preimage, the chromatic number of an interval graph G that has a bipartite homomorphic image is at most two. Thus, G is a caterpillar forest. Therefore, when G is bipartite, the set of all homomorphic images of caterpillar forests in G coincides with the set of all homomorphic images of interval graphs in G. Thus, by Proposition 4(iv) we have $c_i^{\mathcal{I}}(H) = c_i^{\mathcal{C}a}(H)$ for $i \in \{g, \ell, f\}$ for every bipartite graph H. In particular, if H is bipartite then $t(H) = \mathrm{ca}(H)$ and $i(H) = \mathrm{ca}_f(H)$. In the remainder of this section we present graphs with high (folded) caterpillar arboricity. Since all these graphs are bipartite, we obtain lower bounds on the track number and interval number of those graphs. Indeed in all constructions we define a supergraph H of the complete bipartite graph $K_{m,n}$. The track number and interval number of $K_{m,n}$ have already been determined: $t(K_{m,n}) = \mathrm{ca}(K_{m,n}) = \left\lceil \frac{mn}{m+n-1} \right\rceil$ [28] and $i(K_{m,n}) = \mathrm{ca}_f(K_{m,n}) = \left\lceil \frac{mn+1}{m+n} \right\rceil$ [31].

In order to formulate the following lemma, we need to introduce one more notion. For a cover φ of H by $G_1 \dot\cup \dots \dot\cup G_k$ with $G_i \in \mathcal{G}$ and a subgraph H' of H, we define the *restriction of* φ *to* H' as a cover ψ of H' by $G_1' \dot\cup \dots \dot\cup G_k'$, where G_i' comes from G_i by deleting $\{e \in E(G_i) : \ \varphi(e) \notin H'\}$ and then by removing isolated vertices. The resulting mapping ψ is the restriction of φ to $G_1' \dot\cup \dots \dot\cup G_k'$. If $\mathcal{G}$ is closed under taking subgraphs, then ψ is also a $\mathcal{G}$-cover. Note that while restriction of a function normally means its specification on a subset of the domain, here we are restricting the image, which turn induces a restriction of the domain.

To increase readability we refer to the classes of size m and n in the bipartition of $K_{m,n}$ by A and B, respectively.

Lemma 11. *Let H be a graph with an induced $K_{m,n}$ and φ be a $\mathcal{C}a$-cover of H with $s = \max\{|\varphi^{-1}(a)| : \ a \in A\}$. If ψ is the restriction of φ to the subgraph H' of H after removing all edges in $K_{m,n}$, then there are at least $n - 2sm$ vertices $b \in B$ such that $|\psi^{-1}(b)| \le |\varphi^{-1}(b)| - m$.*

Proof. Every $a \in A$ is the image of at most s vertices among $C_1 \dot\cup \dots \dot\cup C_k$. Denote by s' the number of vertices in $\varphi^{-1}(a)$ that are incident to two spine-edges and by s'' the number of vertices in $\varphi^{-1}(a)$ that are leaves. Clearly, $s' + s'' \le s$. Moreover, at most $2s' + s''$ edges incident to a are covered by spine-edges or edges whose degree 1 vertex is mapped to a. Therefore, at least $n - 2s$ edges at a have to be covered under φ by a non-spine edge with a vertex b being the image of a leaf. Thus, for at least $n - 2sm$ vertices $b \in B$ this is the case with respect to every $a \in A$.

Now if $e = ab$ is covered by some edge in C_i with b being a leaf, then in the restriction of φ to $H \setminus e$ the number of preimages of b is one less than in φ. This concludes the proof. □

Theorem 12. *For $k \ge 1$ there is a bipartite graph H such that*

$$2\mathrm{dgn}(H) \le 2k \le \mathrm{ca}(H) = t(H).$$

Proof. To construct H, begin with a copy of $K_{k,n}$ having $|A| = k$ and $|B| = n$ with $n > (k-1)\binom{2k-1}{k-1} + 2k(2k-1)$. For each k-subset S of B, add $(k-1)^2 + 1$ new vertices B_S with neighborhood S. The resulting graph H is bipartite with every vertex in A and B_S for any S having degree k, so $\mathrm{dgn}(H) = k$.

Now consider an injective $\mathcal{C}a$-cover φ of H and its restriction ψ to the subgraph of H after removing all edges in $K_{k,n}$. Assume for the sake of contradiction that the size s of φ is at most $2k-1$, i.e., $\max\{|\varphi^{-1}(v)| : v \in V(H)\} = s \le 2k-1$. Then by Lemma 11, there is a set $W \subset B$ of at least $n - 2(2k-1)k > (k-1)\binom{2k-1}{k-1}$ vertices such that $|\psi^{-1}(b)| \le |\varphi^{-1}(b)| - k \le s - k \le k - 1$ for every $b \in W$. In other words, every $b \in W$ has a preimage under ψ in at most $k-1$ of the $2k-1$ caterpillar forests. Since $|W| > (k-1)\binom{2k-1}{k-1}$, there is a k-set S in W whose preimages are contained in at most $k-1$ caterpillar forests.

This implies that ψ restricted to $H[S \cup B_S]$ is an injective $\mathcal{C}a$-cover of $K_{k,(k-1)^2+1}$ of size at most $k-1$, which is impossible since $\mathrm{ca}(K_{k,(k-1)^2+1}) = \left\lceil \frac{k(k-1)^2+k}{k+(k-1)^2} \right\rceil = k$, due to [9]. □

5.2. Bounded (simple) tree-width

A *k-tree* is a graph that can be constructed starting with a $(k+1)$-clique and in every step attaching a new vertex to a k-clique of the already constructed graph. We use the term *stacking* for this kind of attaching. The *tree-width* $\mathrm{tw}(H)$ of a graph H is the minimum k such that H is a *partial* k*-tree*, i.e., H is a subgraph of some k-tree [51].

We consider a variation of tree-width, called *simple tree-width*. A simple k-tree is a k-tree with the extra requirement that there is a construction sequence in which no two vertices are stacked onto the same k-clique. Now, the *simple tree-width* $\mathrm{stw}(H)$ of H is the minimum k such that H is a partial simple k-tree, i.e., H is a subgraph of some simple k-tree.

For a graph H with $\mathrm{stw}(H) = k$ or $\mathrm{tw}(H) = k$ we fix any (simple) k-tree that is a supergraph of H and denote it by $\tilde{H}$. Clearly, H inherits a construction sequence from $\tilde{H}$, where some edges are omitted.

