

Geometric Representations of Planar Graphs

vorgelegt von
M. Sc.-Mathematikerin
Nieke Aerts
aus Riel

Von der Fakultät II – Mathematik und Naturwissenschaften
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
– Dr. rer. nat. –

genehmigte Dissertation

Promotionsausschuss

Vorsitzender: Prof. Dr. Jörg Liesen
Berichter: Prof. Dr. Stefan Felsner
Prof. Dr. Alexander Wolff

Tag der wissenschaftlichen Aussprache: 23 Januar 2015

Berlin 2015
D 83

Voor mama

Acknowledgements

After almost a year in Australia, Berlin felt pretty close to home to me. After three years in Berlin, it sometimes appears not to be as close as I want it to be. The less-than-hour-airbridge to Eindhoven was a welcome surprise in 2012. It has made the visits (both ways) a lot more convenient. The first thank you should therefore go to the people responsible for the mistakes around the Berlin Brandenburg airport. Without Tegel airport and the TXL bus, my life at Beusselstrasse would have been completely different.

Stefan, thank you for giving me the opportunity to join your workgroup. I have enjoyed the atmosphere, and the freedom of choosing the problems on which I have worked. During the last three years, you have been a patient and thorough teacher, and this thesis definitely would not have existed without the many discussions we have had. Many of my friends have been jealous of how lucky I have been with the choice of my professor, apparently it is not standard that a professor always makes time for his students and I really appreciate that you do.

Thanks to all the (former) members of the workgroup. It has been a pleasure going to the office, discuss, or talk about other things during lunch and cake-Fridays. In particular thanks to Irina, Linda, Tom and Udo for the many discussions about work. Udo, thanks also for the many football evenings in the Oscar Wilde.

Then there are many people who have made my life outside work very enjoyable. Sometimes, an evening away from it all can do more for your productivity than you can imagine. In particular I would like to thank Loes and Karlijn, for always being there, just for everything. And Fanni, it has been great to have you around all this time. It is funny how we barely knew each other while we were living in Eindhoven, and by now I cannot imagine not knowing you. Thank you, for listening, for the distractions, the hilarious conversations, the introduction at Hansa, and a whole lot more. Thank you, Steve, for making life easier and more fun for me, and for your everlasting patience and support.

Papa, Marjon, Joep en Ruud, en natuurlijk ook de aanhang en kinderen, het was me een genoegen jullie, en mama, te verwelkomen in Berlijn. Dank jullie wel voor de afleiding, de steun en het vertrouwen, niet alleen, maar toch vooral ook heel erg, in de afgelopen maanden. Lieve mama, was je nog maar hier. Ik hoop dat je meekijkt en voelt hoe dankbaar ik je ben, voor alles.

*Nieke Aerts
Berlin, November 2014*

Contents

1	Introduction and Preliminaries	1
1.1	Orderings and Orientations	2
1.2	Sparse and Tight Graphs	4
1.3	Spanning Tree Decompositions	5
1.4	Schnyder Woods	8
1.4.1	Drawings Based on Schnyder Woods	12
2	Straight-Line Triangle Representations	17
2.1	Combinatorial Characterization	18
2.1.1	Good Flat Angle Assignments	25
2.1.2	Schnyder Labelings and Flat Angle Assignments	30
2.2	Applications	46
2.2.1	Graphs with few Schnyder labelings	46
2.2.2	Primal-Dual Triangle Contact representation.	47
2.2.3	Planar Generic Circuits	51
2.3	A Flow Network for Corner Compatible Pairs	54
2.4	Conclusion	61
3	Touching Triangle Representations	65
3.1	Triangle Representations (Known Results)	66
3.2	Results on 3TTRs	72
3.2.1	Halin Graphs	73
3.2.2	Biconnected Outerplanar Graphs	75
3.3	Conclusion	93
4	Grid-Paths Contact Representations	96
4.1	A Combinatorial Characterization of VCPGs	98
4.1.1	A ≤ 2 -Orientation representing Edges	101
4.1.2	A Feasible Flow representing Bends	102
4.1.3	Realizable pairs	105


4.1.4	Realizable Pairs are in Bijection with VCPGs	107
4.2	Bounding the Number of Bends	115
4.2.1	B_1 -VCPGs	116
4.2.2	Constructing a strict B_1 -VCPG	117
4.2.3	The class B_1 -VCPG	119
4.2.4	B_2 -VCPGs	122
4.2.5	B_k -VCPGs for $k > 2$	123
4.2.6	Obtaining Better Bounds	124
4.2.7	Local Flow Decreasing Steps	124
4.2.8	Constructive Argument for (2,2)-tight graphs	126
4.3	Conclusion	128
5	List of Open Problems	131
	Bibliography	134
	Samenvatting	140
	Zusammenfassung	141
	Curriculum Vitae	142

A graph is a set of vertices (objects) together with a set of edges (relations between the objects). There are different ways to represent a graph, divided into two types: combinatorial representations and geometric representations. A combinatorial representation can be an adjacency matrix, which states the neighbors of each vertex, or a list of edges, or a list of cycles, and so on. A combinatorial representation is convenient for example when one needs to let a computer investigate some properties of the graph. A geometric representation is a drawing of the graph. A drawing can display some properties of the graph in a nice way. For example when one wants to know whether two vertices are neighbors it may be more convenient to see if there is a connection between the two in a drawing than searching the adjacency lists.

In general, the start of graph theory is contributed to Euler, when he published the solution to the problem of the bridges of Königsberg in 1736. Remarkably, Euler did not use drawings to explain his solution. There are early examples of drawings of graph-like structures, however, graph drawings appear in the context of graph theory only at the end of the 18th, beginning of the 19th century [KMBW02].

In this thesis we will consider different geometric representations of planar graphs, different types of drawings. There are many types of geometric representations, identified for example by the objects used, e.g., points, Jordan curves, polygons, and so on.

In 1962 it was shown by Tutte that every 3-connected planar graph has a convex drawing [Tut63]. In such a drawing the vertices are represented by points in the plane, the edges are represented by straight lines between two points and, every face is a convex polygon. Such a drawing is sometimes referred to as rubber band representation; one can imagine that all the edges are rubber bands, the vertices of one cycle of the graph are pinned onto the plane in convex position. The equilibrium that is formed by the rubber bands is a convex straight-line drawing of the graph.




Another visualization that received a lot of attention over the years is called *intersection representation*. Here, the vertices are objects in the plane, e.g., strings, and the edges are represented as intersections between two objects. A special type of intersection representation is a *contact representation*, i.e., the edges are contacts (instead of intersections) of the two objects. It is clear that a graph that admits a contact representation of objects in two dimensions must be planar. A beautiful result in this area is the Koebe-Andreev-Thurston circle packing theorem [Koe36]: Every planar graph has a representation where the vertices are circles in the plane and the edges are represented as contacts between two circles. Blowing up the circles slightly shows that every planar graph also admits a circle intersection representation. Cutting each blown-up circle shows that every planar graph has a string intersection representation.

Two types of questions arise frequently in the theory of graph drawing. Firstly, for a given type of drawing we would like to know which graphs can be represented and the other way around, for a given graph class, which representations do the graphs in this class admit. Secondly, certain aspects of the drawings are considered. For example, for a grid drawing, i.e., each vertex is drawn on a grid point, one may ask what is the size of the grid needed to represent any graph of a certain class. Or what is the number of slopes needed for a certain representation. A drawing with the low ‘complexity’ with respect to some parameter, could make it easier to quickly see the properties of the graph. In Chapter 2 and 3 we investigate which graphs admit a certain drawing. In Chapter 4 we investigate the complexity of a drawing with grid-paths, i.e., we try to minimize the number of bends.

Thesis Outline

In Chapter 1 basic concepts as well as some more elaborate concepts are introduced.





In Chapter 2 we consider a drawing in the classical setting, i.e., vertices are points in the plane and edges are straight-lines between two points. We ask which graphs admit a straight-line drawing such that all the faces are triangles, this is denoted by *Straight-Line Triangle Representation*. To the best of our knowledge, it is still open whether, given a planar graph G , the question ‘Does G admit a straight-line triangle representation?’ belongs to P . We will give two characterizations of graphs that admit a straight-line triangle representation. However, we are not aware of an efficient way to check whether a given graph satisfies the conditions. On the positive side, we identify several classes of graphs for which all graphs admit a straight-line triangle representation.

Parts of this chapter have appeared in [AF13b, AF13a, AF].

In Chapter 3 the question is considered in its dual form:

we now ask for a representation in which the vertices are triangles and the edges are side-contacts between two triangles. Such a representation is called a *Touching Triangle Representation*. We have been able to characterize the biconnected outerplanar graphs that admit a touching triangle representation in a convex polygon. Secondly, we show that every Halin graph admits a touching triangle representation in a triangle.

Part of this chapter has appeared as a poster at GD2014 [Aer14].



In Chapter 4 the objects that represent vertices are grid-paths and we study contact representations of paths in a grid. The class of graphs that admit a grid-path contact representation is precisely the class of planar $(2,0)$ -sparse graphs. We will give a combinatorial characterization of grid-path contact representations. The interesting question about this representation is whether we can minimize the number of bends, locally as well as globally. Using the combinatorial characterization, we give bounds on the number of bends per vertex, that suffice to represent $(2,\ell)$ -tight planar graphs for different values of ℓ .

Part of this chapter has appeared in [AF14].

In Chapter 5 we will give the current status of the problems addressed in this thesis and the open problems that we have encountered along the way.

*Noem mij, noem mij, spreek mij aan,
o, noem mij bij mijn diepste naam.
Voor wie ik liefheb, wil ik heten.*

Neeltje Maria Min, Voor wie ik liefheb wil ik heten

1

Introduction and Preliminaries

Throughout the thesis a basic knowledge of graph theory is assumed. General notations are used, e.g., $G = (V, E)$ denotes a graph G with vertex set V and edge set $E \subseteq V \times V$. Edges are pairs of vertices and an edge that connects u and v is denoted by uv . In the undirected case uv and vu are the same, in the directed case uv is the edge with u as tail and v as head. The cardinality of a set X is denoted by $|X|$. When we speak about the faces of a graph, we assume that a crossing-free (topological) embedding of the graph is given. The set of faces is often denoted by F . The unbounded face is called outer face. An embedded graph uniquely defines cyclic orders of the neighbors of a vertex. The set of all these cyclic orders is called a rotation system. Two topological embeddings with the same rotation system have the same combinatorial embedding. Such topological embeddings are often considered to be equivalent.

A graph is called k -connected when the removal of any set of k vertices does not disconnect the graph. For a plane graph, i.e., a planar graph with a fixed embedding, there is another notion of connectivity.

Definition 1.1 (Internally k -connected). A plane graph G is said to be internally k -connected when the addition of a new vertex v_∞ in the outer face that is made adjacent to all the boundary vertices of G results in a k -connected graph.

In the remainder of this chapter we will introduce some structures that play an important role throughout this thesis.

1.1 Orderings and Orientations

An ordering of the vertices of the graph can be useful in graph drawing. Instead of having to draw the whole graph at once, using a vertex order, it is often possible to introduce the vertices one by one.

A *canonical order* is a vertex order with some special properties. This order has proven to be very useful for compact grid drawings. An example is depicted in Figure 1.3 (c). Canonical orders were first introduced for triangulations by de Fraysseix, Pach, and Pollack [dFPP90]. Canonical orders can be used to obtain compact grid drawings of graphs. A drawing algorithm that uses canonical orderings is presented in [dFPP90]. The algorithm embeds a maximal planar graph on n vertices on a $(2n - 4) \times (n - 2)$ grid using straight-line edges. Kant generalized canonical orderings to 3-connected graphs and also gives an algorithm to construct a straight-line convex grid embedding on a $(2n - 4) \times (n - 2)$ grid [Kan96].

Definition 1.2 (Canonical Order). Let $G = (V, E)$ be a plane, 3-connected graph and v_1 a vertex on the outer face. Let $\pi = (V_1, \dots, V_k)$ be an ordered partition of the vertices of G , i.e., $V_1 \cup \dots \cup V_k = V$ and $V_i \cap V_j = \emptyset$ for $i \neq j$. Define G_i , for $i = 1, \dots, k$, to be the subgraph of G induced by $V_1 \cup \dots \cup V_i$ and C_i to be the boundary of G_i . Then, π is a canonical order of G if:

- $V_1 = \{v_1, v_2\}$, where v_2 is a neighbor of v_1 on the outer face of G .
- $V_k = \{v_n\}$, where v_n lies on the outer face of G , v_n is a neighbor of v_1 and $v_n \neq v_2$.
- For $i > 1$, C_i is a cycle containing the edge $v_1 v_2$.
- Every G_i is internally 3-connected.

- For $i = 2, \dots, k-1$ one of the following holds:
 - $V_i = \{x\}$ and $x \in C_i$, or,
 - $V_i = \{y_1, \dots, y_j\}$ is a path in G such that y_1 and y_j each have one neighbor on C_{i-1} , these are the only neighbors of V_i in G_{i-1} and every vertex in V_i has at least one neighbor in $G - G_i$.




Figure 1.1: Two options for V_i : $V_i = \{u\}$ or $V_i = \{w_1, w_2, w_3, w_4\}$.

Chrobak and Kant improved the drawing algorithm of Kant to construct a straight-line convex grid embedding on a $(n-2) \times (n-2)$ grid [CK97].

The drawing algorithm based on canonical orders is often denoted as ‘shift-method’. Vertices are added to the drawing one by one, following the canonical order. The vertices that are already drawn, are shifted to the outside with respect to the newly drawn vertex. Along the construction, every boundary edge except for v_1v_2 has slope 1 or -1 . A clear description of the algorithm of de Fraysseix, Pach and Pollack can be found in Chapter 4 of [NR04].




Figure 1.2: Introducing vertex v_k according to the shift method. The leftmost and rightmost neighbor of v_k are denoted by l and r . All vertices left of l and l are shifted 1 unit to the left, every vertex right of r and r is shifted 1 unit to the right. Again the boundary edges (except v_1v_2) all have slope 1 or slope -1 .

An orientation of an undirected graph G is an orientation of all the edges in G . Restrictions on the number of outgoing arcs and incoming arcs are captured by α -orientations. Let $\alpha: V \rightarrow \mathbb{N}$ be a function from the set of vertices to the non-negative integers. An orientation is an α -orientation if every vertex v has outdegree precisely $\alpha(v)$. When α is a constant function, i.e., $\alpha(v) = c$ for some constant c and all vertices v , then the orientation is called a c -orientation. The number of edges induced by a set of vertices plays an important role in deciding whether a graph admits an α -orientation.




Figure 1.3: An α -orientation of the cube for a non-constant function α (a) and for the octahedron with constant $\alpha = 2$ (b). A graph with a canonical ordering in (c). This graph does not have a 2-orientation as it has 7 vertices and 15 edges.

1.2 Sparse and Tight Graphs

Bounds on the number of edges are highly related to the drawings the graph admits. For example, if the vertices are represented by segments and the edges by proper contacts between two segments (see Figure 1.4), then each segment can represent at most two edges, hence, the number of edges is bounded by twice the number of vertices. The following definition captures such properties.

Definition 1.3 (Sparse and Tight Graphs). Let $G = (V, E)$ be a graph and $k, l \geq 0$ integers. G is (k, l) -sparse if

$$\forall W \subseteq V : |E_W| \leq k|W| - l$$

where E_W is the set of edges induced by W and if $k > 0$ then W must have cardinality at least $\lfloor l/k \rfloor$.

G is (k, l) -tight if G is (k, l) -sparse and

$$|E| = k|V| - l$$

A subset U of the vertices of a (k, l) -sparse or tight graph is called a *critical set* if it induces $k|U| - l$ edges.

Planar graphs are $(3,6)$ -sparse. Therefore, a graph admits a drawing in the plane without crossing edges only if it is $(3,6)$ -sparse. However, not all $(3,6)$ -sparse graphs are planar, e.g., the complete bipartite graph $K_{3,3}$ is $(3,6)$ -sparse. Recently it has been shown [CGO10] that every planar graph has a 1-string representation, i.e., an intersection representation by strings such that each pair of strings crosses at most once. Even stronger, Chalopin and Gonçalves [CG09] have shown that every planar graph has a intersection representation by segments.

Outerplanar graphs have a drawing such that all vertices are on the outer face. They are $(2,3)$ -sparse. Not every $(2,3)$ -sparse graph is outerplanar, e.g. a wheel is $(2,3)$ -sparse. Not every $(2,3)$ -sparse graph is planar, e.g. the full subdivision of K_5 is $(2,3)$ -sparse.

Triangle-free graphs are precisely the $(2,4)$ -sparse graphs. It has been shown that planar triangle-free graphs have a segment contact representation by segments in at most three directions [dCCD⁺02]. This class includes the bipartite graphs. Bipartite planar graphs have a segment contact representation by segments in at most two directions (e.g. [dFdMP94]). Not all $(2,4)$ -sparse graphs are bipartite, as odd cycles of size at least five are $(2,4)$ -sparse.




Figure 1.4: A segment representation of C_5 with segments in 3 directions, of C_6 with segments in two directions and of the prism with three directions. On the right an L-contact representation.

A plane graph can be seen as a bar and joint framework. A framework is *flexible* if there is a continuous deformation from one drawing to another while preserving the length of every edge. Minimally rigid generic frameworks in two dimensions are precisely the $(2,3)$ -tight

graphs [Lam70]. A framework is minimally rigid if it is not flexible and the removal of one edge makes it flexible. It is generic if the rigidity does not depend on the chosen lengths of the edges. The result is generally contributed to Gerard Laman, and $(2,3)$ -tight graphs are also called *Laman Graphs*. Long before the characterization due to Laman, Henneberg studied the minimally rigid generic frameworks in two dimensions. He characterized this class by two construction steps [Hen11].

Theorem 1.4 (Henneberg, 1911). *Every $(2,3)$ -tight graph can be constructed from a single edge by the following two construction steps.*

- *Henneberg type 1 step: Add a vertex x and connect it to two (different) vertices in the current graph.*
- *Henneberg type 2 step: Subdivide an edge uv and connect the new vertex x to a third vertex in the current graph, not u or v .*

In Figure 1.5 this construction is visualized.




Figure 1.5: Three Henneberg type 1 steps, followed by a Henneberg type 2 step. The vertex that is introduced in a step is labeled x .

The graphs that admit a segment contact representation are precisely the graphs that are $(2,3)$ -sparse and planar, as was shown by de Fraysseix and Ossona de Mendez [dFdM07b, Corollary 20]. It is easy to see that no graph with more edges can have such a representation: each vertex is represented by a segment and each edge by a proper contact, if there are n vertices, there can be at most $2n$ edges. Moreover, in the outer face there have to be at least three ends of segments that do not represent an edge, hence, the number of edges is at most $2n - 3$.

These graphs also have a so-called *L-contact representation*, i.e., each vertex is represented by an axis-aligned path with precisely one 90° bend, and each edge is represented by a proper contact between two paths [KUV13] (see Figure 1.4).

1.3 Spanning Tree Decompositions

A spanning tree of a graph G is a tree that contains all the vertices of G . A set of subgraphs of G is a decomposition of G if every edge of G appears in precisely one of the graphs. A *spanning tree decomposition* is a set of spanning trees that is a decomposition of the graph.

Every $(1,1)$ -tight graph is a tree. Nash-Williams showed that every $(2,2)$ -tight graph has a decomposition into two spanning trees [NW64]. Nash-Williams also showed that every (k,k) -sparse graph has a decomposition into k edge-disjoint forests. From this it follows that every $(2,3)$ -tight graph has a decomposition into a spanning tree and a spanning forest consisting of two trees. This can also be treated as a pair of spanning trees where precisely one edge of the graph appears in both trees.

For (2,3)-tight graphs a spanning tree decomposition can be obtained from the Henneberg construction. The start edge is the edge that appears in both trees. The two edges of a Henneberg type 1 step are each added to a different tree. The edges introduced by a subdivision of a Henneberg type 2 step replace the original edge in one of the trees, the other new edge is added to the other tree (see Figure 1.6). Note that this also holds for non-planar (2,3)-tight graphs.




Figure 1.6: Obtaining a spanning tree decomposition of a (2,3)-tight graph using the Henneberg construction.

Every (3,3)-tight graph has a decomposition into three spanning trees [NW64]. Every (3,6)-tight graph has a decomposition into three spanning forests. The planar (3,6)-tight graphs are also known as *triangulations* or *maximal planar graphs*, and their edges can be decomposed into three trees such that each tree spans all interior vertices and precisely one of the three boundary vertices. The decomposition into three trees is well known under the name *Schnyder wood*, which we will discuss in the next section.

A *separating decomposition* is the decomposition of a maximal planar bipartite graph into two spanning trees. A maximal planar bipartite graph on n vertices has $2n - 4$ edges and is also denoted by *quadrangulation*.

Definition 1.5 (Separating Decomposition). Let $Q = (V_1 \cup V_2, E)$ be a plane quadrangulation and $s, t \in V_1$ two vertices on the boundary of the outer face Q . A separating decomposition is an orientation and coloring of the edges of Q such that:

1. All edges incident to s are incoming and red, and all edges incident to t are incoming and blue.
2. Every vertex $v \neq s, t$ has precisely one red and one blue outgoing arc. If $v \in V_1$, then, around v in clockwise order, there is a blue outgoing edge, zero or more incoming red edges, a red outgoing edge and zero or more incoming blue edges. If $v \in V_2$, then, around v there is a red outgoing edge, zero or more incoming red edges, a blue outgoing edge and zero or more incoming blue edges.

Separating decompositions have been studied thoroughly and we refer the reader to [dFdM01] for more on the relations between orientations and tree decompositions, [FHKO10] for more on binary labelings and more Schnyder-like properties. Every planar bipartite graph has a segment contact representation by horizontal and vertical segments and these representations are in bijection with separating decompositions.

Theorem 1.6 ([1]). *Separating decompositions of a quadrangulation Q are in bijection to segment contact representations with horizontal and vertical segments of Q .*

To go from a segment contact representation to a separating decomposition is easy. Let Q be a quadrangulation and R a segment contact representation of Q . Let s be represented by the bottommost horizontal segment and t by the topmost horizontal segment, let the ends of s and t be free, as in Figure 1.7. A coloring and orientation of the edges of Q is obtained by setting (see Figure 1.7):

- $u \rightarrow v$ if the segment that represents u ends on the segment that represents v , the color is blue if u is vertical and v is above u , or if u is horizontal and v is left of u , the color is red otherwise.

