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:		Communicated by:	
	Regular paper		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.

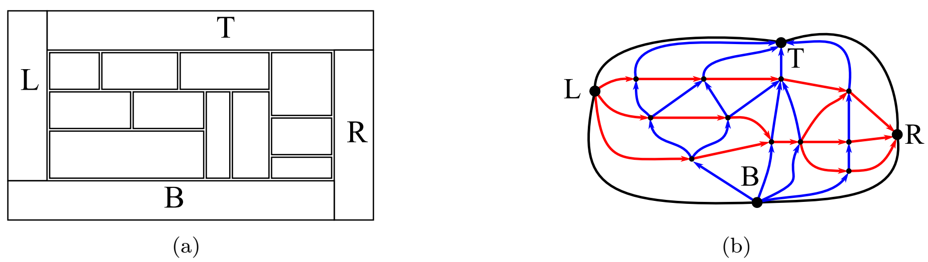

Figure 1: (a) A rectangular dual; and (b) its transversal structure.

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.

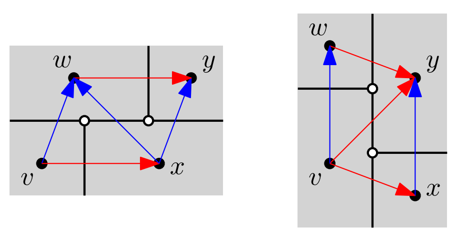

Figure 2: Two configurations that do not appear in the minimal transversal structure.

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

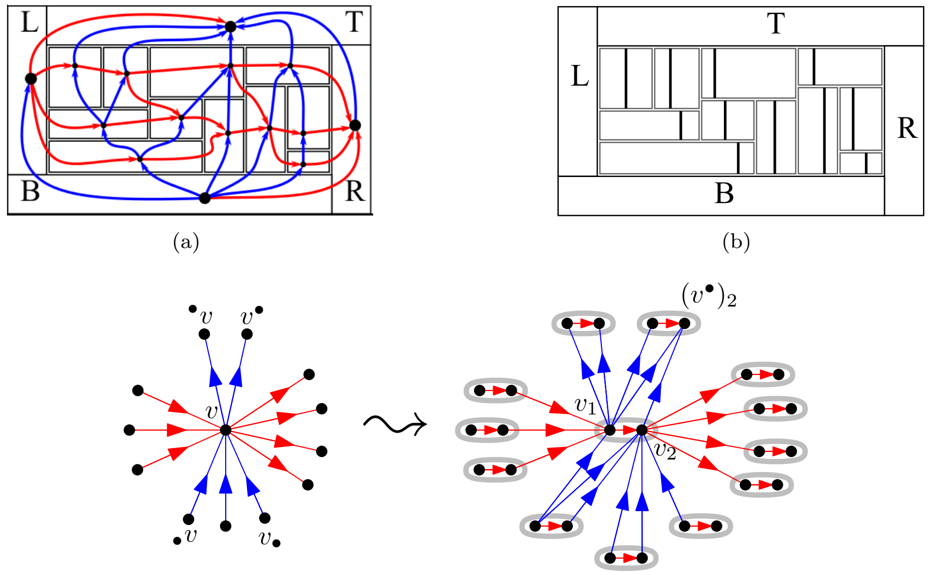

(c)

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

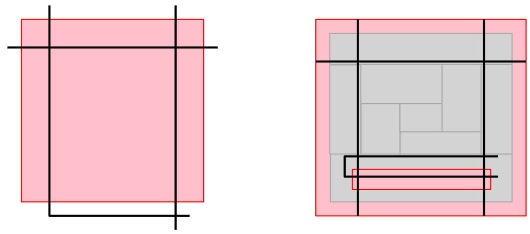

Figure 4: Left: The VPG representation of the outer face of G and its frame. Right: Placing a rectangular dual inside a frame and constructing the path $\mathbf{B}$.

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.

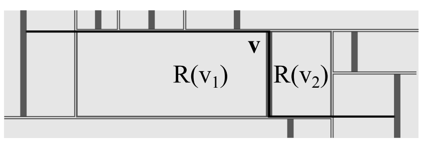

(a)

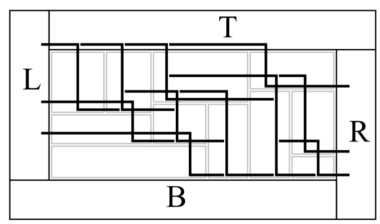

(b)

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 **v** lies on the right side of $R((v_\bullet)_1)$ (or **R** in case $v_\bullet = R$) and the left end of **v** lies on the right side of $R((^\bullet v)_1)$ (or **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 **v** that lie outside $R(v)$. The left extension of **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 **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 **L** and **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.

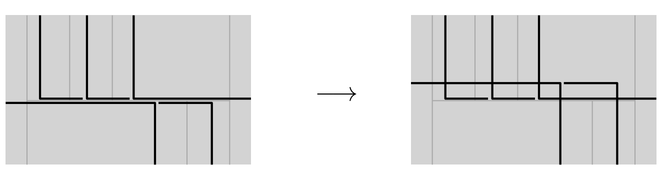

Figure 6: Extending the vertical segments of all paths in B.