Lemma 13. *We have* $\mathrm{tw}(H) \le \mathrm{stw}(H) \le \mathrm{tw}(H) + 1$ *for every graph H.*

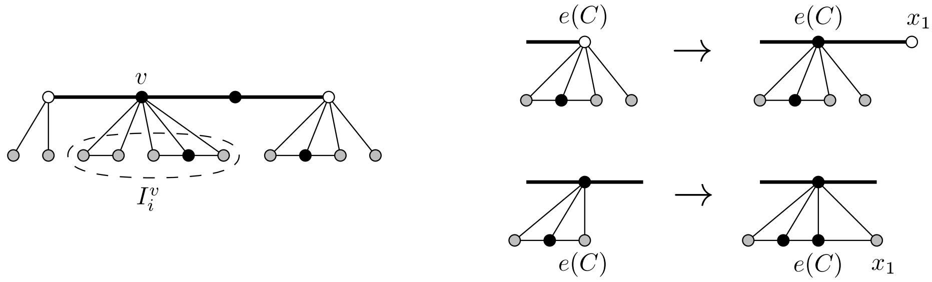

Fig. 2. A slug and its extension.

Proof. The first inequality is clear. For the second inequality we show that every k-tree H is a subgraph of a simple $(k+1)$-tree H. Whenever in the construction sequence of H several vertices $\{v_1, \dots, v_n\}$ are stacked onto the same k-clique C we consider $C \cup \{v_1\}$ as a $(k+1)$-clique in the construction sequence for H. Stacking v_i onto C now can be interpreted as stacking v_i onto $C \cup \{v_{i-1}\}$ and omitting the edge $v_{i-1}v_i$. In this way we can avoid multiple stackings onto k-cliques by considering $(k+1)$-cliques. □

Simple tree-width endows the notion of tree-width with a more topological flavor. For a graph H we have the following: $\mathrm{stw}(H) \leq 1$ if and only if H is a linear forest, $\mathrm{stw}(H) \leq 2$ if and only if H is outerplanar, $\mathrm{stw}(H) \leq 3$ if and only if H is planar and $\mathrm{tw}(H) \leq 3$ [18].

Simple tree-width also has connections to discrete geometry. In [12] a *stacked polytope* was defined to be a polytope that admits a triangulation whose dual graph is a tree. From that paper one easily deduces that a full-dimensional polytope $P \subset \mathbb{R}^d$ is stacked if and only if $\mathrm{stw}(G_P) \leq d$. Here G_P denotes the 1-skeleton of P. See [40,32,33] for more on simple tree-width.

We consider both graphs with bounded tree-width and graphs with bounded simple tree-width as input classes, since (A) most of the results for outerplanar graphs are implied by the corresponding result for $\mathrm{stw}(H) \leq 2$, (B) lower bound results for $\mathrm{stw}(H) \leq 3$ carry over to planar graphs, (C) the extremal results for these two input classes differ when the covering class is that of interval graphs, and (D) when the maximum covering numbers are the same for both classes, the lower bounds are slightly stronger when witnessed by graphs of low simple tree-width.

Theorem 14. *We have $t_\ell(H) \leq \mathrm{stw}(H)$ for every graph H.*

Proof. If $\mathrm{stw}(H) = 1$, then H is a linear forest and hence an interval graph. If $\mathrm{stw}(H) = 2$, then H is outerplanar, and it even has track number at most 2 as shown in [41].

So let $\mathrm{stw}(H) = s \geq 3$. We build an injective cover $\varphi : I_1 \dot\cup \cdots \dot\cup I_k \to H$ with $|\varphi^{-1}(v)| \leq s$ for every $v \in V(H)$ and $I_i \in \mathcal{I}$ for $i \in [k]$. We use as $I_1, \dots, I_k$ only certain interval graphs, which we call *slugs*: A slug is like a caterpillar with a fixed spine, except that the graph I_i^v induced by the leaves at every spine vertex $v \in I_i$ is a linear forest. (In a caterpillar I_i^v is an independent set for every spine vertex v.) The end vertices of the spine are called *spine-ends* and vertices of degree at most 1 in I_i^v are called *leaf-ends*. See the left of Fig. 2 for an example of a slug I_i with the spine drawn thick, spine-ends in white, and leaf-ends in gray. Note that slugs are indeed interval graphs.

We define the cover φ along a construction sequence of H that is inherited from a simple s-tree $\tilde{H} \supseteq H$. At every step let H' be the subgraph of H that is already constructed and hence already covered by φ, and let $\tilde{H}'$ be the corresponding subgraph of $\tilde{H}$. We call an s-clique C of $\tilde{H}'$ *stackable* if no vertex has been stacked to C so far. We maintain the following invariants on φ, which allow us to stack a new vertex onto every stackable C.

Invariant. At all times the following is satisfied for the current graph H'.

(1) For every vertex v in H' there is a unique slug $I(v)$ with $I(v) \neq I(w)$ for $v \neq w$, and a spine vertex $s(v)$ of $I(v)$ in $\varphi^{-1}(v)$.
(2) For every stackable s-clique C there is a vertex $w_1 \in C$, a slug $I(C)$, and a spine-end or leaf-end $e(C)$ of $I(C)$ with $\varphi(e(C)) = w_1$, such that:
  (2a) If $e(C)$ is a spine-end, then $I(C) \neq I(v)$ for all $v \in V(H')$.
  (2b) If $e(C)$ is a leaf-end, then $I(C) = I(w_2)$ for some vertex $w_2 \in C \setminus \{w_1\}$, and the vertices $e(C)$ and $s(w_2)$ are adjacent in $I(C)$.
  (2c) Every leaf-end or spine-end v is $e(C)$ for at most two cliques C with equality only if v has degree 0 or 1 in the slug.

It is not difficult to satisfy the above invariants for an initial s-clique of $\tilde{H}$. Indeed, this clique can be build up in a very similar way to the stacking procedure that we describe now: In the construction sequence of H we are about to stack a vertex w onto a stackable clique C of the current graph H'. Let $C = \{w_1, \dots, w_s\}$. Without loss of generality we assume that $\varphi(e(C)) = w_1$ and that if $e(C)$ is a leaf-end, then $I(C) = I(w_2)$. We never change the preimages of vertices in H under φ. In particular, all vertices we add to the existing or new slugs are mapped by φ onto the new vertex w. We will denote these new vertices by $x_1, \dots, x_s$ to emphasize that no more than s such vertices are introduced. Note that for every $i \in [s]$, the clique C_i in $\tilde{H}$ defined by $C_i = (C \setminus \{w_i\}) \cup \{w\}$ is stackable in $H' \cup \{w\}$, and that all remaining stackable cliques in $H' \cup \{w\}$ are already stackable cliques in H'.

For $i \in \{3, \ldots, s\}$ we do the following. If $ww_i \in E(H)$, then we introduce a new leaf x_i to $I(w_i)$ at $s(w_i)$, and if $ww_i \notin E(H)$ we introduce a new slug consisting only of x_i. Either way, we set $e(C_{i-1}) = x_i$. Additionally we set $e(C_1) = x_s$. Note that (2b) is satisfied since $w_i, w \in C_{i-1}$ and $w_s, w \in C_1$.