Every segment that is not the topmost or bottommost horizontal segment, ends on two other segments. Therefore, each vertex $v \neq s, t$ has outdegree two in the obtained orientation. Each edge is represented by one segment ending on another, hence, each edge is oriented. Clockwise around a vertex that is represented by a horizontal segment, there is an outgoing blue edge, zero or more incoming red edges, an outgoing red edge and zero or more incoming blue edges. Clockwise around a vertex that is represented by a vertical segment, there is an outgoing blue edge, zero or more incoming blue edges, an outgoing red edge and zero or more incoming red edges. The vertex s has only incoming red edges and the vertex t has only incoming blue edges. The obtained orientation and coloring is a separating decomposition of Q .




Figure 1.7: From a segment contact representation to a separating decomposition.

We will not give a proof of the other direction, however, an algorithm to construct a segment contact representation from a 2-orientation will be needed in Chapter 4. We will show how to construct a segment contact representation from a separating decomposition that uses a bipolar orientation. This method is mainly taken from [dFdM01]. There are other methods known, which do not use a bipolar orientation, for example via a 2-page book embedding [dFdMP95, FFNO11].

A theorem of de Fraysseix and Ossona de Mendez is the first step.

Theorem 1.7 ([dFdM01]). *The 2-orientations of a quadrangulation Q are in bijection with separating decompositions of Q .*

Here a 2-orientation is an orientation of the interior edges of the quadrangulation such that each interior vertex has outdegree precisely 2. Let s and t be two non-adjacent vertices on the boundary of Q . Then the 2-orientation can be extended with the orientation of the four boundary edges such that s and t have only incoming edges. The coloring of the edges follows from the choice of s and t , all edges at s are colored red, all the edges at t are colored blue. The other edges can be colored iteratively using the two rules of a separating decomposition.

Definition 1.8 (Bipolar Orientation). Let Q be a quadrangulation and V_1 and V_2 the two color classes of Q . The graph that consists of V_i and the diagonals that connect two vertices of V_i in every face of Q , is denoted by D_i . Let s_i and t_i be the vertices in V_i that are on the outer face of Q . A bipolar orientation is an orientation of D_i such that:

- The orientation is acyclic, and
- the vertex s_i is the only source and t_i is the only sink.

From a separating decomposition to a bipolar orientation. Let Q be a quadrangulation with a separating decomposition. Let V_1 be one of the color classes of Q and s_1 and t_1 the vertices of V_1 that are on the outer face. The edges of D_1 are oriented such that:

- All edges at s_1 are outgoing.
- All edges at t_1 are incoming.
- At every vertex $v \in V_1$, $v \neq s_1, t_1$, the incoming and outgoing edges are separated by the outgoing edges of v in the separating decomposition.

Similarly a bipolar orientation of D_2 can be constructed. An example is depicted in Figure 1.8.




Figure 1.8: From a separating decomposition of Q to the bipolar orientations of D_1 and D_2 . The segment contact representation is constructed using two orders that are obtained from the bipolar orientations of D_1 and D_2 .

From a bipolar orientation to a segment contact representation. From the bipolar orientations of D_1 and D_2 , a vertex order is constructed for the vertices in V_1 and V_2 . The order is such that each arc in the bipolar orientation is oriented from the lower vertex to the higher vertex. The vertex order is obtained by subsequently taking out vertices that have no incoming arcs. Let h_1, \dots, h_l be the resulting order of the vertices in V_1 and let v_1, \dots, v_k be the resulting order of the vertices in V_2 . For a vertex x of Q , let $M(x)$ be the maximum index over the indices of the neighbors of x in Q , and let $m(x)$ be the minimum index over the indices of the neighbors of x . The segment contact representation is obtained by drawing the following segments:

- For $i = 1, \dots, l$ draw a horizontal segment h_i from $(m(h_i), i)$ to $(M(h_i), i)$.
- For $j = 1, \dots, k$ draw a vertical segment v_j from $(j, m(v_j))$ to $(j, M(v_j))$.

Hence, a segment contact representation by axis-aligned segments can be constructed from a 2-orientation.

1.4 Schnyder Woods

Schnyder woods were introduced by Walter Schnyder in the context of the order dimension of planar graphs [Sch89]. A Schnyder wood is an orientation and 3-coloring of the interior edges of a triangulation. In a second publication, Schnyder used this structure to obtain

compact straight-line drawings of planar graphs [Sch90]. Schnyder woods have since found many additional applications to various graph drawing models as well as to the enumeration and encoding of planar maps.

Definition 1.9 (Schnyder Wood). Let G be a triangulation with the vertices s_1, s_2, s_3 in clockwise order on the boundary of the outer face. A Schnyder wood is an orientation and labeling of the interior edges of G with the labels 1, 2, and 3 such that the following conditions are satisfied¹.

- [w1] Around an interior vertex in clockwise order there is one outgoing edge with label 1, zero or more incoming edges with label 3, one outgoing edge with label 2, zero or more incoming edges of label 1, one outgoing edge with label 3 and zero or more incoming edges with label 2.
- [w2] The vertex s_i has only incoming edges, all of which are labeled i .

It is convenient to relate three colors to the labels $\{1, 2, 3\}$. In this thesis label 1 will be represented by the color red, label 2 by the color green and label 3 by the color blue².




Figure 1.9: A canonical order of a triangulation, the steps of obtaining a Schnyder wood from this canonical order, and, the Schnyder wood induced by this canonical order.

The graph induced by the color i of the Schnyder wood is a tree that is rooted in s_i . If the two boundary edges $s_i s_{i+1}$ and $s_i s_{i-1}$ are added, then the tree is a spanning tree, i.e., all the vertices of the graph are in the tree. Hence, a Schnyder wood is a decomposition of a triangulation into three trees, such that each boundary edge appears in two of the trees.

Every canonical order induces a Schnyder wood (e.g. [Bre00, Section 4.2]). The vertices v_1 and v_2 of the canonical order are the green and the blue sink. The vertex v_n is the red sink. At the introduction of v_i , its edges that are now on the outer face are its green and blue outgoing edges, the edges in between will be red incoming edges. For a triangulation, a canonical order can be obtained from a Schnyder wood and a Schnyder wood can be obtained from a canonical order (e.g. [Uec13, Lemma 1.1.6]). Different canonical orders may induce the same Schnyder wood, hence, the two are not in bijection, see Figure 1.10). For 3-connected graphs the relation no longer holds, there are Schnyder woods that do not induce a canonical order (see Figure 1.11 taken from [BBC11]). Badent et al. [BBC11] introduce a variant of canonical orders denoted by *ordered path partitions* and show that these are in bijection with Schnyder woods.

¹The labels are considered in a cyclic structure, such that $(i - 1)$ and $(i + 1)$ are always well defined.

²A useful mnemonic for the relation between the labels and the colors is RGB, the name one of the standard color models.




Figure 1.10: Two canonical orders that induce the same Schnyder wood: the vertices labeled 4/5 can be labeled either way, both induce the Schnyder wood drawn.





Figure 1.11: A Schnyder wood that cannot be obtained from a canonical order.

Felsner generalized the theory of Schnyder to 3-connected planar graphs [Fel01]. Triangulations are (special) 3-connected planar graphs, therefore, the generalization of Schnyder woods to 3-connected planar graphs is sometimes denoted by *generalized Schnyder wood*. Throughout this thesis, we always consider the general case, therefore, we refer to generalized Schnyder woods, simply as Schnyder woods. Many of the following definitions were originally defined for triangulations by Schnyder. We will only introduce the generalized versions.

A 3-connected planar graph may have more than three vertices on its boundary. Three special vertices on the boundary will act as the roots of the trees of the Schnyder wood. These vertices are denoted by *suspensions* of the graph. A plane graph together with three suspensions is denoted by *suspended graph*.

Definition 1.10 (Generalized Schnyder Wood). Let G be a 3-connected plane graph with three suspensions s_1, s_2, s_3 in clockwise order on the boundary of the outer face. A Schnyder wood is an orientation and labeling of the edges of G with the labels 1, 2, and 3 such that the following four conditions are satisfied³:

- [W1] Each edge is either unidirected or bidirected. In the latter case the two directions have distinct labels.
- [W2] At each suspension s_i there is a half-edge into the outer face with label i . A half-edge is an edge with only one endpoint.
- [W3] Every vertex v has outdegree 1 in each label. Around v in clockwise order there is an outgoing edge of label 1, zero or more incoming edges of label 3, an outgoing edge of label 2, zero or more incoming edges of label 1, an outgoing edge of label 3 and zero or more incoming edges of label 2.
- [W4] There is no directed cycle in one color.

Another structure contributed to Walter Schnyder are Schnyder labelings. The generalization from triangulations to 3-connected graphs can also be found in [Fel01].

Definition 1.11 (Schnyder Labeling). Let G be a 3-connected plane graph with three suspensions s_1, s_2, s_3 in clockwise order on the boundary of the outer face. A Schnyder labeling is a labeling of the angles of G with the labels 1, 2, and 3 such that the following three conditions are satisfied.

³The labels are considered in a cyclic structure, such that $(i - 1)$ and $(i + 1)$ are always well defined.




Figure 1.12: The conditions of Schnyder woods and labelings.

- [L1] The labels of a face form in clockwise order: a nonempty interval of 1's, a nonempty interval of 2's and a nonempty interval of 3's.
- [L2] At suspension s_i the outer angles, divided by the half-edge, have labels $(i + 1)$ and $(i - 1)$ in clockwise order. The inner angles at s_i are labeled i . Around each non-suspension vertex the labels form, in clockwise order, a nonempty interval of 1's, a nonempty interval of 2's and a nonempty interval of 3's.

Schnyder labelings are in bijection to Schnyder woods [Fel04, Theorem 2.3]. A Schnyder wood can be obtained from a Schnyder labeling by labeling the edge outgoing from v with label i if it separates an angle at v that is labeled $i - 1$ from an angle labeled $i + 1$. On the other hand a Schnyder labeling can be obtained from a Schnyder wood by labeling all angles clockwise between the outgoing edge with label $i - 1$ and the outgoing edge with label i with label $i + 1$ (see Figure 1.13).




Figure 1.13: A Schnyder wood and Schnyder labeling that are in bijection.

The weak dual G^* of a plane graph G , has the interior faces of G as its vertices and two vertices in G^* are connected by an edge, if the corresponding faces in G , share a boundary edge. The dual Schnyder wood, i.e., a Schnyder wood of the weak dual of G , is in bijection with the primal Schnyder wood, as the following theorem suggests.

Theorem 1.12 (Proposition 3 [Fel04]). *Let G be a suspended graph, the following structures are in bijection:*

- The Schnyder woods of G ,
- The Schnyder woods of the (weak) dual G^* of G .

The bijection between the Schnyder labeling of the primal graph and the Schnyder labeling of the dual graph is very simple. An angle at a vertex has precisely one opposite angle at a face-vertex in the dual graph. The labels are copied from the angle at the vertex to the




Figure 1.14: A primal dual Schnyder wood of G and its dual G^* . The black vertices are the vertices of G , the white vertices are the vertices of G^* and the squares represent the edges. The arcs belonging to the Schnyder wood of G are dashed.

angle at the face-vertex in the dual graph, see Figure 1.15. The bijection between a Schnyder labeling and a Schnyder wood, also works for the dual graph. Following the labeling one can obtain the dual Schnyder wood from the primal Schnyder wood. In Figure 1.15 (c)




Figure 1.15: A primal and dual Schnyder labeling of K_4 , a primal and dual labeling of a face, and the labels around an edge.

two primal-dual labelings around an edge are shown. On precisely one end, at a vertex or a face-vertex, the same label appears twice. This relates to the undirected edge, and it follows that an edge is bidirected in one and unidirected in the other of the two Schnyder woods related to this labeling.

1.4.1 Drawings Based on Schnyder Woods

In a second publication, Schnyder used his theory in order to obtain compact straight-line drawings of planar graphs [Sch90]. In this section we discuss two drawings based on Schnyder woods. The first drawing is the drawing as defined by Schnyder, to obtain compact grid drawings.

A Schnyder wood consists of three spanning trees, rooted in the suspensions [Fel04, Corollary 2.5]. The trees define paths from a vertex to each root. The paths from a given vertex v to the three roots are disjoint, except for v . Therefore, we can speak about the number of interior faces between the path to suspension s_{i-1} and the path to suspension s_{i+1} . Associate to each vertex v a triple (v_1, v_2, v_3) where v_i counts the number of interior faces between

the outgoing $(i + 1)$ -colored path from v to s_{i+1} and the outgoing $(i - 1)$ -colored path from v to s_{i-1} in the Schnyder wood. For each vertex it holds that $v_1 + v_2 + v_3 = |F| - 1$. A compact straight-line drawing can be obtained using this vector. Let $\alpha_1 = (0, 1)$, $\alpha_2 = (1, 0)$ and $\alpha_3 = (0, 0)$. A drawing obtained by *face-counting* is a mapping μ of the vertices to a 2-dimensional space. Two vertices that are adjacent are connected by a straight-line segment.

$$\mu : v \rightarrow v_1\alpha_1 + v_2\alpha_2 + v_3\alpha_3$$

Given a planar graph G and a Schnyder wood of G , the drawing D of G obtained by face-counting is plane and convex. [Fel01, Theorem 3]

We state a nice property of this drawing [BFM07, Lemma 3], the property is depicted in Figure 1.16.

Property 1.13. The vertices of an interior face are placed on the boundary of a triangle with sides on lines $c_i(\alpha_{i-1} - \alpha_{i+1})$ for some constant c_i . There are no vertices in the (open) interior of the bounding triangle of the face. The angles in a face, at the vertices on the line $c_i(\alpha_{i-1} - \alpha_{i+1})$ have label i in the Schnyder labeling.




Figure 1.16: A Schnyder wood and a drawing obtained by face-counting. An example of a triangle as in Property 1.13.

Another method of drawing a planar graph based on a Schnyder wood, is the so-called geodesic embedding (see Figure 1.17). Informally, in a geodesic embedding, the graph is embedded on the surface of a 3-dimensional object. Miller was the first to observe the connection between orthogonal surfaces in \mathbb{R}^3 and Schnyder woods [Mil02].

With a point $p = (p_1, \dots, p_d) \in \mathbb{R}^d$ associate its *cone* $C(p) = \{q \in \mathbb{R}^d : q_i \leq p_i \forall i = 1, \dots, d\}$ where p_i denotes the i -th coordinate of p . The *filter* $\langle \mathcal{V} \rangle$ generated by a finite set $\mathcal{V} \subset \mathbb{R}^d$ is the union of all cones $C(v)$ for $v \in \mathcal{V}$. The *orthogonal surface* $S_{\mathcal{V}}$ generated by \mathcal{V} is the boundary of $\langle \mathcal{V} \rangle$. A point $p \in \mathbb{R}^d$ belongs to $S_{\mathcal{V}}$ if and only if p shares a coordinate with all $v \leq p$, $v \in \mathcal{V}$. The generating set \mathcal{V} is an *antichain* if and only if all elements of \mathcal{V} appear as minima on $S_{\mathcal{V}}$. Figure 1.17 shows an example of an orthogonal surface with an embedded graph. The vertices of the graph are the elements of \mathcal{V} . Each vertex is incident to three ridges, we call them *orthogonal arcs*. The set of all orthogonal arcs of the surface yields the partition into plane patches, we call them *flats*. An *elbow geodesic* is a connection between two vertices u and v , it connects the two vertices with line segments on the surface to a saddle-point s of $S_{\mathcal{V}}$. One or both of the line segments forming an elbow geodesic are orthogonal arcs.

Figure 1.17 shows a geodesic embedding, in fact the geodesic embedding is decorated with the orientation and coloring of a Schnyder wood, precisely the Schnyder wood shown on the left of Figure 1.16.

Definition 1.14 (Geodesic Embedding). Let G be a plane 3-connected graph. A drawing of G onto an orthogonal surface $\mathcal{S}_{\mathcal{V}}$ generated by an antichain \mathcal{V} is a geodesic embedding if the following axioms are satisfied.

- [G1] There is a bijection between the vertices of G and the points in \mathcal{V} .
- [G2] Every edge of G is an elbow geodesic in $\mathcal{S}_{\mathcal{V}}$ and every bounded orthogonal arc in $\mathcal{S}_{\mathcal{V}}$ belongs to an edge in G .
- [G3] There are no crossing edges in the embedding of G on $\mathcal{S}_{\mathcal{V}}$.

Let G be a 3-connected plane graph with suspensions s_1, s_2, s_3 and let T be a Schnyder wood of G . Miller observed that there is an orthogonal surface S such that G has a geodesic embedding on S that induces T [Mil02], a sketch of the proof can be found in [FZ06, Theorem 5].

Together with the primal Schnyder wood, there is also a dual Schnyder wood embedded on the surface. Taking the maxima of S as vertices, we obtain a geodesic embedding of the weak dual G^* of G . Edges of G^* connecting to the vertex that represents the outer face, are the unbounded rays on the unbounded flats. The edges of the dual Schnyder wood use the orthogonal arcs that are incident to the maxima of the surface, whereas, the primal Schnyder wood uses the orthogonal arcs that are incident to the minima of the surface (see Figure 1.17 and Figure 1.18). The geodesic embedding of G^* is naturally decorated with colors and orientations. Adding one suspension for the unbounded rays of each color, yields a Schnyder wood T^* of the dual.




Figure 1.17: A geodesic embedding decorated with a Schnyder wood.




Figure 1.18: A geodesic embedding decorated with the dual Schnyder wood.

An orthogonal surface might support more than one graph, as shown in Figure 1.19. There-

fore, the Schnyder wood presented by the orthogonal surface is not unique.




Figure 1.19: The orthogonal surface on the left supports two graphs, the red outgoing arc of vertex 2 can be $2 \rightarrow 3$ or $2 \rightarrow 4$.

Let F be a flat that is constant in coordinate i . The saddle-points incident to F , can be divided into two types, those that are incident to an orthogonal arc with increasing i -coordinate, denoted by dual-saddles, and those that are incident to an orthogonal arc with decreasing i -coordinate, denoted by primal-saddles. The orthogonal arcs incident to dual-saddles belong to arcs in the dual Schnyder wood and the orthogonal arcs incident to the primal-saddles belong to arcs in the primal Schnyder wood.

Definition 1.15 (Rigid Orthogonal Surface). A flat F is called *rigid*, if every dual-saddle dominates at most one local maximum in F and every primal-saddle is dominated by at most one local minimum in F . An orthogonal surface is rigid if all its bounded flats are rigid.

The flat in Figure 1.20 has a dual-saddle that dominates two local maxima, and a primal-saddle that dominates two local minima. In the sense of Schnyder woods, consider Figure 1.19, the flat containing 3 and 4 is not rigid, as both vertex 3 and vertex 4 dominate the primal-saddle that belongs to the orthogonal arc to vertex 2. In other words, if a primal-saddle is dominated by two local minima, then there are two possibilities to extend the elbow geodesic from the orthogonal arc, hence, the graph on the surface is not uniquely defined.




Figure 1.20: Two flats. The left flat is not rigid as y is dominated by v_1 and v_2 (and x dominates f_1 and f_2). The flat on the right is rigid.




Figure 1.21: Two flats. The left flat is not rigid as the given path is not monotone in direction 2. The flat on the right is rigid.

Another view on rigid flats is obtained using the path that connects the saddles of the flat. A flat F , that is constant in coordinate i , is rigid, if and only if, there is a path P_F connecting all the saddle-points, such that P_F is monotone with respect to coordinates $i-1$ and $i+1$ (see Figure 1.21). The flat shown in the left part of Figure 1.21 is not-rigid, there is no path connecting all saddle-points, that is monotone with respect to coordinate $i+1$ and $i-1$.

It has been shown in [Fel03, Theorem 9] and [FZ06, Theorem 6] that every Schnyder wood has a geodesic embedding on some rigid orthogonal surface.

Theorem 1.16 ([Fel03, FZ06]). *Every Schnyder wood has a geodesic embedding on some rigid orthogonal surface.*

“There is no perfection only life”

Milan Kundera, *The unbearable lightness of being*

2

Straight-Line Triangle Representations

In this chapter we are dealing with representations of graphs in the classical setting, i.e., vertices are represented as points in the Euclidean plane and edges are segments connecting two vertices. The faces are represented as polygons bounded by as many segments as there are vertices in this face. If two consecutive segments bounding a polygon are drawn on the same line in the plane, the polygon can be seen as a polygon of lower degree. The polygon of lowest degree in the Euclidean plane is the triangle. Given a planar graph G the following question arises: *does G admit a planar drawing such that all the faces are triangles?*

Definition 2.1 (SLT Representation). A planar drawing of a graph such that:

- all the edges are straight-line segments, and,
- all the faces, including the outer face, bound a non-degenerate triangle

is called a *straight-line triangle representation* or in short *SLTR*.

A triangulation is a planar graph in which every face is a triangle. A straight-line drawing of a triangulation is an SLTR, therefore, the class of graphs admitting an SLTR is rich. A graph that admits an SLTR cannot have a cut-vertex since the outer face should also bound a triangle. However, being well connected is not sufficient as is shown by the cube graph. In this chapter we will investigate which graphs admit such a drawing. We consider the question: Given G , does G admit an SLTR?

In Section 2.1 we show a combinatorial description of an SLTR. From the combinatorial description, we obtain two necessary conditions (C_v, C_f) on the graph. These conditions are ‘easy to check’ but not sufficient. We introduce another necessary condition (C_o) and in Section 2.1.1 we show that this set of necessary conditions is also sufficient. However, we are not aware of an efficient way to check whether a given graph satisfies these conditions.

In Section 2.1.2 we identify another necessary condition (C_c) and show that condition C_o can be replaced by condition C_c . For a graph G , we can build a two-commodity network for which a feasible integral solution is equivalent to a drawing that satisfies conditions C_v, C_f and C_c . This network is introduced and investigated in Section 2.3. Note that the problem of deciding whether a two-commodity network has an integer feasible solution is NP-complete [EIS76]. Hence, it is still open if it is polynomially tractable to decide whether a graph has an SLTR.

In Section 2.2 we consider some graph classes for which we can prove that every graph in this class admits an SLTR. One of the results uses Henneberg type 2 steps. We will prove that an SLTR can be extended along such a step. The other results are obtained by using the necessary and sufficient conditions directly.

2.1 Combinatorial Characterization

In an SLTR the outer face has different properties than the interior faces: an interior face bounds a triangle, the outer face has a ‘hole’ that is a triangle. There are graphs for which there is no SLTR equivalent to a certain embedding, but there is an SLTR equivalent to another embedding of this graph. To simplify analysis we will consider the plane equivalent of the question ‘does G admit an SLTR?’, namely, ‘Given a plane graph G , does G admit an SLTR equivalent to this embedding?’. There are three vertices of G that will be identified with the corners of the triangular hole in the outer face in an SLTR. These vertices play a special role on many occasions. We call these three vertices *suspension vertices*. When three

vertices of a face are chosen to be the suspension vertices, we call the graph a suspended plane graph (the outer face is also defined by the choice of the suspensions).