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 $\mathbf{B}$ and the vertical segments of $\mathbf{L}$ and $\mathbf{R}$.

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.

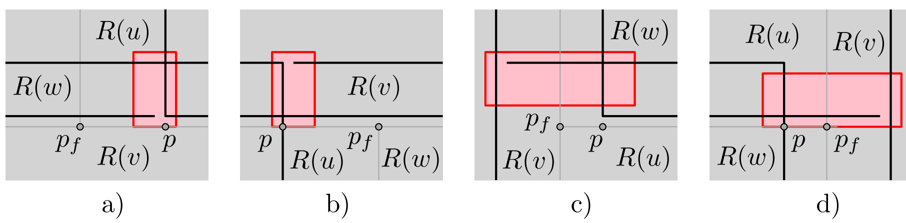

Figure 7: Identifying the frame for an inner face of H_Δ.

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 $\mathbf{u}$ and the right horizontal end of $\mathbf{w}$. Moreover, the upper horizontal segment of $\mathbf{v}$ lies immediately above p, crossing $\mathbf{u}$. 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 $\mathbf{u}$ and the left horizontal end of $\mathbf{w}$, while $\mathbf{v}$ 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 $\mathbf{u}$. The upper horizontal segment of $\mathbf{w}$ lies right above p and ends on the vertical segment of $\mathbf{v}$. 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 $\mathbf{w}$. The lower horizontal segment of $\mathbf{u}$ runs through p and ends on the vertical segment of $\mathbf{v}$. 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.

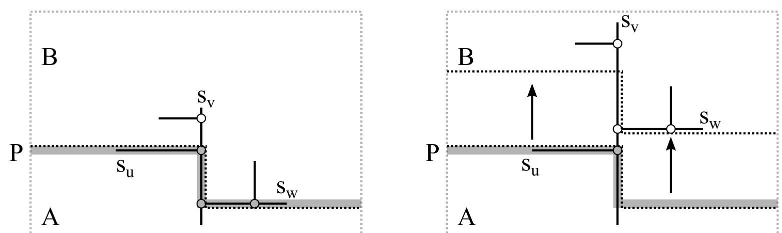

Figure 8: The slicing operation.

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 **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 **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 **C** that do not lie in the interior of P the *outside of* **C**. See Fig. 9 for an example.

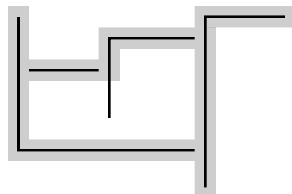

Figure 9: A contact-L-Γ representation of a 4-cycle. Its outside is highlighted.

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 **C** is empty. We again apply induction to $G_2 = G[C \cup V_C]$ with respect to the representation **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 **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 **G** of G. See Fig. 10 for an example.

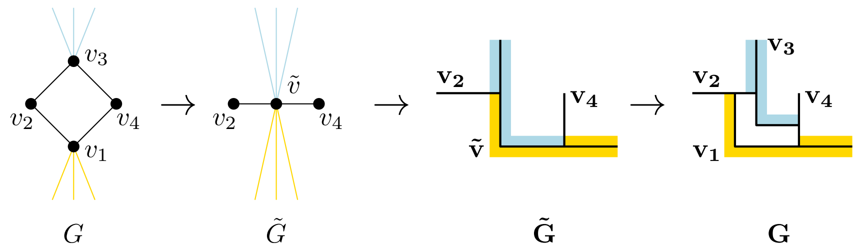

Figure 10: How to split a face in *Case 2*.

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.

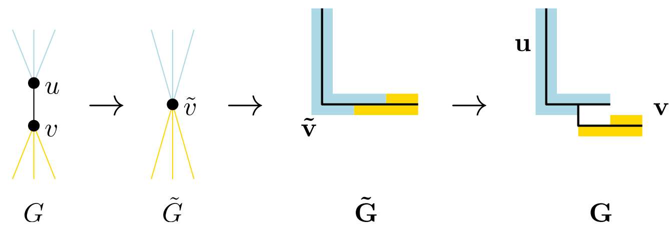

Figure 11: How to split an edge in *Case 3a*.

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.

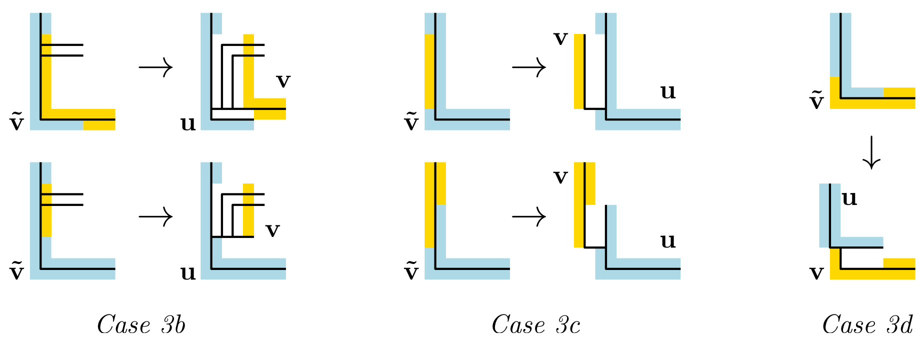

Figure 12: How to split an edge in *Case 3b*, *Case 3c*, and *Case 3d*.

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.