It remains to cover possible edges joining w to $\{w_1, w_2\}$, to find a spine-end or leaf-end $e(C_s)$ for C_s, and to find a slug $I(w)$ for the new vertex w. In doing so we may still introduce two new vertices x_1 and x_2 to our slugs. We distinguish two cases, which are illustrated on the right in Fig. 2.

Case 1: If $e(C)$ is a spine-end of $I(C)$, then we first proceed with w_2 similarly as with w_i for $i \geq 3$ above. That is, we introduce a new leaf x_2 at $s(w_2)$ if $ww_2 \in E(H)$ and a new slug consisting only of x_2 if $ww_2 \notin E(H)$, and we set $e(C_s) = x_2$.

Case 1.1: If $ww_1 \in E(H)$, then we introduce a new spine vertex x_1 to $I(C)$ adjacent to $e(C)$. This covers the edge ww_1, since we assumed that $\varphi(e(C)) = w_1$. We set $I(w) = I(C)$, which satisfies condition (1) of the invariant since (2a) implies $I(C) \neq I(v)$ for every vertex v in H'.

Case 1.2: If $ww_1 \notin E(H)$, then we introduce a new slug I consisting only of x_1 and set $I(w) = I$.

Case 2: If $e(C)$ is a leaf-end of $I(C)$, then by assumption we have $I(C) = I(w_2)$.

Case 2.1: If $ww_2 \in E(H)$, then we introduce a new leaf x_2 to $I(C)$ adjacent to $s(w_2)$ and a new slug I consisting just of a new vertex x_1. If additionally $ww_1 \in E(H)$, then we also introduce an edge joining x_2 and $e(C)$ in $I(C)$. Again, since $\varphi(e(C)) = w_1$ and $\varphi(x_2) = w$, this covers the edge ww_1. Either way, we set $e(C_s) = x_2$ and $I(w) = I$.

Case 2.2: If $ww_2 \notin E(H)$, then we introduce a new slug I consisting only of a new vertex x_2 and set $I(w) = I$. When $ww_1 \in E(H)$ we add a new leaf x_1 to $s(w_1)$ in $I(w_1)$, and when $ww_1 \notin E(H)$, then we introduce a new slug consisting only of x_1. Either way we set $e(C_s) = x_1$.

It is straightforward to check that we obtain a $\mathcal{I}$-cover of $H' \cup \{w\}$ and that the invariants above are satisfied. Note that since $\tilde{H}$ is a simple s-tree, the clique C is no longer stackable and hence condition (2) of the invariant need not be satisfied in $H' \cup \{w\}$. Finally, every stackable clique in H' different from C was not affected by the above procedure, which completes the proof. □

We can prove three lower bounds for covering numbers.

Theorem 15. *For $k \geq 1$, there is a bipartite graph H such that*
$\mathrm{stw}(H) \leq \mathrm{tw}(H) + 1 \leq k + 1 \leq \mathrm{ca}_f(H) = i(H)$.

Proof. Construct H from $K_{k,n}$ with $n = 2k^2 + 1$ by adding a pendant vertex at each vertex of the larger partite set B. It is easy to see that $\mathrm{tw}(H) \leq k$, and then Lemma 13 yields $\mathrm{stw}(H) \leq \mathrm{tw}(H) + 1$.

Consider any $\mathcal{C}a$-cover φ of H with $s = \max\{|\varphi^{-1}(v)| : v \in V(H)\}$ and its restriction ψ to the subgraph H' of H obtained by removing all edges of $K_{k,n}$. By Lemma 11 there are at least $n - 2sk = 2k(k - s) + 1$ vertices $b \in B$ such that $|\psi^{-1}(b)| \leq |\varphi^{-1}(b)| - k$. Any such b is incident to an edge in $H \setminus K_{k,n}$, which should be covered by ψ. Thus, $|\psi^{-1}(b)| \geq 1$. Hence, $s \geq |\varphi^{-1}(b)| \geq k + 1$, so $\mathrm{ca}_f(H) \geq k + 1$. □

Theorem 16. *For $k \geq 3$, there is a bipartite graph H such that*
$\mathrm{stw}(H) + 1 \leq k + 1 \leq \mathrm{ca}(H) = t(H)$.

Proof. The construction of the graph H starts with $H_0 \cong K_{k-1,m_1}$, where $|B| = m_1 = 2(2k^2 - 2k + 1)$. Let $B = \{u_1, \ldots, u_{m_1/2}\} \cup \{v_1, \ldots, v_{m_1/2}\}$. For $i \in [m_1/2]$, add a copy I_i of $K_{2,5k-5}$ with partite sets $\{u_i, v_i\}$ and $\{b_1^{i,j}, \ldots, b_{k-1}^{i,j} : j \in [5]\}$, calling the smaller set A_i and the larger set B_i. Next, let $m_2 = (k - 2)^+1$. For $(i, j) \in [m_1/2] \times [5]$, add a set $B_{i,j}$ of m_2 new vertices and a copy $J_{i,j}$ of K_{k-1,m_2} with partite sets $b_1^{i,j}, \ldots, b_{k-1}^{i,j}$ and $B_{i,j}$. Note that the smaller part $A_{i,j}$ in $J_{i,j}$ is contained in B_i. See Fig. 3 for an illustration.

Assume for the sake of contradiction that φ is an injective $\mathcal{C}a$-cover of H of size at most k. Consider the restriction ψ of φ to the subgraph $H' = H \setminus E(H_0)$ of H. By Lemma 11 there are at least $m_1 - 2k(k - 1) = 2k^2 - 2k + 2 > \frac{m_1}{2}$ vertices in $b \in B$ with $|\psi^{-1}(b)| \leq 1$. In particular there is some $i' \in [\frac{m_1}{2}]$ such that $|\psi^{-1}(u_i')|, |\psi^{-1}(v_i')| \leq 1$. That is, in the covering u_i' and v_i' each appear in only one caterpillar forest, which we call $C_{u_i'}$ containing u_i and $C_{v_i'}$ containing v_i'. Now consider the restriction ϕ of ψ to the subgraph $H'' = H' \setminus E(I_{i'})$ of H'. Again by Lemma 11 there are at least $5(k - 1) - 4$ vertices $b \in B_{i'}$ with $|\phi^{-1}(b)| \leq k - 2$. In particular there is some $j' \in [5]$ such that $|\phi^{-1}(b)| \leq k - 2$ for all $b \in A_{i'j'}$.

In other words, ϕ restricted to $H[A_{i'j'} \cup B_{i'j'}]$ is an injective $\mathcal{C}a$-cover of $K_{k-1,(k-2)^2+1}$ of size at most $k - 2$, which is impossible, since $\mathrm{ca}(K_{k-1,(k-2)^2+1}) = \left\lceil \frac{(k-1)(k-2)^2+k-1}{k-1+(k-2)^2} \right\rceil = k - 1$, due to [9].