We consider the following question: Given a plane suspended graph G , does G admit an SLTR equivalent to this embedding?

When the answer to this question is ‘no’ for all possible choices of suspensions of G then the answer to the more general question, ‘does G admit an SLTR?’, is also negative. When there is a suspension of G for which the answer to the above question is ‘yes’, then also the answer to the more general question is positive. Hence, for the remainder of this chapter we focus on the question: Given a plane suspended graph G , does G admit an SLTR equivalent to this embedding?

We proceed with identifying some properties of an SLTR.

Connectivity of an SLTR. Let \mathcal{R} be an SLTR of a graph G and let f be a face of G that is not of size three. Then in the drawing of f in \mathcal{R} there are three vertices, which are the corners, and the others have an angle of size π in f . If a degree-2 vertex has an angle of size π in one of its incident faces, then it also has an angle of size π in the face on the other side. Hence, this vertex and its two incident edges can be replaced by a single edge connecting the two neighbors of the vertex. Such an operation is called a *vertex reduction*. The only angles of an SLTR whose size exceeds π are the outer angles at the suspensions. Therefore, we can use vertex reductions to eliminate all degree-2 vertices except for the suspensions.

Recall that a plane graph G with suspensions s_1, s_2, s_3 is said to be *internally 3-connected* when adding a new vertex v_∞ in the outer face and making it adjacent to the three suspension vertices yields a 3-connected graph.

Proposition 2.2. *If a graph G admits an SLTR with s_1, s_2, s_3 as corners of the outer triangle and no vertex reduction is possible, then G is internally 3-connected.*

Proof. Consider the SLTR of G . Suppose that there is a separating set U of size 2. It is enough to show that each component of $G \setminus U$ contains a suspension vertex, so that $G + v_\infty$ is not disconnected by U . Since G admits no vertex reduction, every degree-2 vertex is a suspension. Hence, if C is a component and $C \cup U$ induces a path, then there is a suspension in C . Otherwise consider the convex hull of $C \cup U$ in the SLTR. The convex corners of this hull are vertices that expose angles of size at least π . Two of these large angles may be at vertices of U but there is at least one additional large angle. This large angle must be the outer angle at a vertex that is an outer corner of the SLTR, i.e., a suspension. \square

Counting vertices and faces. Consider a plane, suspended, internally 3-connected graph $G = (V, E)$. Suppose that G admits an SLTR. This representation induces a set of *flat angles*, i.e., incident pairs (v, f) such that vertex v has an angle of size π in the face f .

Since G is internally 3-connected every vertex has at most one flat angle. Therefore, the flat angles can be seen as a partial mapping of vertices to faces. Since the outer angle of the suspension vertices exceeds π , suspensions have no flat angle. Since each face f (including the outer face) is a triangle, each face has precisely three angles that are not flat. In other words every face f has all but three of its incident vertices assigned to f . The number of incident vertices of a face f is denoted by $|f|$. This motivates the following definition.

Definition 2.3 (FA Assignment). A *flat angle assignment* (FAA) is a mapping from a subset U of the non-suspension vertices to faces such that

- [C_v] Every vertex of U is assigned to at most one face,
- [C_f] For every face f , precisely $|f| - 3$ vertices are assigned to f .

An FAA is a combinatorial description of an SLTR. A graph that has no flat angle assignment has no SLTR, hence, having a flat angle assignment is a necessary condition. Checking whether a graph has an FAA can be done by a simple counting argument.

Proposition 2.4. Let $G = (V, E)$ be a suspended, plane, internally 3-connected graph and F the set of faces of G . Then G has an FAA if and only if:

$$\forall H \subseteq F : \sum_{f \in H} (\deg(f) - 3) \leq |VH| - 3, \quad (2.1)$$

where VH is the set of all vertices incident to a face in H .

Proof. Construct a bipartite graph with vertex classes W_1 and W_2 . The class W_1 contains $|f| - 3$ duplicates of f for every $f \in F$. The class W_2 contains all non-suspension vertices. A flat angle assignment is a matching such that all elements in W_1 are matched. By Equation 2.1 Hall's marriage condition is satisfied for all subsets of W_1 . Hence, there exists a matching such that all elements of W_1 are matched. We denote this by *one-sided-perfect matching*. If Equation 2.1 does not hold then there does not exist a one-sided-perfect matching for W_1 . It follows that there is no FAA and the plane graph does not admit an SLTR. \square

In the introduction of this chapter we already mentioned that the cube graph (see Figure 2.1) does not admit an SLTR. This can now be verified. Indeed Equation 2.1 does not hold when H is the set of all the faces of the cube graph.

Unfortunately, not every FAA induces an SLTR. An example is given in Figure 2.2. The angles at a, b and c , indicated by the arrows, must be of size π . Then these vertices must lie on a straight-line segment connecting s_3 and s_2 . Therefore, everything inside the cycle s_1, a, b, c, s_3 must lie on this segment. Intuitively, the third condition can already be seen here: To avoid degeneracies, every cycle must have at least three vertices whose angles inside the cycle are smaller than π . We proceed with a more formal introduction of the ‘cycles’ that we consider.




Figure 2.1: The cube graph does not have an FAA.




Figure 2.2: An FAA that does not induce an SLTR.

Let G be a suspended, internally 3-connected plane graph and let H be a connected subgraph of G . The *outline cycle* $\gamma(H)$ of H is the closed walk corresponding to the outer face of H (see Figure 2.3). An *outline cycle* of G is a closed walk that can be obtained as outer cycle of some connected subgraph of G . Outline cycles may have repeated edges and vertices. The interior $\text{int}(\gamma)$ of an outline cycle $\gamma = \gamma(H)$, consists of H together with all vertices, edges and faces of G that are contained in the area enclosed by γ .

Proposition 2.5. *An SLTR obeys the following condition C_o :*

[C_o] *Every outline cycle that is not the outline cycle of a path, has at least three geometrically convex corners.*

Proof. Consider an SLTR. Suppose that there is a connected subgraph that is not a path such that its outline cycle has less than three geometrically convex corners. If the outline cycle has at most two geometrically convex corners, then the subgraph is mapped to a line in the plane. The subgraph must either contain a vertex of degree more than three or a face, as it is not a path. If a vertex v together with three of its neighbors is mapped onto a line, then the boundary of at least one of the faces incident to v is not a triangle. If the subgraph contains a face, then this face is mapped to a line, and therefore, its boundary is not a triangle. In both cases the properties of an SLTR are violated. This shows that C_o is a necessary condition. \square

Condition C_o has the disadvantage that it depends on a given SLTR, hence, it is useless for deciding whether a plane graph G admits an SLTR. The following definition allows to replace C_o by a combinatorial condition on an FAA.

Definition 2.6 (Combinatorial Convex Corners). Let ψ be an FAA of G . A vertex v of an outline cycle γ is a *combinatorial convex corner* for γ with respect to ψ if:

- [A1] v is a suspension vertex.
- [A2] v is not assigned to a face and v is incident to an edge $e \notin \text{int}(\gamma)$.
- [A3] v is assigned to a face $f \notin \text{int}(\gamma)$ and v is incident to an edge $e \notin \text{int}(\gamma)$.




Figure 2.3: Outline Cycles.




Figure 2.4: Combinatorial convex corners.

Proposition 2.7. *Let G admit an SLTR R that induces the FAA ψ , and let H be a connected subgraph of G . If v is a geometrically convex corner of the outline cycle $\gamma(H)$ in R , then v is a combinatorial convex corner of $\gamma(H)$ with respect to ψ .*

Proof. If v is a suspension vertex, it is clearly geometrically and combinatorial convex.

Let v be geometrically convex and suppose that v is not a suspension and not assigned by ψ . In this case, v is interior and, with respect to γ , the outer angle at v exceeds π . Therefore, at least two incident faces of v are outside of γ . These faces can be chosen to be adjacent, hence, the edge between them is an edge e with $e \notin \text{int}(\gamma)$. This shows that v is combinatorial convex.

Let v be geometrically convex and suppose that v is assigned to f by ψ . If $f \in \text{int}(\gamma)$, then the inner angle of v with respect to γ is at least π . This contradicts the fact that v is geometrically convex. Hence, $f \notin \text{int}(\gamma)$. If there is no edge e incident to v such that $e \notin \text{int}(\gamma)$, then v has an angle of size π with respect to γ . This again contradicts the fact that v is geometrically convex. Therefore, if v is geometrically convex and assigned to f , then $f \notin \text{int}(\gamma)$ and there exists an edge e incident to v such that $e \notin \text{int}(\gamma)$. This shows that v is a combinatorial convex corner for γ . \square

The proposition enables us to replace the condition on geometrically convex corners w.r.t. an SLTR by a condition on combinatorial convex corners w.r.t. an FAA.

[C_o^{*}] Every outline cycle that is not the outline cycle of a path, has at least three combinatorial convex corners.

From Proposition 2.5 and Proposition 2.7 it follows that this condition is necessary for an FAA that belongs to an SLTR.

Later in Theorem 2.13 we prove that if an FAA obeys C_o^{*} then it belongs to an SLTR. In anticipation of this result we say that an FAA obeying C_o^{*} is a *good flat angle assignment* and abbreviate it as a *GFAA*.

Definition 2.8 (Contact family of pseudo-segments). A contact family of pseudo-segments is a family $(C_i)_i$ of simple curves

$$C_i: [0, 1] \rightarrow \mathbb{R}^2, \text{ with different endpoints, i.e., } c(0) \neq c(1),$$

such that any two curves C_i and C_j ($i \neq j$) have at most one point in common. If so, then this point is an endpoint of (at least) one of them.

A GFAA ψ on a graph G gives rise to a relation ρ on the edges: Two edges, both incident to v and f are in relation ρ if and only if v is assigned to f . The transitive closure of ρ is an equivalence relation.

Proposition 2.9. *The equivalence classes of edges of G defined by ρ form a contact family of pseudosegments.*

Proof. Let the equivalence classes of ρ be called arcs.

Condition C_v ensures that every vertex is interior to at most one arc. Hence, the arcs are simple curves and no two arcs cross.

Every arc has two distinct endpoints, otherwise it would be a cycle and its outline cycle would have only one combinatorial convex corner. If an arc touches itself, the outline cycle of this equivalence class has at most one combinatorial convex corner. This contradicts C_o^{*}.

If two arcs share two points, the outline cycle has at most two combinatorial convex corners. This again contradicts C_o^{*}.

We conclude that the family of arcs satisfies the properties of a contact family of pseudo-segments. \square




Figure 2.5: An arc with only one endpoint (a). An arc which touches itself (b). Two arcs that share two points (c).

Definition 2.10 (Free Point). Let Σ be a family of pseudosegments and S a subset of Σ . A point p of a pseudosegment from S is a free point for S if the following four conditions hold:

1. p is an endpoint of a pseudosegment in S .
2. p is not interior to a pseudosegment in S .
3. p is incident to the unbounded region of S .
4. p is a suspension or p is incident to a pseudosegment that is not in S .

With Lemma 2.11 below we prove that the family of pseudosegments Σ that arises from a GFAA has the following property¹:

$[C_P]$ Every subset S of Σ with $|S| \geq 2$ has at least three free points.




Figure 2.6: The contact family of pseudosegments is highlighted on an SLTR.




Figure 2.7: Free points (in green circles) and not free points (in red squares).

Lemma 2.11. Let ψ be a GFAA of a plane, internally 3-connected graph G and let S be a subset of the family of pseudosegments associated with ψ . If $|S| \geq 2$ then S has at least three free points.

Proof. Let S be a subset of the contact family of pseudosegments defined by the given GFAA (Proposition 2.9).

Each pseudosegment of S corresponds to a path in G . Let H be the subgraph of G obtained as the union of the paths of pseudosegments in S . We first assume that H is connected. If H itself is not a path, then by C_o^* the outline cycle $\gamma(H)$ must have at least three combinatorial convex corners. Every combinatorial convex corner of $\gamma(H)$ is a free point of S .

If S induces a path, then the two endpoints of this path are free points for S . Moreover, there exists at least one vertex v in this path which is an endpoint for two pseudosegments and not an interior point for any. Now v is a suspension or there must be an edge e incident to v such that $e \notin S$. Therefore, v is a free point for S .

If H is not connected then it has at least two components. Each component has at least two combinatorial convex corners and if these corners are incident to the unbounded region,

¹Note that this property is similar to the condition stated by De Fraysseix and Ossona de Mendez [dFdM07a]. For a more details on the similarity we refer to Section 2.1.1.

we are done. If one component of H , say H_1 , contains all other components in its interior, then only the combinatorial convex corners of the outermost component are incident to the unbounded region. However, if H_1 contains all other components then H_1 is not a path and, hence, H_1 has at least three combinatorial convex corners all incident to the unbounded region of H . This concludes the proof. \square

The proof that a GFAA induces an SLTR (Theorem 2.13) is constructive. In the remainder of this section we set up the proposed drawing of an internally 3-connected, suspended, plane graph G according to a GFAA of this graph. The proof is then completed by showing that this proposed drawing indeed is an SLTR of G (that induces the GFAA).

The proposed drawing. Given an internally 3-connected, suspended, plane graph G and a GFAA of G . To find a corresponding SLTR, we aim at representing every assigned vertex between its two neighbors along the assignment. This property can be modeled by requiring that the coordinates $p_v = (x_v, y_v)$ of an assigned vertex v of G satisfies a harmonic equation. Indeed if uv and vw are edges belonging to a pseudosegment s , then the coordinates satisfy

$$x_v = \lambda_v x_u + (1 - \lambda_v) x_w \quad \text{and} \quad y_v = \lambda_v y_u + (1 - \lambda_v) y_w \quad (2.2)$$

for some λ_v . In our model we can choose λ_v as a parameter from $(0, 1)$. With fixed λ_v the equations of (2.2) are the harmonic equations for v .

In the SLTR every unassigned vertex v is placed in the convex hull of its neighbors. In terms of coordinates this means that there are $\lambda_{vu} > 0$ with $\sum_{u \in N(v)} \lambda_{vu} = 1$ such that

$$x_v = \sum_{u \in N(v)} \lambda_{vu} x_u, \quad y_v = \sum_{u \in N(v)} \lambda_{vu} y_u. \quad (2.3)$$

We can choose the $\lambda_{vu} > 0$ arbitrarily subject to $\sum_{u \in N(v)} \lambda_{vu} = 1$. With fixed parameters the equations (2.3) enforce that v is located in the a weighted barycenter of its neighbors. These are the harmonic equations for an unassigned vertex v .

Vertices whose coordinates are not restricted by harmonic equations are called *poles*. In our case, the suspension vertices are the three poles of the harmonic functions for the x - and y -coordinates. The coordinates for the suspension vertices are fixed as the corners of some non-degenerate triangle.

The theory of harmonic functions and applications to (plane) graphs are nicely explained by Lovász [Lov09]. The proof of the following proposition is inspired by Lovász's proofs [Lov09].

Proposition 2.12. *Let $G = (V, E)$ be a directed graph, $\lambda : E \rightarrow \mathbb{R}^+$ be a weight function, and $P \subset V$ be a set of poles. If every subset Q of $V \setminus P$ has an out-neighbor in $V \setminus Q$, then for all $\psi_0 : P \rightarrow \mathbb{R}$ there is an extension $\psi : V \rightarrow \mathbb{R}$ which is harmonic on all $v \in V \setminus P$, i.e., $\psi(v) = \psi_0(v)$ for all $v \in P$ and $\psi(v) = \sum_{u \in \text{out}(v)} \lambda_{(v,u)} \psi(u)$ for all $v \in V \setminus P$.*

Proof. The proof has three steps. First we show that the maximum and minimum of a harmonic function are attained at poles. Then we show that for every map $\psi_0 : P \rightarrow \mathbb{R}$ from the set of poles to the reals, there exists a unique extension $\psi : V \rightarrow \mathbb{R}$ that is harmonic in all the vertices that are not poles. Last we show that a solution of the system of equations exists.

Let f be a non-constant harmonic function on G . Let $Q = \{v \in V : f(v) \text{ maximum}\}$ and $Q' = \{v \in Q : v \text{ has an out-neighbor not in } Q\}$. Since f is not constant, $Q \neq V$. Suppose Q does not contain a pole. From the connectivity assumption it follows that Q' is not empty. Elements of Q' are not harmonic and, hence, must be poles. This is a contradiction as $Q' \subseteq Q$. Therefore, Q must contain a pole. Similarly we find a pole among the vertices where the minimum is attained.

Consider $\psi_0: P \rightarrow \mathbb{R}$, a map from the set of poles to the reals and suppose there are two extensions $\psi, \psi^*: V \rightarrow \mathbb{R}$ that satisfy the harmonic equations of all non-poles. Then the function $\omega = \psi - \psi^*$ is also harmonic in all vertices which are not in P . As ψ and ψ^* are extensions of ψ_0 , the value of ω at all poles is zero. Since maximum and minimum of a harmonic function are attained at poles, we conclude that ω is zero everywhere, hence, $\psi = \psi^*$.

Prescribed values at poles together with the harmonic equations at non-poles, yield a linear system of n equations in n variables. From the uniqueness (of the extension) it follows that the homogeneous system has a trivial kernel, hence, by the rank-nullity theorem, the system has a unique solution for every $\psi_0: P \rightarrow \mathbb{R}$ prescribing the values for the poles. \square

To make use of Proposition 2.12 we need to show that a system of equations that comes from a GFAA, induces a directed graph and a weight function that satisfy the above properties. The vertices of the directed graph are the vertices of G . For a vertex v that is assigned and between u and w , we add the edges $v \rightarrow u$ and $v \rightarrow w$. For a not assigned vertex, we add a directed edge to each of its neighbors. The weights are given by the chosen parameters λ_v and λ_{vu} . The poles are the suspension vertices. To show that every subset Q of $V \setminus P$ has an out-neighbor in $V \setminus Q$, we consider the contact family of pseudosegments induced by the GFAA.

Suppose that there exists a set $Q \subseteq V \setminus P$ that has no out-neighbor in $V \setminus Q$. For every vertex $v \in Q$, if v is interior to a pseudosegment, the whole pseudosegment must be in Q . If v is not assigned then all of its neighbors must be in Q . Therefore, Q contains vertices of at least two pseudosegments. Moreover, Q is not the whole set, as $Q \subseteq V \setminus P$. Since the contact family of pseudosegments comes from a GFAA, the set of pseudosegments that are incident to Q , denoted by S_Q , must have at least three free points. A free point is on the boundary, not interior to any pseudosegment in S_Q and has at least one neighbor outside S_Q . If a free point is also in S , then, since this vertex has an out-neighbor in $V \setminus Q$, Q must have an out-neighbor in $V \setminus Q$. If a free points of S_Q is not in Q , then there must be a vertex $v \in Q$ that belongs to the pseudosegment to which this free point belongs, such that, v has an out-neighbor in $V \setminus Q$. Therefore, Q must have an out-neighbor in $V \setminus Q$.

2.1.1 Good Flat Angle Assignments

We are now ready to prove that the drawing given by the solution of the system of harmonic equations as defined before is indeed an SLTR if the flat angle assignment satisfies condition C_o^* . This shows that the conditions C_v, C_f, C_o^* are sufficient.

Theorem 2.13. *Given an internally 3-connected, plane graph G and a GFAA of G . The unique solution of the system of equations that arises from the GFAA is an SLTR.*

Proof. The proof consists of seven arguments, which together yield that the drawing induced from the GFAA is a non-degenerate, plane drawing. The proof has been inspired by a

proof for convex straight-line drawings of plane graphs via spring embeddings shown to us independently by Günter Rote and Éric Fusy. Both attribute key ideas to Éric Colin de Verdière.

1. *Pseudosegments become Segments.* Let $(v_1, v_2), (v_2, v_3), \dots, (v_{k-1}, v_k)$ be the set of edges of a pseudosegment defined by ψ . The harmonic conditions for the coordinates force that v_i is placed between v_{i-1} and v_{i+1} for $i = 2, \dots, k-1$. Hence, all the vertices of the pseudosegment are placed on the segment with endpoints v_1 and v_k .

2. *Convex Outer Face.* The outer face is bounded by three pseudosegments and the suspensions are the endpoints of these three pseudosegments. The coordinates for the suspensions (the poles of the harmonic functions) have been chosen as corners of a non-degenerate triangle and the pseudosegments are straight-line segments, therefore the outer face is a triangle and in particular convex.

3. *No Concave Angles.* Every vertex that is not a pole is forced either to be on the line segment between two of its neighbors (if assigned) or in a weighted barycenter of all its neighbors (otherwise). Therefore, every non-pole vertex is in the convex hull of its neighbors. This implies that these are no concave angles at non-poles.

4. *No Degenerate Vertex.* A vertex is degenerate if it is placed on a line, together with at least three of its neighbors. Suppose there exists a vertex v such that v and at least three of its neighbors are placed on a line l . Let S be the connected component of pseudosegments that are aligned with l such that S contains v . The set S contains at least two pseudosegments. Therefore, S must have at least three free points, which we denote by v_1, v_2, v_3 .

By property 4 in the definition of free points, each of the free points is incident to a segment that is not aligned with l . Suppose the free points are not suspension vertices. If v_i is interior to some pseudosegment s_i , then s_i has an endpoint on each side of l . If v_i is not assigned by the GFAA it is in the strict convex hull of its neighbors, hence, v_i is an endpoint of a segment reaching into each of the two half-planes defined by l .

Now suppose v_1 and v_2 are suspension vertices² and consider the third free point, v_3 . If it is interior to a pseudosegment not on l , but then one endpoint of this pseudosegment lies outside the convex hull of the three suspensions, which is a contradiction. Hence, it is not interior to any pseudosegment and at least one of its neighbors does not lie on l . Then v_3 should be in a weighted barycenter of its neighbors, hence, again we would find a vertex outside the convex hull of the suspension vertices. Therefore, at most one of the free points is a suspension and l is incident to at most one of the suspension vertices.

In every case, each of v_1, v_2, v_3 has a neighbor on either side of l .

Let n^+ and $n^- = -n^+$ be two normals for line l and let p^+ and p^- be the two poles, that maximize the inner product with n^+ respectively n^- (see Figure 2.8). Starting from the neighbors of v_i in the positive half-plane of l we can always move to a neighbor with larger³ inner product with n^+ until we reach p^+ . Hence, v_1, v_2, v_3 have paths to p^+ in the upper half-plane of l and paths to p^- in the lower half-plane. Since v_1, v_2, v_3 also have a path to v we can contract all vertices of the upper and lower half-plane of l to p^+ respectively p^- and all inner vertices of these paths to v to produce a $K_{3,3}$ minor of G . This is in contradiction to the planarity of G . Therefore, there is no degenerate vertex.