It remains to show that $\mathrm{stw}(H) \leq k$. In order to describe the construction sequence for a simple k-tree containing H, we introduce some further vertex labels. Let $A_0 = \{a_1, \ldots, a_{k-1}\}$ be the smaller partite set of H_0, recall that $B_i = A_{i1} \cup \cdots \cup A_{i5}$ where $A_{ij} = \{b_1^{ij}, \ldots, b_{k-1}^{ij}\}$ for all $i \in [\frac{m_1}{2}], j \in [5]$, and let $B_{ij} = \{c_1^{ij}, \ldots, c_{m_2}^{ij}\}$ for $i \in [\frac{m_1}{2}], j \in [5]$. We construct a simple k-tree starting with a $(k + 1)$-clique on $A \cup \{u_1, v_1\}$ via the following stackings:

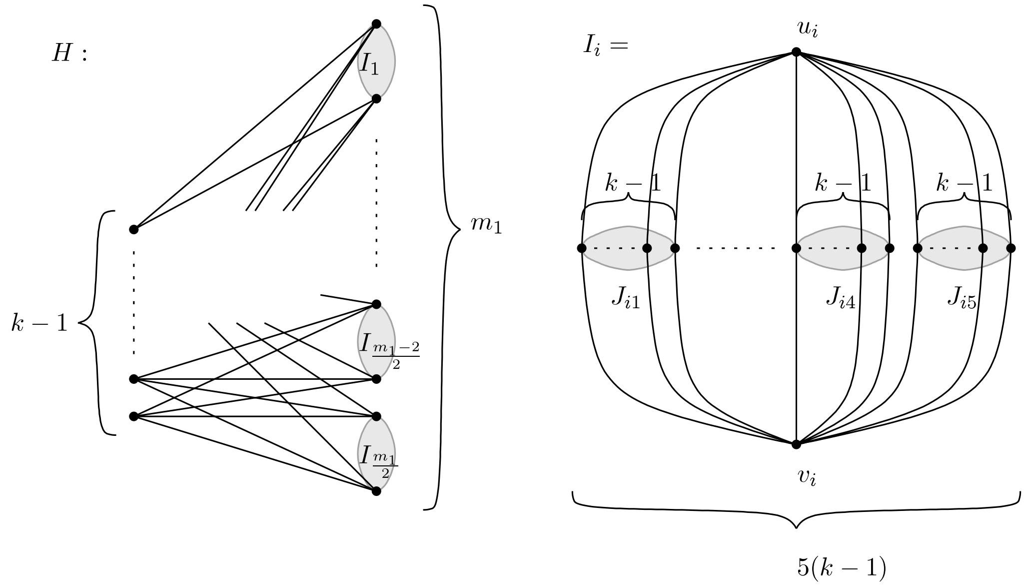

Fig. 3. The graph H and its induced subgraph I_i and J_j.

A Stack u_i onto $A \cup \{v_{i-1}\}$ and v_i onto $A \cup \{u_i\}$ $\forall i \in \{2, \ldots, \frac{m_1}{2}\}$
B Stack b_ℓ^{i1} onto $\{a_1, \ldots, a_{k-\ell-1}, u_i, v_i, b_1^{i1}, \ldots, b_{\ell-1}^{i1}\}$ $\forall i \in [\frac{m_1}{2}], \ell \in [k-1]$
C Stack b_ℓ^{ij} onto $\{u_i, v_i, b_1^{i(j-1)}, \ldots, b_{k-\ell-1}^{i(j-1)}, b_1^{ij}, \ldots, b_{\ell-1}^{ij}\}$ $\forall i \in [\frac{m_1}{2}], \ell \in [k-1], j \geq 2$
D Stack c_1^{ij} onto $A_{ij} \cup \{u_i\}$ and c_ℓ^{ij} onto $A_{ij} \cup \{c_{\ell-1}^{ij}\}$ $\forall i \in [\frac{m_1}{2}], j \in [5], \ell \geq 2$.

One can check that after step A. the entire graph H_0 is contained in the so-far constructed k-tree. Step B. deals with the complete bipartite graphs induced on $\{u_i, v_i\} \cup A_{i1}$ for all $i \in [\frac{m_1}{2}]$, step C. adds the remaining complete bipartite graphs induced on $\{u_i, v_i\} \cup A_{ij}$ for $j \geq 2$, such that afterwards all I_i are contained. In step D. all edges and vertices necessary for the J_{ij} are created. Since no k-clique appears twice we conclude that $\mathrm{stw}(H) \leq k$. □

Theorem 17. *For $k \geq 2$, there is a graph H such that*

$$\mathrm{stw}(H) + 1 \leq k + 1 \leq \mathrm{ca}_f(H).$$

Proof. Fix $k \geq 2$. We construct H starting with a star with $k-1$ leaves $\ell_1, \ldots, \ell_{k-1}$ and center c_1. In the simple partial k-tree containing H this star is a k-clique. For $n = 16k^2 - 16k + 4$ and $2 \leq i \leq n$ stack a new vertex c_i to $\ell_1, \ldots, \ell_{k-1}, c_{i-1}$. Now stack vertices $s_2, \ldots, s_n$ to $\ell_1, \ldots, \ell_{k-2}, c_{i-1}, c_i$. Finally introduce a pendant vertex a_i as a neighbor of s_i, for each i. In the simple partial k-tree containing H, the vertex a_i is stacked to the k-clique on $\ell_1, \ldots, \ell_{k-2}, c_{i-1}, s_i$. By construction $\mathrm{stw}(H) \leq k$. See Fig. 4 for an illustration.

Assume for the sake of contradiction that $\mathrm{ca}_f(H) \leq k$. That is, there is a $\mathcal{C}a$-cover φ of H with $|\varphi^{-1}(v)| \leq k$ for all $v \in V(H)$. We consider three edge-disjoint complete bipartite subgraphs H_1, H_2, H_3 of H with partite sets A_i and B_i for H_i defined as follows:

- $A_1 = \ell_1, \ldots, \ell_{k-1}$ and $B_1 = \{c_{2i} \colon 1 \leq i \leq n/2\}$
- $A_2 = \ell_1, \ldots, \ell_{k-1}$ and $B_2 = \{c_{2i-1} \colon 1 \leq i \leq n/2\}$
- $A_3 = \ell_1, \ldots, \ell_{k-2}$ and $B_3 = \{s_i \colon 2 \leq i \leq n\}$.

Note that H_i and H_j are edge-disjoint for $i \neq j$. Denote by ψ the restriction of φ to $H \setminus (E(H_1) \cup E(H_2) \cup E(H_3))$. We apply Lemma 11 three times, once for each H_i, but the bounds for restrictions of φ to $H \setminus E(H_i)$ clearly also apply to ψ. Thus, we obtain sets $W_i \subset B_i$ (for each $i \in \{1, 2, 3\}$). For $i \in \{1, 2\}$ we get $|W_i| \geq n/2 - 2k(k-1)$ and $\psi^{-1}(b) \leq k - (k-1) = 1$ for $b \in W_i$. Furthermore we have $|W_3| \geq n - 1 - 2k(k-2)$ and $\psi^{-1}(b) \leq k - (k-2) = 2$ for $b \in W_3$. From the choice of n it follows that there exist $c_i, c_{i+1}, c_{i+2}, c_{i+3} \in W_1 \cup W_2$ with consecutive indices such that $s_{i+1}, s_{i+2}, s_{i+3} \in W_3$. Together with the leaves $a_{i+1}, a_{i+2}, a_{i+3}$ these vertices induce a 10-vertex graph H' highlighted in Fig. 4. It is not difficult to check that there is no $\mathcal{C}a$-cover ψ of H' with $|\psi^{-1}(c_{i+j})| \leq 1$ for $j \in \{0, 1, 2, 3\}$ and $|\psi^{-1}(s_{i+j})| \leq 2$ for $j \in \{1, 2, 3\}$—a contradiction. □

5.3. Planar and outerplanar graphs

Determining maximum covering numbers of (bipartite) planar graphs and outerplanar graphs enjoys a certain popularity, as demonstrated by the variety of citations in Table 1. We add three easy new results to the list.

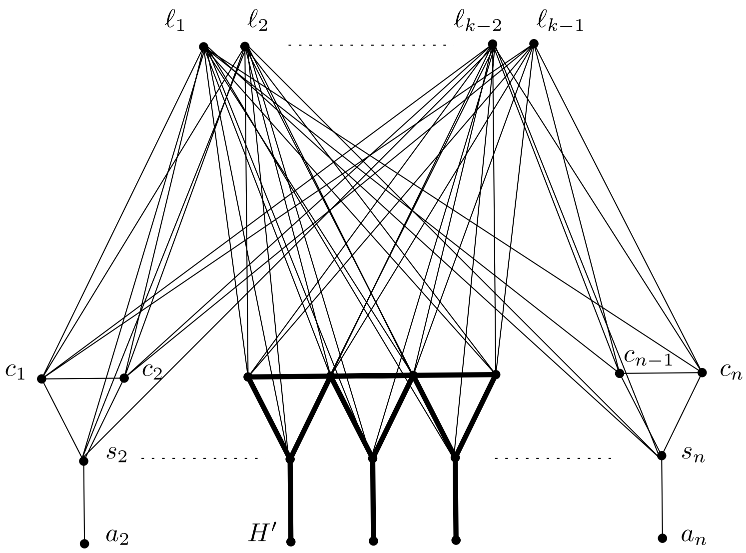

Fig. 4. The graph H and its subgraph H'.

Corollary 18. *The star arboricity of bipartite planar graphs is at most* 4*. The local star arboricity of planar graphs and bipartite planar graphs is at most* 4 *and at most* 3*, respectively.*

Proof. As mentioned in Section 4.2, the arboricity $a(H)$ of every graph H can be expressed as $\max_{S\subseteq V(H)} \left\lceil \frac{|E[S]|}{|S|-1} \right\rceil$ [46]. By Euler's Formula every planar graph has at most $3V(H) - 6$ edges and every bipartite planar graph has at most $2V(H) - 4$ edges and clearly both classes are closed under taking subgraphs. Together it follows that every planar graph has arboricity at most 3 and every planar bipartite graph has arboricity at most 2. With this, the statement about global star arboricity follows since we have $\mathrm{sa}(H) \leq 2a(H)$ by [6]. The statements about local arboricity follow since we have $\mathrm{sa}_\ell(H) \leq a(H) + 1$ by Theorem 9. □

The only question mark in Table 1 concerns the local track number of planar graphs. Scheinerman and West [52] show that the interval number of planar graphs is at most 3, but this is verified with a cover that is not injective. On the other hand, there are bipartite planar graphs with track number 4 [24]. However by Corollary 18 and Theorem 14 every bipartite planar graph and every planar graph of tree-width at most 3 has local track number at most 3. We believe that there are planar graphs with local track number 4, but the following remains open:

Question 19. *What is the maximum local track number of a planar graphs?*

6. Separability and complexity

This section is devoted to different types of questions. First, we investigate how much global, local, and folded covering numbers can differ with respect to the same covering and input class. Second, we look at the complexity of computing these parameters.

In Table 1 we provide several pairs of an input class $\mathcal{H}$ and a covering class $\mathcal{G}$ for which the global covering number and the local covering number differ, i.e., $c_g^{\mathcal{G}}(\mathcal{H}) > c_\ell^{\mathcal{G}}(\mathcal{H})$. Indeed this difference can be arbitrarily large.

Theorem 20. *For the covering class $\mathcal{Q}$ of collections of cliques and the input class $\mathcal{H}$ of line graphs, we have $c_g^{\mathcal{Q}}(\mathcal{H}) = \infty$ and $c_\ell^{\mathcal{Q}}(\mathcal{H}) \leq 2$.*

Proof. By a result of Whitney [58] a graph H is a line graph if and only if $c_\ell^{\mathcal{Q}}(H) \leq 2$.

To prove $c_g^{\mathcal{Q}}(\mathcal{H}) = \infty$, we claim that $c_g^{\mathcal{Q}}(L(K_n)) \in \Omega(\log n)$, i.e., the covering number of the line graph of the complete graph on n vertices is unbounded as n goes to infinity. Assume that $L(K_n)$ is covered by k collections of cliques $C_1, \dots, C_k$. Every clique in $L(K_n)$ corresponds to either a triangle or a star in K_n. Now, every C_i in $L(K_n)$ corresponds to a vertex disjoint collection of triangles and stars in K_n. Together these collections cover the edges of K_n. We will restrict the covering of $L(K_n)$ to a covering of $L(K_m)$ with collections of cliques all of whose cliques correspond to stars in K_m. In the first step delete at most $\frac{1}{3}n$ vertices of K_n such that in the restricted cover of the smaller line graph no clique in C_1 corresponds to a triangle. Repeating this for every C_i, we end up with a clique cover of $L(K_m)$ with $m \geq (\frac{2}{3})^k n$ that corresponds to a cover of K_m with star forests. Since by [4] the star arboricity of K_m is $\lceil \frac{m}{2} \rceil + 1$, we get $k \geq \frac{m+2}{2} > (\frac{2}{3})^{k-1} n$, and thus $k \in \Omega(\log n)$. □

Remark 21. Milans, Stolee, and West [44] proved a similar result with interval graphs as covering class, i.e., they showed that the growth rate of $t(L(K_n))$ is between $\Omega(\log\log n/\log\log\log n)$ and $O(\log\log n)$, while $i(H) \le 2$ for every line graph H.