²Not all three suspension vertices lie on one line, hence, at least one of the three free points is not a suspension.

³If n^+ is perpendicular to another segment this may not be possible. In this case we can use a slightly perturbed vector n_ε^+ to break ties.




Figure 2.8: A vertex with three neighbors on a line.




Figure 2.9: Examples of vertices with their surrounding faces not oriented consistently.

5. *Preservation of Rotation System.* Let $\theta(v) = \sum_f \theta(v, f)$ denote the sum of the angles around an interior vertex. Here f is a face incident to v and $\theta(v, f)$ is the (smaller!) angle between the two edges incident to v and f in the drawing obtained by solving the harmonic system. If the incident faces are oriented consistently around v , then the angles sum up to 2π , otherwise $\theta(v) > 2\pi$ (see Figure 2.9). We do not consider the outer face in the sums so that the b vertices incident to the outer face contribute $(b-2)\pi$ in total.

Now consider the sum $\theta(f) = \sum_v \theta(v, f)$ of the angles of a face f . At each vertex incident to f the contribution $\theta(v, f)$ is at most of size π . A closed polygonal chain with k corners, selfintersecting or not, has a sum of inner angles equal to $(k-2)\pi$. Therefore, $\theta(f) \leq (|f|-2)\pi$. The sum over all vertices $\sum_v \theta(v)$ and the sum over all faces $\sum_f \theta(f)$ must be equal since they count the same angles in two different ways.

$$(|V| - b)2\pi + (b - 2)\pi \leq \sum_v \theta(v) = \sum_f \theta(f) \leq ((2|E| - b) - 2(|F| - 1))\pi$$

This yields $|V| - |E| + |F| \leq 2$. Since G is planar Euler's formula implies equality. Therefore, $\theta(v) = 2\pi$ for every interior vertex v and the faces must be oriented consistently around every vertex, i.e. the rotation system is preserved. Note that the rotation system may be flipped between clockwise and counterclockwise but then it is flipped at every vertex.

6. *No Crossings.* Suppose two edges cross. On either side of both of the edges there is a face, therefore, there must be a point p in the plane which is covered by at least two faces. Outside of the drawing there is only the unbounded face. Move along a ray, that does not pass through a vertex of the graph, from p to infinity. A change of the cover number, i.e. the number of faces by which the point is covered, can only occur when crossing an edge. But then the rotation system at the vertices of that edge must be wrong. This would contradict the previous item. Therefore, a crossing can not exist.

7. *No Degeneracy.* Suppose there is an edge of length zero. Since every vertex has a path to each of the three suspensions, there has to be a vertex a that is incident to an edge of length zero and an edge ab of non-zero length. Following the direction of forces, we can even find such a vertex-edge pair with b contributing to the harmonic equation for the coordinates of a . We now distinguish two cases.

If a is assigned, it is on the segment between b and some other vertex b' . Together with the neighbor of the zero-length edge this makes three neighbors of a on a line. Hence, a is a degenerate vertex. A contradiction.

If a is not assigned, it is in the convex hull of its neighbors. However, starting from a and using only zero-length edges, we eventually reach some vertex a' that is incident to an

edge $a'b'$ of non-zero length, such that b' is contributing to the harmonic equation for the coordinates of a' . Vertex a' has the same position as a and is also in the convex hull of its neighbors. This makes a crossing of edges unavoidable. A contradiction. Hence, there are no edges of length zero.

Suppose there is an angle of size zero. Recall that every vertex is in the convex hull of its neighbors and there are no interior angles of size larger than π . Moreover there are no crossings, hence, the face with the angle of size zero is stretching along a line segment with two angles of size zero. Since there are no edges of length zero and all vertices are in the convex hull of their neighbors, all but two vertices of the face must be assigned to this face. Therefore, there are two pseudosegments bounding this face, which have at least two points in common. This contradicts that Σ is a family of pseudosegments. We conclude that there is no degeneracy.

From the above arguments we conclude that the drawing is plane and thus an SLTR. \square

We obtained equivalence between the existence of an SLTR, the existence of an FAA satisfying C_v , C_f and C_o^* , and a stretchable system of pseudosegments that arises from this FAA.

For later use we will show that it is sufficient if all simple outline cycles have at least three combinatorial convex corners. A simple outline cycle does not contain edges or vertices more than once.

Lemma 2.14. *Given an internally 3-connected, plane graph G and an FAA such that every simple outline cycle has at least three combinatorial convex corners. Then every outline cycle, not the outline cycle of a path, has at least three combinatorial convex corners.*

Proof. Suppose the lemma does not hold. Let $\bar{\gamma}$ be the smallest outline cycle, not the outline cycle of a path, that has at most two combinatorial convex corners. Let γ be the largest simple outline cycle contained in $\bar{\gamma}$ (see Figure 2.10).



Figure 2.10: Examples of outline cycles $\bar{\gamma}$, the largest simple outline cycle contained is colored red.

Suppose γ contains only one vertex. As $\bar{\gamma}$ is not the outline cycle of a path, there exists a $v \in \bar{\gamma}$ which has degree at least 3 in $\bar{\gamma}$, let $\gamma = \{v\}$. Now $\bar{\gamma} \setminus v$ has at least three components, let C be such a component. If $|C| = 1$ then this vertex is a combinatorial convex corner for $\bar{\gamma}$. If C is a path then (at least) the endvertex of C that is not connected to v , is a combinatorial convex corner for $\bar{\gamma}$. If C is not a path, then since it is smaller than $\bar{\gamma}$, it has at least three combinatorial convex corners. At least two of those must also be combinatorial convex corners of $\bar{\gamma}$. We conclude that when γ contains only one vertex, $\bar{\gamma}$ has at least three combinatorial convex corners.

Suppose γ is a cycle of length at least three. As $\bar{\gamma}$ is not a simple outline cycle, $\bar{\gamma} \setminus \gamma$ has at least one component. Such a component is connected to at most one vertex of γ as otherwise γ is not the largest simple outline cycle in $\bar{\gamma}$. Similar as in the previous case, each component in $\bar{\gamma} \setminus \gamma$ contributes at least one combinatorial convex corner. As γ has at least three combinatorial convex corners, it now follows that $\bar{\gamma}$ has at least three combinatorial convex corners. This concludes the proof. \square

Contact Families of Pseudo-segments

As mentioned before, the notion of free points is almost the same as the notion of extremal points as used by de Fraysseix and Ossona de Mendez in a series of publications that consider families of pseudosegments. A contact system of pseudosegments is *stretchable* if it is homeomorphic to a contact system of straight-line segments. De Fraysseix and Ossona de Mendez characterized stretchable systems of pseudosegments [dFdM03, dFdM04, dFdM07a]. In this section we give a new proof of this characterization using Theorem 2.13.

Definition 2.15. Let Σ be a family of pseudosegments and let S be a subset of Σ . A point p is an *extremal point* for S if:

1. p is an endpoint of a pseudosegment in S , and,
2. p is not interior to a pseudosegment in S , and,
3. p is incident to the unbounded region of S .

Theorem 2.16 ([dFdM07a]). *A contact family Σ of pseudosegments is stretchable if and only if every subset $S \subseteq \Sigma$ of pseudosegments with $|S| \geq 2$ has at least three extremal points.*

Our notion of a free point (Definition 2.10 on page 23) is more restrictive than the notion of an extremal point. In the following we show that they are essentially the same. First in Proposition 2.17 we show that in the case of families of pseudosegments that live on a plane graph via an FAA, the two notions coincide. Then we continue by reproving Theorem 2.16 as a corollary of Theorem 2.13.

Proposition 2.17. *Let G be an internally 3-connected, plane graph and Σ a contact family of pseudosegments associated to an FAA such that each subset $S \subseteq \Sigma$ has at least three extremal points or cardinality at most one. The unique solution of the system of equations corresponding to Σ is an SLTR.*

Proof. Note that in the proof of Theorem 2.13 the notion of free points is only used to show that there is no degenerate vertex (Item 4.). We show how to modify this part of the argument for the case of extremal points.

Consider again the set S of pseudosegments aligned with ℓ . We will show that all extremal points are also free points. Let p be an extremal point of S . Assuming that p is not free, we can negate item 4 from Definition 2.10, i.e., all the pseudosegments incident to p are in S . By 3-connectivity p is incident to at least three pseudosegments, all of which lie on the line ℓ . Since all regions are bounded by three pseudosegments and p is not interior to a pseudosegment of S , all the regions incident to p must lie on ℓ . But then p is not incident to the unbounded region of S , hence, p is not an extremal point. Therefore, all extremal points of S are also free points of S . Proposition 2.17 now follows from Theorem 2.13.

Proof of Theorem 2.16. Let Σ be a contact family of pseudosegments that is stretchable. Consider a set $S \subseteq \Sigma$ of cardinality at least two in the stretching, i.e., in the segment contact representation. Endpoints (of segments) on the boundary of the convex hull of S are extremal points. There are at least three of them, unless S lies on a line ℓ . In the latter case, there is a point q on ℓ that is the endpoint of two colinear segments. This is a third extremal point.

Conversely, assume that each subset $S \subseteq \Sigma$ of pseudosegments, with $|S| \geq 2$, has at least three extremal points. We aim at applying Proposition 2.17. To this end we construct an

extended system Σ^+ of pseudosegments in which every region is bounded by precisely three pseudosegments.

First we take a set Δ of three pseudosegments that touch each other as the three sides of a triangle so that Σ is in the interior. The corners of Δ are chosen as suspensions and the sides of Δ are deformed such that they contain all extremal points of the family Σ . Let the new family be Σ' .




Figure 2.11: Protection points (red) and triangulation points (blue) of two faces of some contact family of pseudosegments.

Next we add new *protection points* (see Figure 2.11). These points ensure that the pseudosegments of Σ' will be mapped to straight-lines. For each inner region R in Σ' , and for each pseudosegment s in R , we add a protection point for each visible side of s . The protection point is connected to the endpoints of s , with respect to R from the visible side of s .

Now the inner part of R is bounded by an alternating sequence of endpoints of Σ' and protection points. We connect two protection points if they share a neighbor in this sequence (blue edges in Figure 2.11). Finally, we add a *triangulation point* in R and connect it to all protection points of R .

This construction yields a family Σ^+ of pseudosegments such that every region is bounded by precisely three pseudosegments and every subset $S \subseteq \Sigma^+$ has at least 3 extremal points, unless it has cardinality one. Let V be the set of points of Σ^+ and E the set of edges induced by Σ^+ . It follows from the construction that $G = (V, E)$ is internally 3-connected.

By Proposition 2.17 the graph $G = (V, E)$ together with Σ^+ is stretchable to an SLTR. Removing the protection points, triangulation points and their incident edges yields a contact system of straight-line segments homeomorphic to Σ . \square

2.1.2 Schnyder Labelings and Flat Angle Assignments

In this section we present a second characterization of SLTRs. This characterization is based on Schnyder woods and FAAs. In an SLTR every face is a triangle and the not assigned angles of a face are its *corners*. We also call the angles in a graph that are not assigned by some FAA *combinatorial corners* or simply *corners*, for this FAA. Given a Schnyder labeling of a 3-connected plane graph, the labels of the corners of every face are a subset of $\{1, 2, 3\}$.

Definition 2.18 (Corner compatibility). A Schnyder labeling σ and an FAA ψ are *corner compatible* if

[C1] The Schnyder labeling σ and the FAA ψ use the same suspensions.

- [C2] Every inner face has a corner in ψ that is labeled 1 in σ , a corner in ψ that is labeled 2 in σ and a corner in ψ that is labeled 3 in σ .

The remainder of this section is dedicated to proving that every SLTR has (at least) a corner compatible pair and every corner compatible pair induces an SLTR.

Theorem 2.19. *Let G be a suspended, internally 3-connected graph. Then G has an SLTR if and only if there exists a corner compatible pair of G , i.e., an FAA ψ and a Schnyder labeling σ .*

Showing that having a corner compatible pair is sufficient for having an SLTR is fairly easy. The structure of the drawing obtained by face-counting (see Section 1.4) can be used efficiently to show that every simple outline cycle has at least three combinatorial convex corners.

Lemma 2.20. *Let G be a suspended, internally 3-connected graph. If an FAA ψ and a Schnyder labeling σ are corner compatible, then every simple outline cycle has at least three combinatorial convex corners with respect to ψ .*

Together with Lemma 2.14 this proves that the FAA is a GFAA. From Theorem 2.13 we know that a GFAA induces an SLTR in which the flat angles are prescribed by ψ .

Proof. Let γ be a simple outline cycle and F_{int} be the set of interior faces of G . Let $\alpha_1, \alpha_2, \alpha_3 \in \mathbb{R}^2$ be three independent vectors, e.g., $(0,0)$, $(0,1)$ and $(1,0)$, and let D be the drawing of G obtained by face-counting using these vectors. An example is given in Figure 1.16 on page 13. Consider the outline cycle γ in D . We sweep over D with the line $(\alpha_{i-1} - \alpha_{i+1})$, starting at the suspension s_i (see Figure 2.12). Let M_i be the set of vertices of γ that are met first by these sweep lines (respectively).

Observation 1. All vertices of a face f with label i inside f are met by the sweep line $(\alpha_{i-1} - \alpha_{i+1})$ at the same time. This follows from Property 1.13: ‘The vertices of an interior face are placed on the boundary of a triangle with sides on lines $c_i(\alpha_{i-1} - \alpha_{i+1})$ for some constant c_i ’ (see Figure 2.12).

Observation 2. For a vertex $v \in M_i$, all the angles of v that lie inside γ are labeled i . The three sweep lines divide the angles of a vertex (see Figure 2.13). The angles that lie completely on the opposite side of suspension s_i with respect to the line with direction $(\alpha_{i-1} - \alpha_{i+1})$ through v are all labeled i .

It follows from Observation 2 that the sets M_1, M_2 and M_3 are disjoint. The Schnyder labeling and the FAA are corner compatible, therefore, every face has a corner of label i . Hence, in each set M_i there is a vertex which is not assigned inside of γ . By convexity of the drawing such a vertex has a neighbor outside of γ . A vertex in γ that is not assigned inside γ and has a neighbor outside is a combinatorial convex corner for γ . As the sets M_i are disjoint, each outline cycle has at least three combinatorial convex corners.

Hence, using Theorem 2.13 we obtain that ψ induces an SLTR of G . □




Figure 2.12: A triangle surrounding a face and the sweep lines.




Figure 2.13: The separation of the neighbors of a vertex, Figure 3 from [BFM07]

On the other hand, for every SLTR there exists a Schnyder labeling such that the FAA from this SLTR and the Schnyder labeling are corner compatible. Let G be a suspended, internally 3-connected graph that admits an SLTR. Let R be an SLTR of G and ψ the FAA induced by R .

First we introduce two geometric objects that will be useful. Examples are shown in Figure 2.14.

Definition 2.21 (Separating (Subdivided) Triangle). A separating (subdivided) triangle is a triangle in the drawing, formed by some set of edges and three corners such that:

- Every boundary vertex of the triangle that is not a corner is assigned (either inside or outside the triangle), and,
- There is a vertex, which is not one of the corners, that has no neighbor strictly outside of the triangle and there is a vertex, which is not one of the corners, that has no neighbor strictly inside the triangle.

Note that this triangle may partly be on the boundary of the drawing.

Definition 2.22 (Dividing Segment). A dividing segment is a set of edges that lie on a line such that the union of these edges cuts the drawing into two nonempty parts. Moreover, every interior vertex on this segment is assigned to a face bounded by two edges of the dividing segment.




Figure 2.14: Examples of separating triangles, colored grey in (a) and (b), and of dividing segments, between p and q , colored grey in (c) and (d).

In order to show that for every SLTR and FAA belonging to this SLTR, there exists a Schnyder labeling such that the FAA and the Schnyder labeling are corner compatible, we

assume that it is not the case and aim for a contradiction. Let G be a counterexample with the minimum number of vertices and, subject to this condition, as few edges as possible. Let R be an SLTR of G , ψ the FAA belonging to R and s_1, s_2, s_3 the suspensions. We will first show some properties of R .

Lemma 2.23. *R has no separating (subdivided) triangle.*

Proof. Suppose, to the contrary, R has a separating (subdivided) triangle (a, b, c) . Let R_1 be the part of R that contains everything outside of the separating (subdivided) triangle, and the boundary of the separating (subdivided) triangle. Let R_2 be the part of R that contains the inside and the boundary of the separating (subdivided) triangle (see Figure 2.15). The vertices on the boundary of the separating (subdivided) triangle that have degree two in R_i , are replaced by an edge between their two neighbors in R_i . R_1 and R_2 are SLTRs with less vertices than G . Therefore, they cannot be counterexamples. Hence, there exists a Schnyder labeling that is corner compatible with the FAAs of R_1 and R_2 . For R_2 the vertices a, b, c are the suspensions, the labels of the suspensions are chosen to coincide with their labels in the now empty triangle in R_1 (renaming the labels does not change the Schnyder labeling). The FAAs for the smaller graphs are a subset of ψ , the FAA of G . The Schnyder labelings combined give a Schnyder labeling σ of G . It follows, from the fact that in R_1 and R_2 the Schnyder labelings and FAAs are corner compatible, that ψ and σ are corner compatible, which contradicts the assumption. Hence, R has no separating (subdivided) triangle. \square




Figure 2.15: Splitting when R has a separating (subdivided) triangle.

Lemma 2.24. *R has no dividing segment.*

Remark. This also implies that there is no degree-2 suspension.

Proof. Suppose to the contrary there is a dividing segment with p and q as endpoints. If one of p and q is a suspension, then R has a separating (subdivided) triangle (see Figure 2.14 (c)). Moreover, if on both sides of the dividing segment there is not just a degree-2 suspension, then again R has a separating (subdivided) triangle (see Figure 2.14 (d)). Therefore, we may assume that the dividing segment separates a degree-2 suspension s_1 , and p and q are not suspensions. There are two cases. The dividing segment consists of one edge pq (Case 1) or there is at least one vertex on the dividing segment between p and q (Case 2).

Case 1. If p or q is a degree-3 vertex there must be a separating (subdivided) triangle, as pq is an edge (see Figure 2.16 (a)). This contradicts Lemma 2.23. So both p and q have at least degree 4. Recall that the suspension that is separated, s_1 , has degree 2. Therefore, all neighbors of p and q , other than s_1 , are on the same side of the dividing segment. We claim that one of the edges ps_1 or qs_1 can be contracted such that the resulting graph has an SLTR with the same assignment as G , except for p and q (see Figure 2.16 (b)). This does not come for free as s_2, q, p (or s_3, p, q) becomes a straight-line segment. Contracting one of the edges ps_1 or qs_1 could result in a degeneracy. To obtain a degeneracy in both cases, there have to be two cycles containing p and q which both have precisely three combinatorial convex corners (see Figure 2.16 (c)). Only then does contracting ps_1 as well as contracting qs_1 induce a simple outline cycle with at most two combinatorial convex corners.




Figure 2.16: (a) A separating (subdivided) triangle when p has degree 3, there cannot be edges or vertices in the grey area. (b) Contracting and decontracting the edge ps_1 . (c) The paths that induce a degeneracy after contracting ps_1 or qs_1 ; the angle at z with the red arc is strictly smaller than π .

Let v and w be the third corners of these cycles (see Figure 2.16 (c)). There exists a path from q to v and similarly a path from p to w . These paths belong to the cycles. To induce a degeneracy when ps_1 or qs_1 is contracted, such a path cannot contain a corner besides its endpoints. Hence, it is a straight-line path in the SLTR or ‘steeper’, i.e., it has concave angles in the interior of the cycles p, q, v or p, q, w . Let z be the vertex where the two paths cross. As z should not be a corner for either of the paths, it must be assigned inside both cycles. This implies that the angle at z that is interior to the cycles and between the paths to p and q is at least of size π in the SLTR. This is a contradiction to the assumption that R is an SLTR (see Figure 2.16 (c)).

Therefore, we can contract at least one of these edges, say qs_1 . Let G' be obtained from G by contracting qs_1 and deleting ps_1 . The assignment ψ' is obtained by removing the assignment of q from ψ . The vertex p is assigned to the outer face. As G' has less vertices than G , it is not a counterexample. Therefore, G' has a Schnyder labeling that is corner compatible with ψ' . This labeling is extended to a labeling of G by adding the labels 1, 2, and 3 clockwise in the face s_1, q, p giving the angle at s_1 label 1 (see Figure 2.16 (b)). It is immediate that this Schnyder labeling is corner compatible with ψ .

Case 2. Let x be the first neighbor of p on the dividing segment. The graph G' is obtained by contracting the edge ps_1 . No edges are deleted (see Figure 2.17). The assignment of G' is obtained from ψ by deleting the assignments of x and p . Every simple outline cycle in G' has an equivalent simple outline cycle in G . The assignments of x and p are removed, but no edge of x is deleted. Therefore, if x is a combinatorial convex corner for the equivalent outline cycle in G then it is also a combinatorial convex corner for the outline cycle in G' . Therefore, G' has an SLTR. As G' is smaller than G , it has a Schnyder labeling that is corner compatible with ψ' . This labeling is extended to a labeling of G by adding the label 1 to

the angle at s_1 and 3 to the new angle at p . This Schnyder labeling is corner compatible with ψ . \square




Figure 2.17: Two examples of contracting ps_1 and extending the labeling along the decontraction.

The following lemma shows a property of a corner compatible pair that turns out to be useful later on.

Lemma 2.25. *Let G be a suspended internally 3-connected planar graph with a corner compatible pair ψ and σ , an FAA and a Schnyder labeling. Let v be a neighbor of a suspension. If v is assigned in ψ to a face that contains this suspension then the label of the assigned angle of v is unique among the labels of the angles of v .*

Proof. Without loss of generality, let v be a neighbor of s_1 and such that v is assigned to the left face f of the edge vs_1 . Then the assigned angle of v has label 2 (see Figure 2.18). The face to which v is assigned is denoted by f , f has both v and s_1 on its boundary, and the other neighbor of v in f is denoted by w . The edge vs_1 has two angles labeled 1 (at s_1) and an angle labeled 2, the assigned angle at v . In a Schnyder labeling all three labels must occur around an edge, hence, the fourth angle around vs_1 has label 3. Since ψ and σ are corner compatible, there must be a corner in f with label 2. Hence, the angle of w in f must also have label 2. Around the edge vw there are now two angles labeled 2, hence, the other angles have labels 1 (at w) and 3 (at v). From L2 (of Definition 1.11 on page 10) it follows that around v the label 2 appears only once, precisely at the assigned angle. \square




Figure 2.18: The label at the assigned angle is unique around this vertex.