A case of particular interest to us is the input class of claw-free graphs—a class containing line graphs. It has been shown that this class has unbounded local clique covering number [36]. We conjecture the following stronger statement:

Conjecture 22. *The class of claw-free graphs has unbounded interval number.*

What can be said about local and folded covering number? Table 1 suggests that the separation of the local and the folded covering number is more difficult. Indeed we have $c_\ell^{\mathcal{G}}(\mathcal{H}) = c_f^{\mathcal{G}}(\mathcal{H})$ for every $\mathcal{G}$ and $\mathcal{H}$ in Table 1, except for the local track number of planar graphs, (c.f. Question 19). However, proving upper bounds for $c_\ell^{\mathcal{G}}(\mathcal{H})$ can be significantly more elaborate than for $c_f^{\mathcal{G}}(\mathcal{H})$, even if we suspect that both values are equal; see for example Conjecture 1 and Theorem 3.

Observing that there is no injective cover of a path by cycles of length at least 3 and that every path is the homomorphic image of a cycle one gets:

Observation 23. *For the covering class $\mathcal{C}$ of collections of cycles of length at least 3 and the input class $\mathcal{H}$ of paths, we have $c_\ell^{\mathcal{C}}(\mathcal{H}) = \infty$ and $c_f^{\mathcal{C}}(\mathcal{H}) \le 2$.*

Observation 23 may be considered pathological. However, the local and folded covering number may differ also when $c_\ell^{\mathcal{G}}(H) < \infty$. We gave one example for this when considering coverings of the Petersen graph with disjoint unions of cycles, see Proposition 6. Here is another example: It is known that $i(K_{m,n}) = \lceil \frac{mn+1}{m+n} \rceil$ [31] and $t(K_{m,n}) = \lceil \frac{mn}{m+n-1} \rceil$ [28]. The lower bound on $t(K_{m,n})$ presented in [14] indeed gives $t_\ell(K_{m,n}) \ge \lceil \frac{mn}{m+n-1} \rceil$ and hence we have $t_\ell(K_{m,n}) > i(K_{m,n})$ for appropriate numbers m and n, such as $n = m^2 - 2m + 2$. With Proposition 4 this translates into $\mathrm{ca}_\ell(K_{m,n}) > \mathrm{ca}_f(K_{m,n})$. Apart from these examples, we have no general answer to the following question.

Question 24. *By how much can folded and local covering number differ?*

Another interesting aspect of covering numbers concerns the computational complexity of determining them. Very informally, one might suspect that the computation of $c_f^{\mathcal{G}}(H)$ is easier than of $c_\ell^{\mathcal{G}}(H)$, which in turn is easier than computing $c_g^{\mathcal{G}}(H)$. For example, if $\mathcal{M}$ is the class of all matchings, then $c_g^{\mathcal{M}}(H) = \chi'(H)$, the edge-chromatic number of H. Hence deciding $c_g^{\mathcal{M}}(H) \le 3$ is NP-complete even for 3-regular graphs [34]. On the other hand $c_\ell^{\mathcal{M}}(H)$ equals the maximum degree of H and can therefore be determined very efficiently. As a second example, more elaborate, consider the star arboricity $\mathrm{sa}(H)$ and the caterpillar arboricity $\mathrm{ca}(H)$. Deciding $\mathrm{sa}(H) \le k$ [42,24] and deciding $\mathrm{ca}(H) \le k$ [24,53] are NP-complete for $k = 2, 3$. The complexity for $k \ge 4$ is unknown in both cases. To the best of our knowledge, the complexity of determining the local and folded caterpillar arboricity of a graph is also open. On the other hand, from Theorem 9 we can derive the following.

Theorem 25. *The local star arboricity can be computed in polynomial-time.*

Proof. In [22] a flow algorithm is used that given a graph H and $\alpha : V(H) \to \mathbb{N}$ decides if an orientation D of H exists such that the out-degree of v in D is at most $\alpha(v)$ for all $v \in V(H)$. Moreover, if such a D exists the algorithm finds one minimizing the maximum out-degree. Now by Lemma 7, we may use this algorithm to find $p(H)$ in polynomial-time. Now let $\alpha(v) = p(H)$ whenever v has degree $p(H)$ and $\alpha(v) = p(H) - 1$ otherwise. We use the algorithm of [22] to check if an orientation D of H satisfying the out-degree constraints given by α exists. By Theorem 9 we have $\mathrm{sa}_\ell(H) = p(H)$ if and only if there exists such an orientation and $\mathrm{sa}_\ell(H) = p(H) + 1$ otherwise. □

Finally, consider interval graphs as the covering class. Shmoys and West [54] and Jiang [38] showed that deciding $i(H) \le k$ and deciding $t(H) \le k$ are NP-complete for every $k \ge 2$, respectively. We claim that the reduction of Jiang also holds for the local track number.

Question 26. *Are there a covering and an input class for which the computation of the folded or local covering number is NP-complete while the global covering number can be computed in polynomial-time?*

7. Concluding remarks

We have presented new ways to cover a graph and given many example covering classes. Also, we highlighted some conjectures and questions on the way, such as the question whether the maximum track number of planar graphs is 3 or 4 (Question 19).

One conjecture important to us is LLAC (Conjecture 1), which is a weakening of the linear arboricity conjecture (LAC). Besides LLAC, there are several more weakenings of LAC that are still open. For example it is open, whether the caterpillar arboricity of graph H of maximum degree $\Delta(H)$ is always at most $\left\lceil \frac{\Delta(H)+1}{2} \right\rceil$. Yet a weaker, but still open, question asks whether the track number of H is always at most $\left\lceil \frac{\Delta(H)+1}{2} \right\rceil$. As a positive result, by Theorem 9 one obtains that for a regular graph H of even degree the local star-arboricity is $\left\lceil \frac{\Delta(H)+1}{2} \right\rceil$, which in particular settles the question for local caterpillar arboricity and local track number for such input graphs. On the other hand, Theorem 9 also tells us that in a regular graph H of odd degree the local star arboricity is larger than $\left\lceil \frac{\Delta(H)+1}{2} \right\rceil$. To the best of our knowledge, it is open whether the local caterpillar arboricity or local track number of such a graph H is always at most $\left\lceil \frac{\Delta(H)+1}{2} \right\rceil$.

Apart from the problems already mentioned throughout the paper, it is interesting to consider the local and folded variants for more graph covering problems from the literature. For example the covering number with respect to planar and outerplanar graphs is known as the *thickness* and *outerthickness* [11], respectively, and the folded covering number with respect to planar graphs is called the *splitting number* [35]. The local covering number in these cases seems unexplored. Further interesting covering classes include linear forests of bounded length [8], forests of stars and triangles [20], and chordal graphs.

A concept dual to covering is *packing*. For an input graph H and a class $\mathcal{G}$ of *packing graphs*, we define a *$\mathcal{G}$-packing of H* to be an edge-injective homomorphism φ to H from the disjoint union $G_1 \dot\cup G_2 \dot\cup \cdots \dot\cup G_k$ with $G_i \in \mathcal{G}$ for $i \in [k]$. The *size* of a packing is the number of packing graphs in the disjoint union. A packing φ is *injective* if $\varphi|G_i$, that is, φ restricted to G_i, is injective for every $i \in [k]$.