Figure 2.19: There are no two vertices (z and z') with a straight-line path to x and to s_1 . When z is not a neighbor of s_1 , then there must exist a vertex x' and z' .

Lemma 2.26. *In a minimal counterexample R , no neighbor x_i of a suspension s_i is assigned to a face incident to the edge $s_i x_i$.*

Remark. In particular this implies that the three suspensions form a triangle.

Proof. Suppose, to the contrary, there is a neighbor x of s_1 which is assigned to a face incident to the edge s_1x . We again construct a smaller graph, show that it has an SLTR and that we can extend its Schnyder labeling to a Schnyder labeling of G . Ideally we contract the edge s_1x to obtain G' , however, this is not always possible.

Consider the drawing R . Let z be the vertex that is the third corner of the face which has x and s_1 as corners. Note that, if there are two vertices z, z' connected by a straight-line path to s_1 as well as x , then there is a separating (subdivided) triangle (see Figure 2.19 (a)).

We may assume that z is a neighbor of s_1 . Suppose not, then there must be another vertex x' , between z and s_1 and assigned to the face with corners s_1, z and x . Let z' be the third corner of the face with corners x' and s_1 (see Figure 2.19 (b)). Either z' is a neighbor of s_1 or we find x'' , and so on. As we are moving over the neighbors of s_1 , this process must end. Hence, we find the desired x and z , both neighbors of s_1 .

Here is a summary of how we obtain G' in different cases, the cases are depicted in Figure 2.20.

1. If x and z are neighbors:
 - (a) And if z is assigned to the other face bounded by xz , then the edge xz is contracted.
 - (b) Otherwise, if z has degree 3, then z is removed and an appropriate edge is added.
 - (c) Otherwise, if xs_1 can be contracted, it is contracted.
 - (d) Otherwise, the edge zs_1 is deleted.
2. If x and z are not neighbors, then the edge xs_1 is contracted.




Figure 2.20: How to obtain G' in the different cases. The solid red edges represent the edge that is to be contracted, the red vertex represents the vertex that is to be removed and the dashed red edge represents the edge to be deleted.

Note that x may be a boundary vertex of G . Similarly as before, the resulting graph G' is shown to have an SLTR, and as it has less vertices than G , it admits a corner compatible pair. The obtained Schnyder labeling is shown to be extendable to G , and this results in a corner compatible pair for G , contradiction. We will now go into the different cases in more detail.

Let xz be an edge and thus (s_1, z, x) forms a triangle.

Case 1a. The graph G' is obtained by deleting the edge s_1z and contracting xz , the new vertex is called x' (see Figure 2.21). The assignment ψ' of G' is obtained from ψ by deleting the assignment of z . Every simple outline cycle c' in G' has an equivalent simple outline cycle in G . The combinatorial convex corners of such an equivalent cycle that are not x or z , must also be combinatorial convex corners of c' . If one of x and z is a corner, then x' is a corner for c' . Suppose x and z are both corners of the equivalent cycle c . Note that c is not

just the face z, x, s_1 , as c' is a simple cycle for which c is the equivalent cycle. Therefore, at least for one of x and z another angle at this vertex must be interior to c (c cannot have a cutvertex). Suppose an angle of z not in z, x, s_1 belongs to the interior of c . Then z cannot see x through this angle, hence, c must connect z to x in a way that goes around the assigned angle of z (c cannot have a cutvertex, connecting through the other side would ensure that s_1 is a cutvertex). But then c has to have at least four combinatorial convex corners in G . A similar argument shows that if an angle of x not in z, x, s_1 belongs to the interior of c then c must also have at least four combinatorial convex corners. It follows that every simple outline cycle in G' has at least three combinatorial convex corners. Therefore, G' has an SLTR. Since G' has less vertices than G , it cannot be a counterexample. Take a Schnyder labeling σ' of G' that is corner compatible with ψ' . From Lemma 2.25 it follows that the assigned angle at x' has a unique label for this vertex.

We reverse the contraction of xz (see Figure 2.21). The angles of s_1, x, z in $f_{s_1, z, x}$ are 1, 2, 3 in clockwise order starting at s_1 . At x all labels occur. The assigned angle of z gets label 1, then all labels occur around z . In the face $f_{s_1, z, x}$ all labels occur and the relabeled angle at z is not a corner. Further nothing has changed with respect to σ' . Therefore, the obtained Schnyder labeling is corner compatible with ψ . This contradicts the assumption that G is a counterexample.




Figure 2.21: Contracting the edge xz . The label at the corner with a $*$ is 1 or 3, in both cases the label 1 can be chosen for the assigned angle of z .

Case 1b. Let z have degree 3. Suppose z is assigned, then it is not assigned to one of the faces bounded by xz (due to Case 1a). Suppose z is assigned along the face to s_1 as in Figure 2.22 (a), then there is a separating (subdivided) triangle, as z has degree 3 and the lower face bounded by xz must be a triangle. This contradicts Lemma 2.23. Therefore, we may assume that z is not assigned.




Figure 2.22: Deleting z when w is assigned to the face with s_1 or when w has degree 3.

The third neighbor of z , denoted by w , is assigned (otherwise there is be a separating (subdivided) triangle, containing z in its interior). In this case we obtain G' by deleting z and adding an appropriate edge. When w is assigned to the face with s_1z (see Figure 2.22 (b,c)) or when w has degree 3 (see Figure 2.23), we proceed as follows. After z is deleted, the

edge xw or s_1w , depending on the assignment of w , is added. The assignment ψ' is obtained from ψ by deleting the assignment of w . It is immediate that G' has an SLTR, take the drawing of G , erase z and add the appropriate edge to obtain an SLTR of G' . Therefore, there must be a Schnyder labeling of G' that is corner compatible with ψ' , the labelings are depicted in Figure 2.22 (c) and Figure 2.23 (c). Note that, if w and s_1 are not neighbors, the labeling at w is the same as depicted in the figures. Adding z is equivalent to subdividing the newly added edge of w and connecting z to the other of x and s_1 . The labels at the ends of the subdivided edge do not change. At x or s_1 the label of the new angles are the same as the label of the bigger angle in G' . The labels around z follow. There is one new assignment, that is the assignment of w . The angle of z along the assignment gets the same label as that angle of w . Therefore, in the incident face all labels occur as corners. The obtained Schnyder labeling of G is corner compatible with ψ . This contradicts the assumption that G is a counterexample.




Figure 2.23: Deleting z when w is assigned to the face with s_1 or when w has degree 3.

If w is assigned to the face with zx and w has degree at least 4, the above procedure does not necessarily work. Consider the drawing in Figure 2.23 (c). Suppose w has degree at least 4. Then the labeling shown is not the only possibility, namely the labels of the face x, w, p could be rotated one step counterclockwise. Then w has a unique label in this face after introducing z . As w is assigned, this Schnyder labeling is not corner compatible with ψ . Hence, we have to take different measures. The procedure is depicted in Figure 2.24.

Instead of removing z , we map z onto the edge xs_1 . To be precise, in the SLTR of G , z is mapped to its projection onto the edge xs_1 as seen from w . This is possible since z is not assigned, i.e., all the angles of z are strictly smaller than π , and both x and w have a straight-line path to s_1 . Therefore, there is no obstruction between z and its projection. The edge s_1x is removed, the resulting graph is denoted by G' . The assignment ψ' of G' is ψ together with the new assignment of z to the same face as where x is assigned. That G' has an SLTR follows from the construction.




Figure 2.24: The vertex z is mapped to its projection as seen by w .

Since G' has the same number of vertices as G but fewer edges than G (namely one edge less), G' is not a counterexample. A Schnyder labeling in G' that is corner compatible with ψ' , yields a Schnyder labeling of G that is corner compatible with ψ (see Figure 2.24).

Hence, we have obtained a Schnyder labeling of G that is corner compatible with ψ , this contradicts the assumption.

This concludes the proof for the case that z has degree 3.

Case 1c. Due to Case 1a and Case 1b we may assume that z is not assigned to a face bounded by xz and z has degree at least 4. G' is obtained by contracting the edge s_1x and deleting the edge xz (see Figure 2.25). Here we use that z has degree at least 4 in G and hence, degree at least 3 in G' . The assignment of flat angles ψ' for G' is taken over from the assignment ψ of G . The assignment of x is removed.




Figure 2.25: Contracting the edge between a suspension s_1 and a neighbor x which is assigned to a face along the edge s_1x when xz is an edge.

Every simple outline cycle in G' has an equivalent simple outline cycle in G . When x is a combinatorial convex corner of the equivalent cycle, then s_1 is its replacement in G' . Suppose z is a corner of the outline cycle in G but not in G' . Then all neighbors of z except x are interior to or on the outline cycle. Suppose this outline cycle has only two combinatorial convex corners in G' . Confer Figure 2.26 (a). The path from s_1 to p has at least one interior vertex, as z has degree at least 4. If all these vertices are assigned on the outside then the outline cycle looks as in Figure 2.26, and it is a separating (subdivided) triangle in G . If the path is concave, as in Figure 2.26 (b), then this does not hold. In particular in this situation we say that s_1x cannot be contracted (this is processed by Case 1d). Obviously, if the path is strictly convex then the outline cycle has at least 4 combinatorial convex corners in G .




Figure 2.26: If the path between s_1 and p is a straight-line segment, then there is a separating (subdivided) triangle (a). If the path is concave, then the edge s_1x is called non-contractable (b). If this holds, then the neighbor of z clockwise before x must see s_1 without being obstructed by a neighbor of z , indicated by the dashed red line (c). This ensures that there is no such visibility between x and the clockwise first neighbor of z after s_1 .

We proceed with the case where there is no outline cycle of which z is the third combinatorial convex corner in G , but z is not a combinatorial convex corner in G' for this cycle. This is

equivalent to saying that s_1x is contractable. In this case, every simple outline cycle in G' has at least three combinatorial convex corners and, therefore, G' has an SLTR.

As G is a minimal counterexample, we find a Schnyder labeling σ' of G' such that ψ' and σ' are corner compatible. A Schnyder labeling can be extended along the reverse of a contraction that involves a suspension (see Figure 2.25). The face bounded by s_1, x and z is a 3-face and it follows from the Schnyder labeling that all three labels occur. Secondly, the assigned angle at x does not have a unique label for this face. The corner compatible pair ψ' and σ' is extendible to a corner compatible pair for G , which contradicts the assumption.

Case 1d. We may assume that z is not assigned to a face bounded by xz , z has degree at least 4 and s_1x is not contractable. Since s_1x is not contractable, we know that there must be a cycle c in our, with z is on the boundary, all neighbors of z besides x are on or inside c and c has precisely three combinatorial convex corners, of which z is one. Confer Figure 2.26 (b), we know that the path between s_1 and p is strictly concave with respect to the interior of the cycle c . From this it follows that p can see s_1 without being obstructed by any of the neighbors of z . Since z is not assigned to a face incident to xz and z is a combinatorial convex corner for c , which includes all neighbors of z except x , we can conclude that z is not assigned.

We obtain G' by deleting the edge s_1z . The assignment ψ' is equal to ψ except for the assignment of vertex z to the face with s_1, x and z on its boundary (see Figure 2.27 (b) and (c)). Suppose in G' under ψ' there is an outline cycle γ' that has at most two combinatorial convex corners. Then γ' must contain z and z is a corner for the equivalent outline cycle γ in G . If γ' contains all neighbors of z in G' then γ has at least four combinatorial convex corners or z is not a combinatorial convex corner for γ . This follows from the fact that in G , the clockwise first neighbor of z after s_1 cannot see x because the clockwise last neighbor of z before x must obstruct its view (see Figure 2.26 (c)). Therefore, between these two neighbors there must be a combinatorial convex corner on γ . Suppose γ' contains the assigned angle of z in G' and z is a combinatorial convex corner for the equivalent outline cycle γ in G . Then the equivalent outline cycle in G has at least 4 combinatorial convex corners, since the paths from z to s_1 bounding this cycle must both be strictly convex with respect to this cycle (see Figure 2.27 (a))




Figure 2.27: If z was a convex corner for the grey cycle in G , then this cycle has at least 4 combinatorial convex corners since both paths from z to s_1 are strictly convex with respect to the cycle (a). Deleting the edge s_1z (b).

We conclude that G' has an SLTR and since G' has fewer edges than G , it is not a counterexample. We obtain a Schnyder labeling σ' which is corner compatible with ψ' . But this can be changed into a Schnyder labeling in G which is corner compatible with ψ (see Figure 2.27 (b)).

This concludes the proof in the case that x and z are neighbors.

Case 2. When xz is not an edge, there is at least one vertex p , between z and x . Such a vertex is assigned to the face containing x, p, z and s_1 (see Figure 2.28). Let p be such a vertex closest to x . We obtain G' by contracting s_1x . The assignments of p and x are removed from ψ to obtain the assignment ψ' for G' .




Figure 2.28: Contracting the edge between a suspension s_1 and a neighbor x which is assigned to a face along the edge s_1x when xz is not an edge.

Every simple outline cycle in G' has an equivalent simple outline cycle in G . If x is a combinatorial convex corner for the equivalent outline cycle, then s_1 or p is its replacement for the cycle in G' . The assignment has changed for the vertex p . Whenever p is a combinatorial convex corner of an outline cycle in G then it must have had a neighbor outside. No edge incident to p has been removed, therefore, p also has a neighbor outside of the outline cycle in G' . As p is no longer assigned, it is a combinatorial convex corner for the outline cycle in G' as well. It follows that every simple outline cycle in G' has at least three combinatorial convex corners and this implies that G' has an SLTR prescribed by ψ' .

As G is a minimal counterexample, we find a Schnyder labeling σ' of G' such that ψ' and σ' are corner compatible. We have to check whether all three labels occur on corners of the face containing x, p, z and s_1 , call this face f . In σ' the angle of p in f' has label 2 (or 3) since it is clockwise first (or last) around the edge ps_1 . In G the same holds for the angle of x in f . Hence, the corner compatible pair ψ' and σ' is extendable to a corner compatible pair for G .

In all cases we have obtained a contradiction to the assumption that G is a minimal counterexample, this concludes the proof of the lemma. \square

Now we are ready to show that there cannot be a counterexample G that has an SLTR but no Schnyder labeling that is corner compatible with the FAA belonging to the SLTR.

Lemma 2.27. *Let G be a suspended, internally 3-connected graph that admits an SLTR for which ψ is an appropriate FAA. Then there exists a Schnyder labeling σ of G such that ψ and σ are corner compatible.*

Proof. Suppose the lemma does not hold and let G be a counterexample with the minimum number of vertices and, subject to this condition, as few edges as possible. Let R be an SLTR of G , ψ the FAA belonging to R and s_1, s_2 and s_3 the suspensions. We aim for a contradiction by constructing a smaller graph G' , from G , such that:

- G' has an SLTR based on the FAA from G , and,
- The corner compatible pair of G' gives a Schnyder labeling for G , and,
- The obtained Schnyder labeling is corner compatible with ψ .

This contradicts the assumption that G is a minimal counterexample.

From Lemma 2.23, Lemma 2.24 and Lemma 2.26 it follows that

- [B1] R has no separating (subdivided) triangle.
- [B2] R has no dividing segment. This implies that there is no degree two suspension.
- [B3] No neighbor x_i of a suspension s_i is assigned to a face incident to the edge $s_i x_i$. In particular this implies that the three suspensions form a triangle.

From B3 it follows that the outer face is a triangle and the neighboring faces of the outer face are also triangles (see Figure 2.29 on the left). The third vertex of the inner face along the edge $s_i s_{i+1}$ plays an important role. We denote this vertex q (it is not important which i is considered).

First we show that if (for some i) the vertex q is not assigned (has no flat angle) then G cannot be a minimal counterexample. Then we show that if q is assigned for each i , we can change G into another graph G' which has some vertex q' as in Lemma 2.26. Finally we show that a Schnyder labeling σ' of G' can be changed into a Schnyder labeling σ of G in such a way that if σ' and ψ' are corner compatible then σ and ψ are corner compatible.




Figure 2.29: Creating a graph with fewer edges.

Let f be the inner face bounded by the edge $s_1 s_2$. Let q be the third vertex of this face. Assume that q is not assigned. We remove the edge $s_1 s_2$ and assign q to the outer face to obtain G' and an FAA ψ' . No simple outline cycle in G' contains the assigned angle of q in its interior. If both s_1 and s_2 are in the outline cycle γ' , then q is not a combinatorial convex corner for the equivalent outline cycle γ in G , as it must lie in the strict interior. Therefore, the combinatorial convex corners of γ are also combinatorial convex corners for γ' . Thus, every outline cycle has at least three combinatorial convex corners and G' has an SLTR for which the stretched angles are prescribed by ψ' . We have obtained a graph with fewer edges. Therefore, G' is not a counterexample. Let σ' be a Schnyder labeling of G' such that σ' and ψ' are corner compatible. It is obvious how this implies the existence of a Schnyder labeling σ of G such that σ and ψ are corner compatible (see Figure 2.29).

Suppose that for every pair of suspensions, the common neighbor q is assigned. Every suspension has at least one neighbor that is not assigned. This follows from the fact that G is not a triangle, and $G \setminus s_i$ must have at least three angles larger than π in the outer face. One of these angles does not belong to a suspension and this vertex must be a not assigned neighbor of s_i . Without loss of generality, let q be the common neighbor of s_1 and s_2 , and let p_k be the clockwise first neighbor of s_1 , seen from q , which is not assigned (e.g., p_3 in Figure 2.30).

We will construct a graph G' from G , in which p_k is assigned to a face bounded by the edge $s_1 p_k$. Then, we will use Lemma 2.26 to obtain a Schnyder labeling σ' for G' , which is corner compatible with ψ' , the assignment of G' . This Schnyder labeling can be changed into a Schnyder labeling of G , which is corner compatible with ψ .

Let p_1 be the neighbor of s_1 clockwise first after q . We obtain a new graph by deleting the edge $q s_1$ and adding the edge $p_1 s_2$, we denote this change by a *flip*. This change is

possible even in the drawing of the SLTR, as s_1, s_2, p_1 and q form a convex 4-gon with a diagonal (see Figure 2.30 (a)). Convexity makes changing the diagonal possible in the current drawing (see Figure 2.30 (b)). Hence, the new graph has an SLTR. If p_1 is assigned, the step is repeated (see Figure 2.30 (c)). Since s_1 must have a neighbor that is not assigned, we meet it eventually. Let p_k be the first such vertex that is not assigned, and p_1, \dots, p_k the set of vertices processed. The graph obtained after flipping the edge qs_1 to p_1s_2 and the edges $p_i s_1$ to $p_{i+1} s_2$, for $i = 1, \dots, k-1$, is denoted by G^* and its assignment by ψ^* . As all the steps are possible even in the drawing of the SLTR of G , therefore, the assignment ψ^* is a Good FAA.




Figure 2.30: Flipping edges.

In the last change the edge $p_k s_2$ is added. Now, this edge is removed⁴ and p_k is assigned to the face with s_1, s_2, p_{k-1} and p_k on its boundary (see Figure 2.31 (a)). Let the obtained graph be G' . The graph G' comes with the assignment ψ' and we have to argue that this again is a good FAA. Let γ' be a simple outline cycle in G' and γ^* the equivalent outline cycle in G^* . If γ' does not have p_k on its boundary, then it has the same combinatorial convex corners as γ^* . So suppose p_k is on the boundary of γ' , and not a combinatorial convex corner. If γ' has the assigned angle of p_k in its interior, but not all neighbors of p_k , then γ^* has at least four combinatorial convex corners in G^* , namely s_1, s_2, p_k and one on the boundary path between p_k and s_2 . Therefore, γ' has at least three combinatorial convex corners. Suppose γ' encloses all neighbors of p_k , but not its assigned angle. In the SLTR of G^* the angle at p_k is a concave corner for γ^* , therefore, there must be at least three other combinatorial convex corners for γ^* . These are the combinatorial convex corners of γ' . We conclude that every simple outline cycle has at least three combinatorial convex corners and, therefore, ψ' is a good FAA.

As G' has a vertex p_k that is assigned to a face bounded by the edge $s_1 p_k$, it cannot be a counterexample, due to Lemma 2.26. Therefore, there is a Schnyder labeling σ' such that ψ' and σ' are corner compatible. It remains to show that σ' can be changed into a Schnyder labeling of G which is corner compatible with ψ .

The reason why the Schnyder labeling can be changed without violating corner compatibility depends highly on the property that an assigned angle does not have a unique label for this face in a corner compatible pair. The Schnyder labeling σ' is unique for the triangles incident

⁴The reason that we do not use the argument as in Figure 2.29 is that in order to do the reverse flipping, we need two angles labeled 1 at p_k around the edge towards p_{k-1} . If we would push p_k to the boundary between s_1 and s_2 then this is not ensured.




Figure 2.31: A labeling of G^* obtained from a labeling of G' .

to s_1 or s_2 (see Figure 2.31 (c)). The first step is to add the edge $p_k s_2$ to obtain G^* . The labeling is extended trivially, see Figure 2.31 (b).

Let q, p_1, \dots, p_k be the set of processed vertices (see Figure 2.30). By definition, the vertices q, p_1, \dots, p_{k-1} all have an assigned angle in G' . Two consecutive vertices in this sequence, p_{i-1} and p_i , are not necessarily neighbors, but they are the corners of the face containing p_{i-1}, p_i and s_2 . Hence, the vertices on the path from p_{i-1} to p_i on the boundary of this face, are all assigned to this face (see Figure 2.32 (a)).

Let $p_0 = q$. We consecutively do the reverse flip: for $i = k, \dots, 1$, the edge $p_i s_2$ is removed and the edge $p_{i-1} s_1$ is added. Note that the degree of the vertex p_i is at least 4 in step i as G is 3-connected and p_i has one neighbor more in G^* than in G . Therefore, the angle clockwise before the edge to s_1 and the angle clockwise after the path to p_{i-1} are not the same angle. Along the steps, the following invariant is maintained.

Invariant: At step i the vertex p_i has labels 2,3,1,1, clockwise, around the edges to s_1 and s_2 and towards p_{i-1} (see Figure 2.32 (a)).

We also maintain that the Schnyder labeling is corner compatible with the FAA.

In the base case the invariant holds, as shown in Figure 2.31 (b). There is no other way to assign the three labels of p_k in G' then the option given in the figure. The addition of the edge $p_k s_2$ also does not yield other options.