Definition 2. For a packing class $\mathcal{G}$ and an input graph $H = (V, E)$ define the *(global) packing number* $p_g^{\mathcal{G}}(H)$, the *local packing number* $p_\ell^{\mathcal{G}}(H)$, and the *folded packing number* $p_f^{\mathcal{G}}(H)$ as follows:

$$p_g^{\mathcal{G}}(H) = \max \left\{\text{size of } \varphi : \varphi \text{ is an injective } \mathcal{G}\text{-packing of } H\right\}$$
$$p_\ell^{\mathcal{G}}(H) = \max \left\{\min_{v\in V} |\varphi^{-1}(v)| : \varphi \text{ is an injective } \mathcal{G}\text{-packing of } H\right\}$$
$$p_f^{\mathcal{G}}(H) = \max \left\{\min_{v\in V} |\varphi^{-1}(v)| : \varphi \text{ is a } \mathcal{G}\text{-packing of } H \text{ having size } 1\right\}.$$

Let us rephrase $p_g^{\mathcal{G}}(H)$, $p_\ell^{\mathcal{G}}(H)$, and $p_f^{\mathcal{G}}(H)$: The packing number is the maximum number of packing graphs that can be packed into the input graph, where packing means identifying edge-disjoint subgraphs in H that lie in $\mathcal{G}$. The local packing number does not measure the number of packing graphs in a packing; instead the minimum number of graphs packed at any one vertex is maximized. The folded packing number is the maximum k such that every vertex v of H can be split into k vertices, distributing the incident edges at v arbitrarily (not repeatedly) among them, such that the resulting graph is in $\mathcal{G}$. Two classical packing problems are given by $\mathcal{G}$ being the class of non-planar graphs or non-outerplanar graphs. In this case the global packing numbers are called *coarseness* and *outercoarseness* [11], respectively.

Acknowledgments

We thank Marie Albenque, Daniel Heldt, and Bartosz Walczak for fruitful discussions and two anonymous referees and Douglas B. West for useful comments improving the presentation of the paper. Kolja Knauer was partially supported by DFG grant FE-340/8-1 as part of ESF project GraDR EUROGIGA and PEPS grant EROS and Torsten Ueckerdt by GraDR EUROGIGA project No. GIG/11/E023.

References

[1] Alok Aggarwal, Maria Klawe, Peter Shor, Multilayer grid embeddings for VLSI, Algorithmica 6 (1991) 129–151.
[2] Jin Akiyama, Geoffrey Exoo, Frank Harary, Covering and packing in graphs. III. Cyclic and acyclic invariants, Math. Slovaca 30 (4) (1980) 405–417.
[3] Jin Akiyama, Geoffrey Exoo, Frank Harary, Covering and packing in graphs. IV, Linear arboricity, Networks 11 (1) (1981) 69–72.
[4] Jin Akiyama, Mikio Kano, Path factors of a graph, in: Graphs and Applications (Boulder, Colo., 1982), Wiley-Intersci. Publ., Wiley, New York, 1985, pp. 1–21.
[5] Ilan Algor, Noga Alon, The star arboricity of graphs, Discrete Math. 75 (1–3) (1989) 11–22. Graph theory and combinatorics (Cambridge, 1988).
[6] Noga Alon, Colin McDiarmid, Bruce Reed, Star arboricity, Combinatorica 12 (4) (1992) 375–380.
[7] Noga Alon, Joel H. Spencer, The Probabilistic Method, third ed., in: Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons Inc., Hoboken, NJ, 2008, With an appendix on the life and work of Paul Erdős.
[8] Noga Alon, Vanessa J. Teague, Nick C. Wormald, Linear arboricity and linear k-arboricity of regular graphs, Graphs Combin. 17 (1) (2001) 11–16.
[9] Lowell W. Beineke, Decompositions of complete graphs into forests, Publ. Math. Inst. Hung. Acad. Sci., Ser. A 9 (1965) 589–594.
[10] R. Bar-Yehuda, M.M. Halldorsson, J. (S.) Naor, H. Shachnai, I. Shapira, Scheduling split intervals, SIAM J. Comput. 36 (1) (2006) 1–15.
[11] Lowell W. Beineke, A survey of packings and coverings of graphs, in: The Many Facets of Graph Theory (Proc. Conf., Western Mich. Univ., Kalamazoo, Mich., 1968), Springer, Berlin, 1969, pp. 45–53.
[12] Alexander Below, Jesús A. De Loera, Jürgen Richter-Gebert, The complexity of finding small triangulations of convex 3-polytopes, J. Algorithms 50 (2) (2004) 134–167. SODA 2000 special issue.
[13] Ayelet Butman, Danny Hermelin, Moshe Lewenstein, Dror Rawitz, Optimization problems in multiple-interval graphs, ACM Trans. Algorithms 6 (2010) 40:1–40:18.
[14] Narsingh Deo, Nishit Kumar, Multidimensional interval graphs, in: Proceedings of the Twenty-fifth Southeastern International Conference on Combinatorics, Graph Theory and Computing (vol. 102, Boca Raton, FL, 1994), 1994, pp. 45–56.