In step i we delete the edge $p_i s_2$ and add the edge $p_{i-1} s_1$. Before the reverse flip, the labels at p_i satisfy the invariant and the Schnyder labeling is corner compatible with the assignment. The angles at p_{i-1} , clockwise around the edge towards p_i , have labels 2 and 3 (see Figure 2.32 (a)). The label 3 follows from the fact that the edge goes to s_2 and the label 2 follows from the fact that either there is another angle labeled 3 (Figure 2.32 (a)) or there are two angles labeled 1 along the edge (Figure 2.32 (b)).

Now we change from $p_i s_2$ to $p_{i-1} s_1$ as shown in Figure 2.32 (b). For most of the angles there is a unique choice for the label in order to maintain a Schnyder labeling, these labels are given in Figure 2.32 (b) on the right. To maintain the invariant the angle with the question mark has to get label 1. We have to show that this is a valid choice. If this angle is already labeled 1 it is trivially a valid choice. Suppose not, then this must be the assigned angle of p_{i-1} . This follows from the fact that the neighboring angles of an assigned angle, at a vertex, cannot all three have the same label, this would violate corner compatibility in the face to which the angle is assigned (see Figure 2.33 (a), the face with a question mark does




Figure 2.32: Labeling between p_i and p_{i-1} and the changing of a labeling along a reverse flipping of $p_i s_2$ to $p_{i-1} s_1$.

not have an corner with label 2.).

Suppose the angle labeled with a question mark in the rightmost drawing of Figure 2.32 is not labeled 1. As the labeling is a Schnyder labeling, the only other option is label 2. Consider Figure 2.33 (b), $j = 2$. From the rule around an edge of a Schnyder labeling, it follows that $k = 1$. Then it follows that $l = 2$ as this face must have a corner with label 2. Changing labeling of the assigned angle to 1 again gives a Schnyder labeling. Moreover, as this is an assigned angle, this Schnyder labeling is also corner compatible with the FAA.




Figure 2.33: Changing label j .

Hence, we have changed from $p_i s_2$ to $p_{i-1} s_1$, the invariant now holds for p_{i-1} and the Schnyder labeling is corner compatible with the FAA. By induction, we obtain a Schnyder labeling of G which is corner compatible with ψ which contradicts the assumption that G is a counterexample.

Therefore, there does not exist a counterexample. This proves the theorem. \square

We have now seen all the ingredients needed to prove Theorem 2.19.

Theorem 2.19. *Let G be a suspended, internally 3-connected graph. Then G has an SLTR if and only if there is a corner compatible pair consisting of an FAA ψ and a Schnyder labeling σ .*

Proof. Let G have a corner compatible pair, then due to Lemma 2.20, this FAA is a GFAA. From Theorem 2.13 it now follows that G has an SLTR which agrees with this FAA.

Let G have an SLTR. By Lemma 2.27 there must be a Schnyder labeling that is corner compatible with the FAA that comes from this SLTR. \square

This characterization has not yet led to a recognition algorithm. However, for particular graph classes, namely those graphs with ‘few’ Schnyder woods, the characterization does give a recognition algorithm that runs in polynomial time. This will be discussed in the next section.

2.2 Applications

In this section we will show some applications of the characterizations of SLTRs.

2.2.1 Graphs with few Schnyder labelings

Given an internally 3-connected, suspended, graph G and a Schnyder labeling σ . We will build a bipartite graph $B_{(G,\sigma)}$, such that, a one-sided-perfect matching in this graph, is a GFAA of G .

For each interior face f of G , let f_i be the number of angles with label i . The first vertex class, W_1 of $B_{(G,\sigma)}$ consists of $f_i - 1$ copies of f , for each face f and for each color i , we say that such a copy of f in W_1 is *produced* by color i . The second vertex class W_2 of $B_{(G,\sigma)}$ contains the interior vertices of G . There is an edge uv , for $u \in W_1, v \in W_2$, if and only if, u is copy of f produced by color i and vertex v has an angle with color i in f .

A matching M of $B_{(G,\sigma)}$ is called *one-sided-perfect* if $W_1 \cap M = W_1$, i.e., all the elements of W_1 are matched. Another interpretation of the matching M is that it assigns a vertex of G to a face of G , in such a way, that every vertex is assigned to at most one face.

Lemma 2.28. *Let G be an internally 3-connected, suspended, graph G and σ a Schnyder labeling of G . Let M be a one-sided-perfect matching of $B_{(G,\sigma)}$. Let ψ_0 be the assignment of non-suspension vertices of G , which are on the boundary of G , to the outer face. Then ψ_0 together with M is a GFAA of G .*

Proof. Consider the matching M as an assignment of vertices of G to faces of G . Every vertex is assigned to at most one face, therefore, C_v holds. Moreover, an interior face f of G , has precisely $f_i - 1$ vertices assigned to it, for every color i . Therefore, an interior face f , has precisely $|f| - 3$ vertices assigned to it. In ψ_0 the outer face has all but the three suspensions assigned to it, therefore, C_f must also hold. Let ψ be the FAA of G that contains ψ_0 and M .

An edge uv in $B_{(G,\sigma)}$, for $u \in W_1, v \in W_2$, implies that u is copy of f produced by color i and vertex v has an angle with color i in f . As f has precisely $f_i - 1$ vertices assigned to it, for every color i , it must hold that for every color i there is a vertex with an angle labeled i in f , which is not assigned to f . Therefore, ψ is corner compatible with σ , and from Lemma 2.27 it follows that ψ is a GFAA. \square

On the other hand, due to Hall’s marriage theorem, we can certify the situation where there is no FAA that is corner compatible with a certain Schnyder labeling σ . The certificate is a subset $U \subseteq W_1$, such that, $|U| > |N(U)|$ in $B_{(G,\sigma)}$, where $N(U)$ is the set of neighbors of U .

Graphs with few Schnyder labelings Felsner and Zickfeld have defined the class of 3-connected planar graphs that have a unique Schnyder labeling [FZ08, Zic07]. These are the graphs that can be obtained from a triangle using a sequence of operations chosen from a set of six operations. For this class there is a polynomial time algorithm that outputs a GFAA or a certificate that shows that there is no FAA that is good. The algorithm is based on finding a maximum cardinality matching in $B_{(G,\sigma)}$.

If a graph has only polynomially many Schnyder labelings, then the above can be repeated for every Schnyder labeling, until a GFAA is found, or the output is a set $\{U_j\}_j$, such that U_j violates Hall's marriage condition in the graph $B_{(G,\sigma)}$ where σ is the j -th Schnyder labeling. There exist planar 3-connected graphs on n vertices, which have 3.209^n Schnyder labelings [FZ08]. Therefore, an algorithm checking for a compatible assignment among all Schnyder labelings does not have polynomial complexity in general.

2.2.2 Primal-Dual Triangle Contact representation.

A triangle contact representation is a drawing of a graph in which every vertex is represented by a triangle and every edge by a point contact. De Fraysseix, Ossona de Mendez and Rosenstiehl, proved that every planar graph has a triangle contact representation [dFdMR94]. The algorithm is described in the next chapter, on page 66.

In a *primal-dual contact representation* both the vertices and faces of a graph are represented by triangles. Together these triangles form a tiling of a triangle by triangles. Two triangles share a point if and only if the represented (dual) vertices are adjacent. Two triangles, one representing a vertex v and the other a dual vertex f , share a line segment if and only if $v \in f$.




Figure 2.34: A triangle contact representation and a primal-dual contact representation by triangles.

Gonçalves et al. [GLP12] have shown that every planar 3-connected graph has a primal-dual contact representation by triangles. They use a Schnyder wood of the primal graph to define a contact family of pseudosegments and then they show that this system is stretchable, using the result of de Fraysseix and Ossona de Mendez [dFdM07a]. Moreover, Gonçalves, Lévêque and Pinlou show that these representations are in one-to-one correspondence with Schnyder woods of planar 3-connected graphs. In this section we will give a simpler proof of the first part, by showing that the FAA as constructed by Gonçalves, Lévêque and Pinlou is a GFAA, using a geodesic embedding and the characterization of SLTRs with outline cycles.

Let a 3-connected plane graph G and a primal-dual Schnyder wood for G be given. Following the approach of Gonçalves, Lévêque and Pinlou, we first construct an auxiliary graph H . The SLTR of H will be the dissection of a triangle which is the primal-dual contact representation of G . In contrast to Gonçalves et al. [GLP12] we work with an FAA on H and not with a contact family of pseudosegments.

The vertices of H are the edges of G , including the half-edges at the suspensions. The vertices corresponding to the half-edges are the suspensions of H . The edges of H correspond to the angles of G , i.e., if e and e' are both incident to a common vertex v and a common face f , then (e, e') is an edge of H . The faces of H are in bijection to vertices and faces (dual vertices) of G . In the context of knot theory this graph H is known as the *medial graph* of G .

The graph H inherits a plane drawing from G . Let a Schnyder wood of G be given and consider the geodesic embedding of G and this Schnyder wood onto an orthogonal surface. The faces of H are in bijection to the vertices and faces of G . In an SLTR of H , we need three corners in every face, moreover, every vertex of H (except the three suspensions) has to be the corner for three of its four incident faces. A corner assignment with these two properties is obtained from the orthogonal arcs of the surface, i.e., if s is a vertex and g is a face of H , then s is one of the three designated corners for g if and only if in g there is an orthogonal arc ending in s . The corner assignment is equivalent to an FAA, an angle of s is to be flat if the two edges of H forming the angle belong to the same flat of the orthogonal surface. An example is shown in Figure 2.37.

The family of pseudosegments corresponding to this FAA is precisely the family defined by Gonçalves, Lévêque and Pinlou. This family of pseudosegments also has a nice description in terms of the flats. In fact, there is a bijection between the pseudosegments and bounded flats. A flat F , whose boundary consists of $2k$ orthogonal arcs, contains k saddle-points of the surface, these are the vertices of H on F . These vertices induce a path P_F in H (see Figure 2.35). Every internal vertex of P_F has a flat angle in F and is, hence, assigned. If F is a flat which is constant in coordinate i , then within P_F one of the endpoints is maximal in coordinate $i - 1$ and the other is maximal in coordinate $i + 1$. We call them the *left-end* and the *right-end* of P_F , respectively. In each of the three unbounded flats we have two suspensions of H as end-vertices for the path.

Recall that a flat is rigid if P_F is a monotone path with respect to the coordinates $i - 1$ and $i + 1$ (see Definition 1.15 on page 15). An orthogonal surface is rigid if all its bounded flats are rigid. It has been shown [Fel03, FZ06] that every Schnyder wood has a geodesic embedding on some rigid orthogonal surface. From now on we assume that the given orthogonal surface is rigid, this assumption will be critical in the proof of Proposition 2.29.

To prove that the above defined FAA is a good FAA, we use the structure of the flats. First we note that the flats are naturally partitioned into three classes, let \mathcal{F}_i be the set of flats of color i , i.e., of the flats that are constant in coordinate i . The boundary of the flats in \mathcal{F}_i consists of orthogonal arcs in directions $i - 1$ and $i + 1$.




Figure 2.35: A flat and its right-end r , left-end l , primal-saddles p and dual-saddles d .




Figure 2.36: The candidates, c_1 and c_2 , of γ in color i .

Proposition 2.29. *The flat angle assignment in H as defined above is a Good FAA.*

Proof. It is enough to show that every simple outline cycle has at least three combinatorial

convex corners (Lemma 2.14). Let γ be a simple outline cycle in H . We consider γ with its embedding into the rigid orthogonal surface. Recall that a vertex v of an outline cycle γ is a combinatorial convex corner for γ if:

- [A1] v is a suspension vertex, or,
- [A2] v is not assigned to a face and there is an edge e incident to v with $e \notin \text{int}(\gamma)$, or,
- [A3] v is assigned to a face f , $f \notin \text{int}(\gamma)$ and there exists an edge e incident to v with $e \notin \text{int}(\gamma)$.

On γ we specify some special combinatorial convex vertices, they will be called *candidates*. The candidates are not necessarily distinct but we can show that at least three of them are pairwise distinct. This is sufficient to prove the proposition.

The candidates come with a color. We now describe how to identify the candidates of color i . If γ contains the suspension of color i , then by A1 this is a combinatorial convex vertex for γ and we take it as the candidate. Otherwise, consider the flat F that has the maximal i coordinate among all flats in \mathcal{F}_i that contain a vertex from γ . Let I be a path in $\gamma \cap F$. As candidates of color i , we take the endpoints of I . Of course, if I consists of just one vertex we only have one candidate.

Claim. The candidates are combinatorial convex.

A primal-saddle of F is a corner between two vertices of G and a dual-saddle is a corner between two dual vertices. The vertices of H in F come in four types, left-end, right-end, primal-saddle and dual-saddle (see Figure 2.35).

A primal-saddle of F has two edges in H , that reach to a flat in \mathcal{F}_i with i coordinate larger than F (see Figure 2.36). From the choice of F_i , we know that these two edges do not belong to γ . Therefore, with a primal-saddle in I , both neighbors in P_F also belong to γ and hence, to I . Therefore, a primal-saddle is not an end of I and thus not a candidate.

If an end z of I , is a dual-saddle, then it has an edge e of P_F that does not belong to $\text{int}(\gamma)$ (see Figure 2.36). The edge e is part of the angle at z that belongs to the face to which z is assigned, i.e., z is assigned to a face outside of γ . This shows that z is combinatorial convex by A3.

If z is an end of P_F . Consider the flat F' that contains two H -edges incident to z . The rigidity of F' implies that $P_{F'}$ contains an edge e incident to z that reaches to a flat in \mathcal{F}_i with i coordinate larger than F . Hence, edge e does not belong to γ and not to $\text{int}(\gamma)$. The edge e is part of the angle at z that belongs to the face to which z is assigned. Again z is combinatorial convex by A3.

This concludes the proof of the claim. \triangle

It can happen that a candidate z_i of color i and a candidate z_j of color j coincide. We have to show that in total we obtain at least three different candidates.

Suppose there is only one candidate in color i , z_i . Let F_i , F_{i-1} and F_{i+1} be the three flats around z_i . As there is no other candidate in color i , two edges of z_i are on F_i and these edges do not belong to γ . It follows that the flats F_{i-1} and F_{i+1} are not maximal in their respective colors. Hence, z_i is a candidate only in color i .

Suppose z is a candidate in all three colors. Then there is no edge in γ , incident to z , as otherwise at least one of the three flats incident to γ is not maximal in its respective color. This implies that γ is a single vertex and not a simple cycle. Hence, if γ is a simple cycle, then every candidate is a candidate in at most two colors.

This is enough to show that there are at least three pairwise different candidates. \square




Figure 2.37: The graph H (in blue) is drawn on top of an orthogonal surface (in dashed grey). The flat angles of an FAA are given by the red arrows.




Figure 2.38: An outline cycle γ (in blue) together with the three maximal flats (in red), the candidates a, c and d are a candidate in one color, b is a candidate in two colors.

As every 3-connected plane graph G has a Schnyder wood, we can define the auxiliary graph H and an FAA of H can be obtained as described. Proposition 2.29 shows that this FAA is good. We have thus proved the theorem:

Theorem 2.30 ([GLP12]). *Every 3-connected plane graph admits a primal-dual triangle contact representation.*

In the proof we have worked with the skeleton graph H of the primal-dual triangle representation. We continue by asking which graphs H can serve as skeleton graphs for a primal-dual representation of some graph.

If a dissection of a triangle is a primal-dual triangle contact representation of some graph, then there is a 2-coloring of the triangles. Hence, the skeleton graph H is Eulerian, i.e., all the vertex degrees are even. It is also evident that only degrees 4 and 2 are possible.

Definition 2.31 (Almost 4-regular). A plane graph is almost 4-regular⁵ if:

- There are three vertices of degree 2 on the outer face, and,
- All the other vertices have degree 4.

With the following theorem, we show that deciding whether a plane, almost 4-regular graph has an SLTR, is equivalent to deciding whether the *underlying graph* is 3-connected. With the underlying graph we refer to the graph G such that the almost 4-regular graph is the medial graph of G .

Theorem 2.32. *An almost 4-regular plane graph H has an SLTR if and only if H is the medial graph of an internally 3-connected graph or $H = C_3$.*

⁵Almost 4-regular graphs are (2,3)-tight graphs, or Laman graphs. The number of edges is twice the number of vertices minus three and this is an upper bound for each subset of the vertices.

Proof. Let $H \neq C_3$ be an almost 4-regular plane graph and let R be a SLTR of H . The three suspensions in R are the three degree-2 vertices. Since all the vertices in H have even degree, the dual is a bipartite graph. We abuse notation and denote the bounded faces in R that contain the suspension vertices, by *suspension of the dual*. Since they are all adjacent to the outer face of R , the suspensions are all in the same color class of the bipartition, say in the white class.




Figure 2.39: An SLTR of an almost 4-regular graph H with a 2-coloring of its faces and its underlying graph G .

Let G^Δ be the graph whose vertices correspond to the white triangles of R together with an extra vertex v_∞ . The edges of G^Δ are the contacts between white triangles together with an edge between each of the suspensions and v_∞ . The degree of v_∞ is 3 and each corner of a white triangle is responsible for a contact, hence, every vertex of G^Δ has degree at least 3.

Claim. G^Δ is 3-connected.

Suppose there is a separating set U of size at most 2. Let C be a component of $G^\Delta \setminus U$ such that $v_\infty \notin C$. The convex hull H_C of the corners of triangles in C has at least 3 corners. Covering all the corners of H_C with only two triangles results in a corner p of H_C that has a contact to a triangle $T \in U$ such that p has an angle larger than π in the skeleton of $C + T$. Since p is a vertex of H and angles larger than π do not occur at vertices of degree 4 of an SLTR, this is a contradiction. \triangle

Therefore, if $H \neq C_3$ has an SLTR then G^Δ is 3-connected. On the other hand, if there is an internally 3-connected planar graph G such that H is its medial graph, then by Lemma 2.29 it follows that H has an SLTR. \square

2.2.3 Planar Generic Circuits

A graph $G = (V, E)$ is a *generic circuit* if $|E| = 2|V| - 2$ and for all subsets $H \subseteq V$ the induced graph $G[H]$ has at most $2|H| - 3$ edges. In other words, a generic circuit is (2,2)-tight but all its proper subgraphs are (2,3)-sparse. The generic circuit with the smallest number of vertices is the complete graph on four vertices (K_4). The following result of Berg and Jordan will be of use.

Theorem 2.33 (Berg and Jordan [BJ03]). *A 3-connected, generic circuit can be constructed with Henneberg type 2 steps from K_4 .*

Recall that a Henneberg type 2 step consists of subdividing an edge and connecting the new vertex to a third vertex (see page 5). To show that every generic circuit admits an SLTR, we will show that an SLTR can be extended along a plane Henneberg type 2 step. A plane Henneberg type 2 step is a Henneberg type 2 step that takes place ‘within one face’, i.e., an edge to subdivide is selected and the new vertex is connected to a vertex in one of the two faces bounded by this edge. For a planar generic circuit, we fix an embedding. The construction that we know to exist due to Theorem 2.33 gives us a reverse order. Reverse steps never violate planarity, hence, a 3-connected, plane generic circuit can be constructed from K_4 with plane Henneberg type 2 steps. We start by showing how to get an assignment of the extended graph. Then we will show that this assignment is indeed a GFAA. Since we only consider planar graphs and thus plane Henneberg type 2 steps, we omit the word ‘plane’ in the sequel.

The assignment. Given a graph G and a GFAA ψ of G . Let uv be the edge that is subdivided, x the new vertex and w the third vertex to which x is connected (see Figure 2.40). We denote the face incident to uv and w by f . After the Henneberg step, the new face incident to u is denoted by f_u and the other new face by f_v . The third face incident to uv is denoted by f_x . The resulting graph is denoted by G^+ . We will construct an assignment ψ^+ for G^+ and prove that ψ^+ is a GFAA.

There are three vertices not assigned to f under ψ : we will call them *corners* of f . We consider two cases, firstly f_u is incident to all corners of f , secondly, f_u is incident to precisely two corners of f . The cases are depicted in Figure 2.40. Note that, if w is a corner of f , it will be a corner for both f_u and f_v . The vertices different from u, v, w, x that are assigned to f under ψ will be assigned in the trivial way under ψ^+ , i.e., such a vertex is assigned to f_u (or f_v) if in G^+ it is incident to f_u (or f_v).

Case 1: f_u is incident to all corners of f . If u or w is assigned to f under ψ , it is assigned to f_u under ψ^+ . The vertex v is assigned to f_x and x to f_u under ψ^+ .

Case 2: f_u is incident to precisely two corners of f . If u or w is assigned to f under ψ , it is assigned to f_u under ψ^+ , if v was assigned to f it is assigned to f_v under ψ^+ and x is assigned to f_x .

This yields an assignment ψ^+ for G^+ .

Theorem 2.34. *Given a 3-connected, plane graph G with a GFAA ψ . Let G^+ be the result of a Henneberg type 2 step applied to G and let ψ^+ be the updated assignment. Then ψ^+ is a GFAA and G^+ admits an SLTR.*

Proof. It is trivial that ψ^+ satisfies C_v and C_f and hence, is an FAA.

We consider the induced families of pseudosegments, Σ and Σ^+ of ψ and ψ^+ respectively. Since ψ is a Good FAA, we know that every subset of Σ has at least three free points or cardinality at most one. We will show that every subset of pseudosegments of Σ^+ of cardinality at least two has at least three free points. Let $S \subseteq \Sigma^+$ be a subset of pseudosegments of cardinality at least two. An endpoint of a pseudosegment in S that is not a free point of S is said to be *covered*.

Case 1: Let s_x and s_v be the pseudosegment that has x respectively v as interior point and let s_w the pseudosegment containing the edge vw . If S does not contain s_x, s_v or




Figure 2.40: A stretched representation of the original face, and of the results of a Henneberg type 2 step in Case 1 and Case 2.

s_w then S is also a subset of Σ , hence, it must have at least three free points. Suppose $S \subseteq \{s_x, s_v, s_w\}$, then S has three free points, since, no two pseudosegments of $\{s_x, s_v, s_w\}$ touch twice and if $S = \{s_x, s_v, s_w\}$ there are precisely three of the six endpoints covered. So suppose S contains at least one pseudosegment not of $\{s_x, s_v, s_w\}$. Consider the comparable set S' of Σ , that is

- If $s_x \in S$, then, replace s_x by the pseudosegment s'_x of Σ that has u and v as interior points.
- If $s_v \in S$, then, replace s_v by the pseudosegment s'_v of Σ that ends in v and contains all the edges of s_v but the edge vx .
- If $s_w \in S$, then, delete s_w .

Now we have $S' \in \Sigma$, thus S' has three free points unless $|S'| = 1$.

- If $s_x \in S$, then, s_x contributes the same free points to S as s'_x to S' .
- If $s_v \in S$, then, if v was a free point for S' , then x is for S , since, in this case $s_x \notin S$. Hence, s_v contributes the same number of free points to S as s'_v to S' .
- If $s_w \in S$ and $|S'| = 1$, then, s_w contributes at least one free point to S and it covers no other points, thus S has three free points. If $|S'| > 1$ then S' has at least three free points, adding s_w does not cover any of them, and therefore, S has at least three free points.

We conclude that S has at least three free points in this case.

Case 2: If w is a corner of f , then, in this case it must be opposite to the subdivided edge, and, it follows immediately that ψ^+ is a GFAA.

So suppose w is not a corner of f . Let s_x and s_w be the pseudosegments that have x respectively w as an interior point and let s_c be the pseudosegment containing the edge from w to the corner of f which is incident to f_v . If S does not contain s_x, s_w or s_c then S is also a subset of Σ , hence, it must have at least three free points. Suppose $S \subseteq \{s_x, s_w, s_c\}$, then S has three free points, since, no two pseudosegments of $\{s_x, s_w, s_c\}$ touch twice and if $S = \{s_x, s_w, s_c\}$ there are precisely three of the six endpoints covered.

So suppose S contains at least one pseudosegment not of $\{s_x, s_w, s_c\}$. Consider the comparable set S' of Σ , that is:

- If $s_x \in S$, then, replace s_x by the pseudosegment s'_x of Σ that has u and v as interior points,
- If $s_w \in S$, then, replace s_w by the pseudosegment s'_w of Σ that has w as an interior point,
- If $s_c \in S$, then, if $s_w \notin S$, replace s_c by the pseudosegment of Σ that has w as an interior point, otherwise, delete s_c .

Now we have $S' \in \Sigma$, thus S' has three free points unless $|S'| = 1$. If $|S'| = 1$ then $S \subseteq \{s_w, s_c\}$ which contradicts the assumption that S contains at least one pseudosegment not of $\{s_x, s_w, s_c\}$, thus $|S'| > 1$.

- If $s_x \in S$, then, s_x contributes the same free points to S as s'_x to S' .
- If $s_w \in S$, then, suppose x is covered in S . Note that c , the corner of f that is now in f_v , is a vertex of the pseudosegment of Σ that contains w as an interior point. Since x is covered, $s_x \in S$, and either c is not free in S' or it is an endpoint of two pseudosegments in S' . In the latter case, c is also a free point of S as $s_x \in S$. The other endpoint of s_w is an endpoint of s'_w . We conclude that replacing s'_w by s_w leaves the number of free points intact.
- If $s_c \in S$ and $s_w \notin S$, then, s_c contributes at least as many free points to S as s'_w to S' , so assume also $s_w \in S$. The free points that s'_w contributes to S' are also free points of S as the endpoints of s'_w are included in the endpoints of the set $\{s_w, s_c\}$. Hence, S has at least three free points.

We conclude that S has at least three free points, hence, ψ^+ is a GFAA. \square

Theorem 2.35. *Every 3-connected, plane generic circuit admits an SLTR.*

Proof. A 3-connected, generic circuit can be constructed with Henneberg type 2 steps from K_4 (Berg and Jordan [BJ03]) and K_4 admits an SLTR. Every plane 3-connected generic circuit can be constructed with Henneberg type 2 steps from K_4 such that all intermediate graphs are plane. Now it follows from Theorem 2.34 that every 3-connected, plane generic circuit admits an SLTR. \square

2.3 A Flow Network for Corner Compatible Pairs

In this section, we design a two-commodity directed network \mathcal{N}_G for a given internally 3-connected, suspended graph G . An integral feasible flow in this network corresponds to an SLTR of G . Moreover an SLTR corresponds to (at least one) integral feasible flow. We start by showing a one-commodity network for which an integral feasible flow encodes a Schnyder labeling. We then show a one-commodity network for which an integral feasible flow encodes an FAA. In the end the two networks are combined in such a way that an integral feasible flow encodes a corner compatible pair. We have not been able to avoid using two commodities in the combined network. Solving the integral feasible flow problem for two-commodity networks is known to be NP-complete [EIS76].

Every network considered is directed, i.e., arcs can only be traversed in one direction. Unless stated otherwise, an arc will have capacity 1.

Encoding a Schnyder wood. We consider the primal-dual graph $G + G^*$ of G (see Figure 2.41). Here G^* is the weak dual of G together with a half-edge into the outer face for each edge that is incident to the outer face. The graph $G + G^*$ is bipartite: one vertex class consists of the edges of G , the other vertex class contains the vertices and inner faces of G . Two vertices, x, e , in $G + G^*$ are connected if x is a vertex that is an end of the edge e in G , or if x is a face that is bounded by e in G . For each edge on the boundary, a half-edge into the outer face is added. A half-edge has only one endpoint. Adding a vertex v_∞ in the unbounded face, and extending all half-edges to end in this vertex, is called *the closure of $G + G^*$* .

We will use the third part of the following theorem of Bonichon, Felsner and Mosbah [BFM07] to design a network for which an integral feasible flow encodes a Schnyder wood.

Theorem 1.12. ([BFM07]) *Let G be a suspended graph, the following structures are in bijection:*

- *The Schnyder woods of G ,*
- *The Schnyder woods of the (weak) dual G^* of G ,*
- *The α_s -orientations of the closure of $G + G^*$, where $\alpha_s(v) = 3, \alpha_s(e) = 1, \alpha_s(v_\infty) = 0$ for each primal and each dual vertex v and each edge e .*




Figure 2.41: The primal-dual graph of K_4 , a Schnyder labeling of K_4 and a Schnyder labeling of its dual drawn in $K_4 + K_4^*$.

Each vertex in $G + G^*$ that represents an edge of G has degree 4 in $G + G^*$. Moreover, in the primal and the dual Schnyder wood together in the primal-dual graph, each edge has three incoming arcs and one outgoing arc. In other words, the edge is bidirected in the primal Schnyder wood and unidirected in the dual Schnyder wood or the other way around. This follows immediately from the relation with Schnyder labelings. Recall that the primal Schnyder labeling is mapped to the dual Schnyder labeling by “moving” the label from the angle at the vertex to the angle at the face (see Figure 1.15 on page 12). From the labeling around an edge in the primal-dual graph, it follows that in one of the Schnyder woods the edge is bidirected and in the other it is unidirected. We exploit this property.

Every edge in the primal-dual graph has outdegree 1 in the orientation that represents the primal and dual Schnyder wood. For each vertex and each face of G in $G + G^*$, all but three of its incident edges are oriented inwards. In other words, a vertex v can absorb $\deg(v) - 3$ incoming arcs from its incident edges. This orientation can be encoded as a flow in a network. The half-edges on the boundary of the graph are always oriented towards the outer face. Therefore, only the interior edges are considered. The network contains a source and a sink, a node for every interior edge, a node for every vertex and a node for every interior face of G (see Figure 2.42). From the source there is an arc to every edge-node.

From an edge-node there is an arc to the nodes representing its endpoints and the nodes representing the faces it bounds. From a face-node and a vertex-node there is an arc to the sink. The arcs to the sink do not have capacity 1, but the capacity is the degree of the vertex or face minus 3. The suspensions are special as they have one outgoing half-edge into the unbounded face. Therefore, the capacity of the arc from a suspension to the sink is the degree of the suspension minus 2.




Figure 2.42: The paths in the Schnyder wood network through the edge uv .

The demand of the network is the number of edges of G . The boundary edges of G are always bidirected in the primal graph, hence, their orientation is known beforehand. A 3-orientation can be obtained from an integral feasible integral flow: the unit of flow through the edge represents its outgoing arc. All other edges of $G + G^*$ are oriented towards the edge-vertex. As the flow is integral, this implies that each edge-vertex in $G + G^*$ has outdegree precisely 1. It is immediate that every Schnyder wood yields an integral feasible flow in the network as well.

Encoding an FAA. An FAA is an assignment of vertices to faces. This can also be seen as a labeling of the angles of G . An angle is either flat or it is a convex. If an angle is flat, this is an assignment of the vertex to the face. If the angle is convex, it is a corner for the face. Hence, we want a labeling such that each face gets precisely three corners and each vertex gets at most one flat angle.

The network has a source and a sink, a node for each inner angle, a node for each non-suspension vertex, and a node for each inner face (see Figure 2.43). There is an arc from the source to each angle. From an angle there is an arc to the incident vertex and to the incident face. From each vertex there is an arc to the sink, with capacity 1, representing the assignment. From each face f there is an arc to the sink with capacity 3, representing the three corners. A unit of flow using a vertex encodes an assignment and a unit of flow using a face encodes a corner.

The demand of the network is the number of inner angles of G , as each angle has to be either a corner or a flat angle. An integral feasible flow selects a label for each angle in such a way that the conditions of an FAA (C_v and C_f) are satisfied. Therefore, there is a one-to-one correspondence between flat angle assignments and integral feasible flows in this network.

The Combined Network. In this section we explain how to build a combined network, for which a feasible integral flow encodes a corner compatible pair. Unfortunately, this will be a 2-commodity network. The flow representing the Schnyder wood and the flow representing the corners will be of type 1 (source 1 to sink 1). The flow representing the assignment will be of type 2. The combined network for a graph G is denoted by \mathcal{N}_G . An example of a graph and the network belonging to this graph is given on page 63 in




 Figure 2.43: The paths in the FAA network that go through the angle uf .

Figure 2.47. An integral feasible flow in this network is depicted in Figure 2.48 on page 64. We abuse notation and denote a path(s) which can be followed by the flow and which then represents part of the Schnyder wood (or a corner or part of the flat angle assignment) by *Schnyder wood flow* (*corner flow* and *assignment flow*, respectively).

Recall that in the network that encodes a Schnyder wood, every interior face f gets $|f| - 3$ units of flow from the edges. However, it has $|f|$ outgoing edges. Therefore, there is “space” to add the 3 units of corner flow to the Schnyder wood flow. Moreover, the unused edges will be the outgoing arcs of f in the dual Schnyder wood. The Schnyder labeling belonging to this Schnyder wood, labels all angles between two (consecutive) outgoing edges of a face with the same label. To encode the corner compatibility, we need to ensure that between every two outgoing edges of a face there is a corner selected.

We encode corner compatibility with a cyclic structure around a face-node. We use Figure 2.44 to introduce this structure. Recall that the network is directed, i.e., all arcs can only be used in one direction. For a *Schnyder wood flow* through a vertex nothing has changed, i.e., such a path will look as in Figure 2.42.


 Figure 2.44: The subnetwork of a face f_{tuvw} is depicted. In red a Schnyder wood path, in blue a corner path and in green an assignment path.

A *Schnyder wood flow* through a face f_{tuvw} is depicted in red in Figure 2.44. On such a path

there is one extra node compared to the network in Figure 2.42. This extra node is denoted by “small square”. This node ensures that the arc into the face can be used by Schnyder wood flow as well as corner flow.

The corner flow comes into the subnetwork via an angle (see the blue path in Figure 2.44). Angles are drawn as triangles in the figures. An angle consists of three arcs. The first arc ensures that the angle is either assigned or a corner, but not both. The second arc has no special task. The third arc ensures that no two corners with the same label are selected.

The *corner flow* uses all three arcs in the triangle through which it enters the subnetwork of the face. Then it proceeds to the first small square, or it proceeds to the third arc of the next angle. It has to go into the face via a small square. The fact that the network is saturated is used to prove that an integral flow encodes a corner compatible pair. Informally, every small square has to be used. Every third arc of an angle, can only be used by one unit of integral flow. So a small square can only be “skipped” by corner flow if it is used by the Schnyder wood flow, which in turn means that the clockwise next angle has the same label. This implies that there is a corner selected between every two “available” small squares, as otherwise the network would not be saturated.

The *assignment flow* uses the first arc of an angle (see the green path in Figure 2.44, detailed view in Figure 2.45). From this arc it goes into a vertex dummy. This is denoted by *dummy vertex* as it is not the vertex-node that is used by Schnyder wood flow. From the vertex dummy there is an arc to the sink. The dummy vertices ensure that at most one angle of a vertex is assigned.




Figure 2.45: Possible assignment paths for a vertex v . The node that ensure that at most one of the angles of v is assigned is denoted by v^* or by “dummy of v ”.

The network that represents an FAA is splitted. To ensure that the corner flow and the assignment flow together form an FAA, we introduce some extra nodes. A *face-bag* is added for each inner face. From source 1 there is an arc of capacity 3 into each face-bag. From the face-bag there is an arc to each of the angles of the face.

The assignments are encoded by type 2 flow. To ensure that the correct number of assignments is made, there is a dummy source added before source 2. The arc from the dummy source to source 2 has capacity equal to the number of assignments needed. Note that, in the network that encodes only an FAA, this is not necessary, as the demand of network and the sum of the capacities of the arcs to the sinks are equal.

The Demands. Let E_{int} be the set of interior edges and F_{int} the set of interior faces of the graph. To represent a Schnyder wood there should be a Schnyder wood flow of value $|E_{\text{int}}|$. To make sure every face has three corners there should be a corner flow of value $3|F_{\text{int}}|$. The number of flat angles needed is $\sum_{f \in F_{\text{int}}} (|f| - 3)$. A union of a type 1 (from source 1

to sink 1) and a type 2 (from source 2 to sink 2) flow $\phi = (\phi_1, \phi_2)$ in this network is called feasible if

- $\text{value}(\phi_1) = |E_{\text{int}}| + 3|F_{\text{int}}|$ and
- $\text{value}(\phi_2) = \sum_{f \in F_{\text{int}}} (|f| - 3)$.

Remark. If the sources and sinks were unified, we would obtain a one-commodity network. However, in such a case, the assignment flow and corner flow could switch places. That is, from a face-bag, a unit of flow can go to the sink via a vertex (it behaves as corner flow and then as assignment flow). Or from the unified source it goes via the assignment source to an angle and then via the face to the sink (it behaves as assignment flow and then as corner flow). This implies that a solution might not be corner compatible.

Another option is to not control the flow before going into an angle (i.e., from the source it goes immediately to an angle) and define it as an assignment if it leaves through a vertex and as a corner if it leaves through a face. In this case, the Schnyder wood flow and the corner flow cannot be controlled. An integral feasible flow could have a face with too much Schnyder wood flow and too little corners and another one with the opposite. In other words the property C_f of the FAA and the property of the flow given by the Schnyder wood might be violated.

Theorem 2.36. *Let G be an internally 3-connected, suspended graph. Then G has an SLTR if and only if there is an integral feasible flow $\psi = (\psi_1, \psi_2)$ in \mathcal{N}_G .*

Proof. Suppose G be an internally 3-connected suspended graph and $\phi = (\phi_1, \phi_2)$ be an integral feasible flow in \mathcal{N}_G . First we will show that from the feasible flow we can extract a Schnyder wood σ and an FAA ψ for G and then that this is a corner compatible pair. From ϕ_1 we have to extract the Schnyder wood flow and the corner flow.

The total amount of flow from source 1 to sink 1 is bounded from above by the sum of the capacities of vertex-to-sink and face-to-sink arcs. Recall that the arc from a suspension vertex s_i to the sink, has capacity $\deg(s_i) - 2$. The total amount of flow between source 1 and sink 1 adds up to the value of ϕ_1 , since, ϕ is a feasible flow. From the following calculation it follows that every feasible flow saturates all the arcs to sink 1.

$$\sum_{f \in F_{\text{int}}} \deg(f) + \sum_{v \in V} (\deg(v) - 3) + 3 = |E_{\text{int}}| + 3|F_{\text{int}}|$$

The capacities of the arcs leaving source 1, also add up to the value of ϕ_1 . Recall that the arcs leaving source 1 are the arcs to the edges, $|E_{\text{int}}|$, and the arcs to the face-bags, $3|F_{\text{int}}|$. It follows, that through each face-bag 3 units of flow are routed, and the only way to reach the sink is through this face. Therefore, each face f has at most $\deg(f) - 3$ of Schnyder wood flow routed through it.

As the arcs from the vertices to sink 1 are saturated, there must be $\deg(v) - 3$ units of type 1 flow through a non-suspension v and there must be $\deg(s_i) - 2$ units of type 1 flow through s_i . The rest of the flow that is routed through the edges must go via a face, and since every face has at most $\deg(f) - 3$ units of flow of this type routed through, it must be precisely this amount.

$$|E_{\text{int}}| - \sum_{v \in V} (\deg(v) - 3) - 3 = \sum_{f \in F_{\text{int}}} \deg(f) - 3|F_{\text{int}}| = \sum_{f \in F_{\text{int}}} (\deg(f) - 3)$$

To extract a Schnyder wood from ϕ_1 , the Schnyder wood flow must be of value $|E_{\text{int}}|$, i.e., the number of interior edges. Moreover, it must be such that:

- precisely $\sum_{v \in V} \deg(v) - 3$ edges are appointed to a vertex,
- precisely $\sum_{f \in F_{\text{int}}} \deg(f) - 3$ edges are appointed to a face, and,
- one more edge is appointed to each suspension.

The calculations above show that this is exactly the case, hence, we can extract the Schnyder labeling σ .

The number of interior angles is equal to the amount of corner flow plus the amount of assignment flow. Therefore, the first arc of each angle (i.e., the first arc drawn inside the triangles) must be saturated as well.

Integrality of the flow now implies that each angle in the graph is either a corner or an assigned angle. The corner flow ensures that each face has precisely three corners. The flow of type 2 ensures that each vertex has at most one assigned angle. This gives us an FAA ψ .

Left to show is that σ and ψ are corner compatible.

Consider the subnetwork at a face f . Its three corners cannot use the small squares that are used by the flow that represents the Schnyder wood. There are precisely three small squares which are not used by the flow that represents the Schnyder wood, we call these three small squares *available*. We trace the units of flow backwards from the face to the available small squares. When it leaves a small square (backwards) it goes to the first angle (triangle) counterclockwise. It can move to the next counterclockwise angle, but we claim this only occurs when the small square between them is not available. In other words, between every two available small squares there is an angle through which a unit of flow comes into the subnetwork.

Claim 1. Let Q_1 , Q_2 and Q_3 be the available small squares in the subnetwork of f in clockwise order. The unit of flow that leaves the subnetwork via small square i , enters the subnetwork at an angle between Q_{i-1} and Q_i (clockwise).

Suppose not, without loss of generality, let there be a unit of flow leaving via Q_3 which does not enter between Q_2 and Q_3 . Suppose the flow leaving via Q_3 enters the network at angle α which is between Q_1 and Q_2 . Consider the clockwise last angle before Q_2 . The third arc in this angle is used by the flow which goes through Q_3 , hence, there is no possibility to enter Q_2 . This contradicts the fact that the network is saturated (as no flow is routed through Q_2 , see Figure 2.46). The same argument applies when the angle α is between Q_3 and Q_1 . This proves Claim 1. \triangle

Claim 2. All angles between two consecutive available small squares are labeled with the same label in σ .

This follows immediately from the definition of the Schnyder labeling in the dual graph. Between the outgoing edges of label i and $i + 1$ of a face, all angles are labeled $i - 1$ (see Figure 1.15 on page 12). The available small squares are located on the place of the outgoing edges of the face in the Schnyder wood. \triangle

Claim 1 and Claim 2 together prove that σ and ψ are corner compatible.

On the other hand, suppose G admits an SLTR. By Theorem 2.19 there is a Schnyder labeling which satisfies the rule of the corners for the FAA that belongs to this SLTR (note that this is not necessarily unique). Consider a complying Schnyder labeling, and the FAA.




Figure 2.46: Suppose the red arrows are part of the Schnyder wood flow, the blue arrows of the corner flow. The top small square is skipped by a unit of corner flow and can no longer be saturated.

Set up the network \mathcal{N}_G . Start with the empty flow and add the following units of flow:

- For each interior edge with labels $i+1, i, i, i-1$ in clockwise order at its ends, a unit of flow from source 1, to this edge, to the neighbor in \mathcal{N}_G which is between the two labels i in G , and then to sink 1 is added.
- For each interior face f , for each corner α of f a unit of flow is added that goes from source 1, into the face-bag of f , then to the angle α , to the first clockwise available small square, into f and then to sink 1.
- For each flat angle β , a unit of flow is added that goes from the dummy source to source 2, to angle β to the appropriate dummy vertex, to sink 2.

It is easy to check that the flows add up to the appropriate values and that no capacity constraint is violated. This concludes the proof. \square

2.4 Conclusion

We have given necessary and sufficient conditions for a 3-connected planar graph to have an SLT Representation. Given an FAA and a set of rational parameters $\{\lambda_i\}_i$, the solution of the harmonic system can be computed in polynomial time. Checking whether a solution is degenerate can also be done in polynomial time. Hence, we can decide in polynomial time whether a given FAA corresponds to an SLTR.

Given a 3-connected planar graph and a GFAA, interesting optimization problems arise, e.g.

find the set of parameters $\{\lambda_i\}_i$ such that the smallest angle in the graph is maximized, or the set of parameters such that the length of the shortest edge is maximized. The solution of the system of equations will give the coordinates of the vertices in a SLTR, however there is no restriction on what they may look like. Is there a way to select the parameters and locations of the suspensions such that all coordinates are integers?

The setup of FAA and the discrete harmonic system works for any assignment, e.g., also for a flat angle assignment in which each face gets four not assigned vertices, or five, or it is a different value for each face. Does a relaxation like “each face has at most four corners” make it easier to find a GFAA?


We have given another characterization of graphs that admit straight-line triangle representation in terms of flat angle assignments and Schnyder labelings. For graphs that have a unique Schnyder labeling (these graphs are identified by Felsner and Zickfeld [FZ08]), the problem of deciding whether the graph has an SLTR can be translated into a matching problem in a bipartite graph. For graphs with very few Schnyder labelings the problem also becomes polynomially tractable. However, there are planar 3-connected graphs on n vertices that have 3.209^n Schnyder labelings [FZ08].

With the new conditions we have shown a translation of the drawing problem into a flow optimization problem. However, finding an integral feasible flow in a two-commodity network is known to be NP-complete. An interesting question is whether for this particular network a feasible solution always implies the existence of an integral feasible solution.

Gonçalves, Lévêque and Pinlou conjectured that every 3-connected planar graph admits a primal-dual contact representation by right triangles, where all triangles have a horizontal and a vertical side and the right angle is bottom-left for primal vertices and top-right otherwise [GLP12]. To the best of our knowledge this is still open. Perhaps the new proof could give more insight into this problem.

Unfortunately, we have to leave the main problem open:

Is the recognition of graphs that have an SLTR (GFAA) in P ?