[15] Guoli Ding, Bogdan Oporowski, Daniel P. Sanders, Dirk Vertigan, Partitioning graphs of bounded tree-width, Combinatorica 18 (1) (1998) 1–12.
[16] Jinquan Dong, Yanpei Liu, On the decomposition of graphs into complete bipartite graphs, Graphs Combin. 23 (3) (2007) 255–262.
[17] Vida Dujmović, David R. Wood, Graph treewidth and geometric thickness parameters, Discrete Comput. Geom. 37 (4) (2007) 641–670.
[18] Ehab S. El-Mallah, Charles J. Colbourn, On two dual classes of planar graphs, Discrete Math. 80 (1) (1990) 21–40.
[19] Genghua Fan, Covers of Eulerian graphs, J. Combin. Theory Ser. B 89 (2) (2003) 173–187.
[20] Jiří Fiala, Van Bang Le, The subchromatic index of graphs, Tatra Mt. Math. Publ. 36 (2007) 129–146.
[21] Peter C. Fishburn, Peter L. Hammer, Bipartite dimensions and bipartite degrees of graphs, Discrete Math. 160 (1–3) (1996) 127–148.
[22] András Frank, András Gyárfás, How to orient the edges of a graph? in: Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), vol. I, in: Colloq. Math. Soc. János Bolyai, vol. 18, North-Holland, Amsterdam, 1978, pp. 353–364.
[23] Daniel Gonçalves, Caterpillar arboricity of planar graphs, Discrete Math. 307 (16) (2007) 2112–2121.
[24] Daniel Gonçalves, Pascal Ochem, On star and caterpillar arboricity, Discrete Math. 309 (11) (2009) 3694–3702.
[25] Niall Graham, Frank Harary, Covering and packing in graphs. V. Mispacking subcubes in hypercubes, Comput. Math. Appl. 15 (4) (1988) 267–270.
[26] Jerrold R. Griggs, Douglas B. West, Extremal values of the interval number of a graph, SIAM J. Algebraic Discrete Methods 1 (1) (1980) 1–7.
[27] Filip Guldan, Some results on linear arboricity, J. Graph Theory 10 (4) (1986) 505–509.
[28] András Gyárfás, Douglas B. West, Multitrack interval graphs, in: Proceedings of the Twenty-sixth Southeastern International Conference on Combinatorics, Graph Theory and Computing (vol. 109, Boca Raton, FL, 1995), 1995, pp. 109–116.
[29] Frank Harary, Covering and packing in graphs. I, Ann. New York Acad. Sci. 175 (1970) 198–205.
[30] Frank Harary, Allen J. Schwenk, Evolution of the path number of a graph: Covering and packing in graphs. II, in: Graph Theory and Computing, Academic Press, New York, 1972, pp. 39–45.
[31] Frank Harary, William T. Trotter Jr., On double and multiple interval graphs, J. Graph Theory 3 (3) (1979) 205–211.
[32] Daniel Heldt, Kolja Knauer, Torsten Ueckerdt, Edge-intersection graphs of grid paths: the bend-number, Discrete Appl. Math. 167 (2014) 144–162.
[33] Daniel Heldt, Kolja Knauer, Torsten Ueckerdt, On the bend-number of planar and outerplanar graphs, Discrete Appl. Math. 179 (2014) 109–119.
[34] Ian Holyer, The NP-completeness of edge-coloring, SIAM J. Comput. 10 (4) (1981) 718–720.
[35] Brad Jackson, Gerhard Ringel, Splittings of graphs on surfaces, in: Graphs and Applications (Boulder, Colo., 1982), Wiley-Intersci. Publ., Wiley, New York, 1985, pp. 203–219.
[36] R. Javadi, Z. Maleki, B. Omoomi, Local Clique Covering of Graphs, arXiv:1210.6965 (2012).
[37] Minghui Jiang, Approximation algorithms for predicting RNA secondary structures with arbitrary pseudoknots, IEEE/ACM Trans. Comput. Biol. Bioinformatics 7 (2010) 323–332.
[38] Minghui Jiang, Recognizing d-interval graphs and d-track interval graphs, in: Proceedings of the 4th International Conference on Frontiers in Algorithmics, FAW'10, Springer-Verlag, Berlin, Heidelberg, 2010, pp. 160–171.
[39] Deborah Joseph, Joao Meidanis, Prasoon Tiwari, Determining DNA sequence similarity using maximum independent set algorithms for interval graphs, in: Algorithm Theory—SWAT'92, vol. 621, 1992, pp. 326–337.
[40] Kolja Knauer, Torsten Ueckerdt, Simple tree-width, in: Midsummer Combinatorial Workshop 2012, 2012, http://kam.mff.cuni.cz/workshops/work18/mcw2012booklet.pdf, pp. 21–23.
[41] Alexandr V. Kostochka, Douglas B. West, Every outerplanar graph is the union of two interval graphs, in: Proceedings of the Thirtieth Southeastern International Conference on Combinatorics, Graph Theory, and Computing, (vol. 139, Boca Raton, FL), 1999, pp. 5–8.
[42] S. Louis Hakimi, John. Mitchem, Edward F. Schmeichel, Star arboricity of graphs, Discrete Math. 149 (1–3) (1996) 93–98.
[43] László Lovász, On covering of graphs, in: Theory of Graphs (Proceedings of the Coloquium held at Tihanx, Hungary, September, 1966), 1968, pp. 231–236.
[44] Kevin G. Milans, Derrick Stolee, Douglas B. West, Ordered Ramsey theory and track representations of graphs, 2015, to appear in J. Combin.
[45] Petra Mutzel, Thomas Odenthal, and Mark Scharbrodt, The thickness of graphs: a survey, Graphs Combin. 14 (1) (1998) 59–73.
[46] Crispin St.J. A. Nash-Williams, Decomposition of finite graphs into forests, J. Lond. Math. Soc. 39 (1964) 12.
[47] Julius Petersen, Die Theorie der regulären Graphen, Acta Math. 15 (1) (1891) 193–220.
[48] Jean-Claude Picard, Maurice Queyranne, A network flow solution to some nonlinear 0-1 programming problems, with applications to graph theory, Networks 12 (2) (1982) 141–159.
[49] Trevor Pinto, Biclique covers and partitions, Electron. J. Combin. 21 (1) (2014) research paper, p1.19, 12.
[50] Subramanian Ramanathan, Errol L. Lloyd, Scheduling algorithms for multi-hop radio networks, SIGCOMM Comput. Commun. Rev. 22 (1992) 211–222.
[51] Neil Robertson, Paul D. Seymour, Graph minors—a survey, in: Surveys in combinatorics 1985 (Glasgow, 1985), in: London Math. Soc. Lecture Note Ser., vol. 103, Cambridge Univ. Press, Cambridge, 1985, pp. 153–171.
[52] Edward R. Scheinerman, Douglas B. West, The interval number of a planar graph: three intervals suffice, J. Combin. Theory Ser. B 35 (3) (1983) 224–239.
[53] Thomas C. Shermer, On rectangle visibility graphs. iii. External visibility and complexity, CCCG'96, 1996, pp. 234–239.
[54] David B. Shmoys, Douglas B. West, Recognizing graphs with fixed interval number is NP-complete, Discrete Appl. Math. 8 (3) (1984) 295–305.
[55] P.V. Skums, S.V. Suzdal, R.I. Tyshkevich, Edge intersection graphs of linear 3-uniform hypergraphs, Discrete Math. 309 (11) (2009) 3500–3517.
[56] Douglas B. West, A short proof of the degree bound for interval number, Discrete Math. 73 (3) (1989) 309–310.
[57] Douglas B. West, Introduction to Graph Theory, Prentice Hall Inc., Upper Saddle River, NJ, 1996.
[58] Hassler Whitney, Congruent graphs and the connectivity of graphs, Amer. J. Math. 54 (1) (1932) 150–168.
[59] Jian-Liang Wu, On the linear arboricity of planar graphs, J. Graph Theory 31 (2) (1999) 129–134.
[60] Jian-Liang Wu, Yu-Wen Wu, The linear arboricity of planar graphs of maximum degree seven is four, J. Graph Theory 58 (3) (2008) 210–220.

Dr. Torsten Ueckerdt
Klosestr. 41
76137 Karlsruhe
E-Mail: torsten.ueckerdt@kit.edu

Karlsruhe, December 1, 2016

Erklärung

Hiermit versichere ich, dass ich die beigefügte Habilitationsschrift selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel genutzt habe. Alle wörtlich oder inhaltlich übernommenen Stellen habe ich als solche gekennzeichnet.

Dr. Torsten Ueckerdt