Source 1 to every edge (1)
 Source 1 to every bag (3)
 Every bag to every of its angles (≤ 1)
 Source 2 to every angle (≤ 1)
 Every face f to Sink 1 ($|f|$)
 Every vertex v to Sink 1 ($\deg(v) - 3$)
 Every dummy vertex v^* to Sink 2 (1)





Figure 2.47: Example of the inner network belonging to the graph on the bottom right. On the bottom left the descriptions of the arcs that are not drawn. The grey disk below a vertex represents the dummy vertex.



- Source 1 to every interior edge (1)
- Source 1 to every bag (3)
- Every bag to every of its angles (≤ 1)
- Source 2 to every angle (≤ 1)
- Every face f to Sink 1 ($|f|$)
- Every interior vertex v to Sink 1 ($\deg(v) - 3$)
- Every suspension s to Sink 1 ($\deg(s) - 2$)
- Every dummy vertex v^* to Sink 2 (1)




Figure 2.48: A feasible flow in the inner network belonging to the graph on the bottom right. The flow results in the Schnyder wood and the assignment given in the graph on the bottom right, the vertex e is assigned to the face I . On the bottom left the descriptions of the arcs that are not drawn. The grey disk below a vertex represents the dummy vertex.

“Even if we could turn back, we’d probably never end up where we started.”

Haruki Murakami, 1Q84

3

Touching Triangle Representations

In this chapter we are dealing with graphs in the dual setting. Given a graph G , is there a representation of G such that each vertex is a triangle and two vertices are adjacent if and only if their triangles touch. Many variants of such a representation have been studied. The edges can be represented by point contacts or by side contacts. In the later case one can require that there are no holes between the triangles, i.e., there is no region enclosed by some triangles which does not represent a vertex of the original graph. In this case there can be a requirement on the shape of the boundary of the union of the triangles. This representation has been studied when there is no restriction to the boundary, when the boundary is a convex polygon and when the boundary is restricted to be a triangle. Here are the main structures of this chapter.

Definition 3.1 (Touching Triangle Representation). A planar graph G admits a *touching triangle representation* (TTR) if it admits a drawing such that every vertex is represented by a triangle and a side contact between two triangles exists if and only if the two vertices are adjacent. Moreover there is no region enclosed by triangles which does not represent a vertex of the original graph. A drawing that is a tiling of a convex polygon by triangles is a *convex touching triangle representation* (cTTR). If the convex polygon is of size k , then the representation is denoted by k TTR. In particular, if the convex polygon is a triangle, then the representation is denoted by 3TTR. Note that the latter has appeared in the literature under the name *proper touching triangle representation*.




Figure 3.1: A TTR, a convex-TTR and a 3TTR.

There is an obvious connection between SLTRs and TTRs, 3TTRs are SLTRs and the underlying graph of a TTR is the weak dual of an SLTR. However, the study of the two objects has been mostly independent. In this chapter there will be clear connections between SLTRs and TTRs.

In Section 3.1, we will discuss the known results about representations with triangles.

In Section 3.2 we will present some new results. First we show that Halin graphs have a 3TTR, Halin graphs are a subclass of 2-outerplanar graphs. We will also give a complete characterization of the biconnected outerplanar graphs that have a cTTR.

3.1 Triangle Representations (Known Results)

De Fraysseix, Ossona de Mendez and Rosenstiehl proved that every planar graph has a triangle contact representation [dFdmR94]. In a triangle contact representations the edges are represented as point contacts (see Figure 3.2).

Theorem 3.2 ([dFdmR94]). *Every planar graph has a triangle contact representation.*

The proof is constructive; we will describe the algorithm that produces the drawing. First the graph is augmented to a maximal planar graph (triangulation) by adding dummy vertices. Consider the canonical order (v_1, \dots, v_n) and a Schnyder wood that comes from this canonical order, of the triangulation. The triangles of v_1, v_2 and v_n are set first, as shown in Figure 3.2, with their bases at height 1, 2 and n , respectively. At step $k = 3, \dots, n-1$ the triangle of v_k is added in the following way. For $i = 1, 2, 3$ let ϕ_i be the parent of v_k in the tree T_i of the Schnyder wood. The base of the triangle of v_k is placed on height k , reaching from the neighbor ϕ_1 to the neighbor ϕ_2 . The top is placed in the middle of the base, the height is the placement of ϕ_3 in the canonical order. Afterwards the triangles that represent dummy vertices are removed to obtain a triangle contact representation of the original graph.




Figure 3.2: The base triangles on the left and an example of a triangle contact representation constructed by the algorithm of De Fraysseix, Ossona de Mendez and Rosenstiehl.

From this construction it follows that every planar graph has a contact representation by:

- Isosceles triangles (as in the construction above).
- Right-angled triangles (by placing the top above the left end of the base, instead of above the middle of the base).
- T-shapes (by inscribing a \perp in each triangle and extending the sides of each horizontal bar).
- Y-shapes (similar as T-shapes).

Every planar graph has a triangle contact representation, however, not all planar graphs admit a *touching* triangle representation. An example is given in Figure 3.3. Gansner, Hu and Kobourov give necessary conditions for graphs that admit a TTR [GHK10].

Lemma 3.3 ([GHK10]). *A planar graph that has a touching triangle representation satisfies the following two conditions.*

- (a) *Two neighbors have at most three common neighbors and the graph induced by these common neighbors has at most one edge.*
- (b) *Two non-neighbors have at most four common neighbors and the graph induced by these common neighbors has at most two edges.*

From (a) it follows that the graph in Figure 3.3 has no TTR. Intuitively, three common neighbors need to use all the sides of at least one of the triangles that represent x and y , hence, there is no possibility to add a fourth common neighbor.




Figure 3.3: The graph on the left does not have a touching triangle representation. On the right a way to represent two vertices with three common neighbors.

These conditions are necessary but not sufficient. The graph in Figure 3.4 (a) satisfies conditions (a) and (b) of Lemma 3.3, however it does not have a touching triangle representation. Suppose it does have a touching triangle representation, then the graph in (b) is a subdivision of the skeleton of the TTR of (a). The red and the blue area in (b) together need four vertices with a straight angle inside, as there are two 4-faces and one 5-face. Therefore, at least one of e, f, h and i is assigned inside this area. On the other hand, i, e and h cannot be assigned inside the red area, as otherwise the red area has less than three combinatorial convex corners. Moreover, e, h, f cannot be assigned inside the blue nor the red area, as otherwise the union of the blue and red area has less than three combinatorial convex corners. But then none of e, f, h, i may be assigned inside the red or blue area, contradiction.




Figure 3.4: A graph that satisfies conditions (a) and (b) of Lemma 3.3, but it does not have a touching triangle representation

On the positive side, Gansner et al. have shown that that every biconnected outerplanar graph, every grid graph and every hexagonal grid graph admits a TTR.

Theorem 3.4 ([GHK10]). *Every biconnected outerplanar graph admits a TTR.*

The proof is constructive, we will describe the algorithm that produces the drawing.

The algorithm depends on a peeling order which is well-defined for biconnected outerplanar graphs.

Definition 3.5 (Reversed Peeling Order). Let G be a biconnected outerplanar graph. A reversed peeling order of G , is a decomposition of the vertex set into subsets and an ordering of these subsets, s_1, \dots, s_k , such that

1. The graph induced by the first i sets, $G[\bigcup_{j=1}^i s_j]$, is connected.
2. In $G[\bigcup_{j=1}^i s_j]$ the vertices of s_i have degree two and form a path on the boundary, which bounds one interior face.

3. $G[s_1]$ is a K_2 with an edge on the outer face of G .

A reversed peeling order is used to construct a TTR.

1. Compute an outerplanar embedding of G .
2. Compute a reverse peeling order of G .
3. Insert triangle(s) corresponding to the current set of vertices in the peeling order, while maintaining a concave upper envelope.

While inserting the triangles, the following invariant is maintained: Every triangle has an exposed side in the upper envelope and the upper envelope is concave. The algorithm is depicted in Figure 3.5. As the upper envelope is concave and a new set of vertices according to the reversed peeling order is added between two current triangles, the angle looks like Figure 3.5 (a). In Figure 3.5 (b) it is shown how to add four triangles in this concave angle, such that all the new angles in the upper envelope are again concave. In Figure 3.5 (d) the TTR of the graph in Figure 3.5 (c) is drawn according to this algorithm. Unfortunately, the triangles become skinny very quickly.




Figure 3.5: Adding four triangles in a concave angle (a and b), a biconnected outerplanar graph G (c) and a TTR of G according to the algorithm of Ganser, Hu and Kobourov (d).

The two algorithms presented above depend on an order in the primal graph. In the first algorithm there is space to add the next set of triangles according to the order by construction, in the second algorithm the space is ensured by the concave angles in the upper envelope. We proceed with an algorithm by Fowler, who adds the next set of triangles inside an existing triangle. This way, the shape of the outer face is maintained.

Let G be a biconnected outerplane graph. We denote by $\text{VEIN}(G)$ the graph consisting of all strictly interior edges of G (the chords) and their endpoints. If $\text{VEIN}(G)$ is connected and G is biconnected, then G is called *strongly connected*.

Definition 3.6 (Chord-to-Endpoint Assignment). A *chord-to-endpoint* assignment is an assignment of chords to endpoints, in such a way that

- at most one chord is not assigned and
- each endpoint has at most one chord assigned to it.




Figure 3.6: The edges of the VEINs are solid, the other edges dashed. A strongly connected graph whose VEIN has two interior faces (a), a strongly connected graph whose VEIN has more than two interior faces and a graph that is not strongly connected.

In the leftmost drawing of Figure 3.9, a chord-to-endpoint assignment is given by the blue arrows, i.e., the head of the arrow on a chord points toward the endpoint this chord is assigned to. The chord-to-endpoint assignment is an important tool for constructing triangle representations of outerplanar graphs. It was first used by Fowler in the proof of the following theorem. Note that Fowler used the term proper TTR instead of 3TTR.

Theorem 3.7 ([Fow13]¹). *A strongly connected outerplanar graph G has a 3TTR if and only if $\text{VEIN}(G)$ has at most two interior faces.*

Proof. Let G be a strongly connected outerplanar graph such that $\text{VEIN}(G)$ has at most two interior faces. We will show that G has a chord-to-endpoint assignment and construct the 3TTR using the chord-to-endpoint assignment. It is obvious that G has a chord-to-endpoint assignment. Since $\text{VEIN}(G)$ has (at most) two interior faces, it has (at most) one edge more than it has vertices. In a chord-to-endpoint assignment, all but one of the chords have to be assigned. If $\text{VEIN}(G)$ has two interior faces, select one chord of one of the interior faces of $\text{VEIN}(G)$ to be the not assigned chord, call this chord the *starting chord*. The graph $\text{VEIN}(G)$ deleted the starting chord has (at most) one cycle and therefore as many vertices as edges. Hence, for every chord c there is an endpoint v to which we can assign c such that each vertex has at most one chord assigned to it. If G has at most one interior face, then $\text{VEIN}(G)$ has a chord-to-endpoint assignment such that all chords are assigned to one of their endpoints.

From the chord-to-endpoint assignment we construct an order of the faces of the graph. If G has two interior faces, start with the two faces incident to the starting chord. In each step a chord on the boundary of the current set of faces is chosen and the face on the other side of this chord is added next. As the graph is biconnected and $\text{VEIN}(G)$ is connected, this process ends when all the faces have been added. If G has at most one interior face, then any face of G can be chosen as the starting face.




Figure 3.7: The base cases when $\text{VEIN}(G)$ has at most one interior face (a) and when G has two interior faces (b).

First a 3TTR is drawn for the starting face(s), as shown in Figure 3.7. The faces are added

¹The ‘only if’ part of the proof slightly differs from the original proof.

in the order chosen before. To add the face f , consider the chord that disconnects it from the already added faces. The new vertices of the face are added into the triangle that represents the vertex to which the chord is assigned. A vertex is called *processed* if another vertex has been added into the triangle that represents this vertex. During the construction the following invariant is maintained.

Invariant: The triangles of vertices that have been added but are not (yet) processed share a boundary segment with the outer region.

Obviously the invariant holds in the base cases, as all triangles share a boundary segment with the outer region. The addition of a face starts by dividing the triangle of the endpoint vertex into two triangles such that one of the triangles has side-contacts only with the triangles that represent the chord. In Figure 3.8 such a step is shown where the new vertices of the face are l_1, \dots, l_k , the chord connecting this face is n_1x and the vertices are added into the triangle of x . After dividing the triangle of x into two triangles, the triangle that does not share a boundary segment with the outer region and becomes the triangle that represents x . The vertex x is now processed. The other triangle is divided into appropriately many triangles that all share a boundary segment with the outer region. Hence, every vertex that has not yet been processed, shares a boundary segment with the outer region. All the faces can be processed this way, and we end up with a 3TTR of G .

Note that all the steps are actually Henneberg type 2 extensions of the representation.




Figure 3.8: Adding the path l_1, \dots, l_k into the triangle of x , where x has neighbors n_1, n_2, n_3 and the chord to which the new face is attached, is xn_1 .

In the “only if” part of this proof, we use a counting argument to show that if a strongly connected outerplanar graph $G = (V, E)$ has three interior faces in $\text{VEIN}(G)$, then G does not admit a 3TTR. Let H be the graph that is the supposed 3TTR of G . The number of interior vertices of H is the number of interior faces of G , denoted by ϕ_G . The maximum number of flat angles that there can be in H is the number of interior vertices of H , ϕ_G . Moreover, there is one more interior face in G than that there are chords in G . This follows from the fact that $\text{VEIN}(G)$ is connected and G is biconnected. The number of chords is denoted by c_G . Let $V_{\geq 3}$ be the set of vertices of G that have degree at least three, $v_{\geq 3}$ the cardinality of $V_{\geq 3}$, v_2 the number of degree-2 vertices in G and q the number of interior faces of $\text{VEIN}(G)$. Firstly, we express ϕ_G in terms of $v_{\geq 3}$ and q . Note that $v_{\geq 3} - c_G + q + 1 = 2$ is given by Euler’s formula on $\text{VEIN}(G)$.

$$\phi_G = c_G + 1 = v_{\geq 3} - 1 + q + 1 = v_{\geq 3} + q$$

Each interior face corresponds to an interior vertex of H , i.e., a vertex that may admit a

flat angle. The number of flat angles needed, is bounded below by:

$$\sum_{v \in V_{\geq 3}} (\deg(v) - 3) = 2e_G - 2v_2 - 3v_{\geq 3} = 2e_G - 2v_G - v_{\geq 3} = 2\phi_G - 2 - v_{\geq 3}$$

where v_G and e_G are the number of vertices and edges in G . Hence, we obtain as a constraint

$$\phi_G \geq 2\phi_G - 2 - v_{\geq 3} \quad \Rightarrow \quad \phi_G \leq v_{\geq 3} + 2$$

and, therefore, q is at most 2, i.e., $\text{VEIN}(G)$ has at most two interior faces. \square




Figure 3.9: A strongly connected outerplanar graph, with a chord-to-endpoint assignment given by the blue arrows, and a 3TTR according to the algorithm of Fowler.

Recently, Chang and Yen showed that deciding whether a biconnected internally cubic graph has a convex-TTR can be done in polynomial time, and their algorithm will also output a cTTR if one exists [CY]. This is the first result that gives a decision algorithm that is not a characterization. Chang and Yen describe how to obtain a graph H such that H has a convex-SLTR if and only if the input graph has a cTTR. To obtain the cTTR from H , they use flat angle assignments of the inner graph of H . Such an assignment is denoted by cFAA.

Definition 3.8 (cFAA). A cFAA is an assignment of vertices to faces such that:

- $[C_v]$ Every interior vertex is assigned at most once, the vertices on the boundary are not assigned and
- $[C'_f]$ Every interior face f has $|f| - 3$ vertices assigned to it.

Chang and Yen explored what cFAAs of H may look like if the input graph is biconnected and internally cubic. They define obstructions to be subsets which are not stretchable or do not have sufficiently many combinatorial convex corners. Strict obstructions are obstructions that cannot be avoided. Chang and Yen prove that, in the case of biconnected internally cubic graphs, if there is no strict obstruction, then the cFAA can be changed to a good cFAA. Checking whether there is a strict obstruction can be done in polynomial time. Later, we will characterize biconnected outerplanar graphs that have a k TTR, and we use an auxiliary graph H similar to the graph defined by Chang and Yen. Moreover, it may be clear that the result of Chang and Yen also applies to biconnected outerplanar graphs.

3.2 Results on 3TTRs

In this section we will present two new results. The first result considers Halin graphs, a subclass of 2-outerplanar graphs. Every Halin graph admits a proper touching triangle

representation and the proof constructs the representation similarly as the construction of 3TTRs of strongly connected outerplanar graphs. The second result is a characterization of biconnected outerplanar graph that admit a cTTR, i.e., a TTR in a convex polygon. To show that a biconnected outerplanar graph that satisfies the conditions admits a cTTR, structures of the previous section are used, e.g., chord-to-endpoint assignments and flat angle assignments of the inner graph.

3.2.1 Halin Graphs

In the previous chapter, Theorem 2.34, we have shown that an SLTR can be extended with a Henneberg type 2 step. Given an SLTR or equivalently a GFAA, one can extend the graph with a Henneberg type 2 step, and the resulting graph again has an SLTR (GFAA). Recall that a Henneberg type 2 step consists of subdividing an edge uv and connecting the new vertex x to a third vertex in the current graph, not u or v (see page 5).

Definition 3.9 (Halin Graph). A graph is a Halin graph if it is formed by embedding a tree, that has no vertices of degree 2, in the plane, and connecting its leaves by a cycle that crosses none of its edges.

Notations. Before we proceed to show that every Halin graph admits a 3TTR, we introduce some notations (see Figure 3.10). Let G be a plane Halin graph. The underlying tree of G , i.e. the graph that remains after removing all edges incident to the unbounded face, is denoted by T . The tree induced by the branching nodes of T , i.e. the nodes of degree at least three in T , is denoted by \mathcal{Z} . With $t(v)$ we denote the triangle that represents v in the 3TTR.

We will prove that every Halin graph has a 3TTR. The proof is based on induction. The base case is the tree induced by one leaf node of \mathcal{Z} . This graph is a wheel (with at least three spokes). Therefore, the first step is to show that a wheel on at least 4 vertices admits a 3TTR.

Lemma 3.10. *The wheel W_k on $k \geq 4$ vertices admits a 3TTR.*

Proof. $W_4 = K_4$ has a 3TTR as shown in Figure 3.11. The 3TTR of W_k can be constructed from the 3TTR of W_{k-1} by one Henneberg type 2 step as follows. Let v be the axle vertex of the wheel and consider the boundary of the triangle $t(v)$ that represents v in the 3TTR. Subdivide an edge on the boundary of $t(v)$ and connect the new vertex to one of the corners of the adjacent triangle that is not $t(v)$. This is a Henneberg type 2 step and by Theorem 2.34 the new graph has an SLTR, hence, by induction, W_k admits a 3TTR. \square




Figure 3.10: The Halin graph, T and \mathcal{Z} .




Figure 3.11: A 3TTR of W_4 and W_8 .

Theorem 3.11. *Every Halin graph admits a 3TTR.*

Proof. Let G be a plane Halin graph. We construct a 3TTR of G by induction on the vertices of \mathcal{Z} . Consider a leaf of \mathcal{Z} , label the leaf 1. Label the other vertices of \mathcal{Z} according to a DFS of \mathcal{Z} , starting from the leaf with label 1. Let L_i be the set of vertices of \mathcal{Z} with label at most i . By T_i we denote the subtree of T , induced by L_i and the neighbors of L_i in T , denoted by $N(L_i)$. By $H(T)$ we denote the Halin graph obtained from the tree T (see Figure 3.12). We construct a 3TTR of $H(T_i)$ iteratively.




Figure 3.12: A Halin graph $H(T)$ with the branching nodes tree \mathcal{Z} in black, the additional edges of T in dashed black and the cycle connecting the leaves in double dashed grey. On the right $H(T_1)$, $H(T_2)$ and $H(T_3)$ are shown.

In every step we maintain the following invariant:

Every leaf of T_i is represented in the 3TTR of $H(T_i)$, by a 3-face, which has a vertex incident to the outer region, or by a 4-face, which has an edge incident to the outer region and the vertex with a stretched angle is an interior vertex in the 3TTR.

Base Case. The tree T_1 consists of the root vertex v_1 of \mathcal{Z} according to the labeling, and all its neighbors in T . Note that v_1 has degree at least 3 in T , therefore $H(T_1)$ is a wheel with at least 4 vertices. By Lemma 3.10, $H(T_1)$ has a 3TTR. The 3TTR obtained by following the construction of the proof of Lemma 3.10, satisfies the invariant.

Induction Step. Consider a 3TTR Δ of $H(T_{i-1})$ such that the invariant holds. Let v_i be the vertex of \mathcal{Z} with label i and $t(v_i)$ the face that represents v_i in Δ . Since v_i is a leaf of T_{i-1} and the invariant holds, $t(v_i)$ is either a 3-face or a 4-face with one stretched angle.

Case 1. Suppose $t(v_i)$ is a 3-face and $t(v_{i-1})$, $t(n_1)$ and $t(n_2)$ are the neighboring triangles (see Figure 3.13 (a)). Since v_i has at least degree 3 in T_i and precisely degree 1 in T_{i-1} , there are at least two vertices introduced in this step. Let l_1, \dots, l_k be the leaf neighbors of v_i in H such that $n_1, l_1, \dots, l_k, n_2$ appear consecutively on the boundary of the outer face of $H(T_i)$. Stack a degree three point in $t(v_i)$ and connect the point to the three corners of $t(v_i)$. Let the region that shares an edge with $t(v_{i-1})$ be the new $t(v_i)$. The region that shares an edge with $t(n_1)$ is the triangle that represents l_1 and the third region is the triangle of l_k .

For each $j = 2, \dots, k-1$ subdivide the edge between $t(l_k)$ and $t(v_i)$ and connect the new vertex to the vertex on the outer boundary, the region that shares an edge with $t(l_{i-1})$ becomes $t(l_i)$ and the other region becomes $t(l_k)$. Note that we stacked three triangles into a triangle, and furthermore only used Henneberg type 2 steps.

Hence, we have obtained a 3TTR of $H(T_i)$. Consider a leaf l of T_i . If this is also a leaf in T_{i-1} then for $t(l)$ the invariant still holds. If l is not a leaf in T_{i-1} , then it has been constructed in this step. From the construction it follows that $t(l)$ is a 3-face with one point on the boundary of the 3TTR.





Figure 3.13: Extension of a 3-face (a) and 4-face (b)

Case 2. Suppose $t(v_i)$ is a 4-face and $t(v_{i-1})$, $t(n_1)$ and $t(n_2)$ are the neighboring triangles such that the point s , between $t(v_i)$, $t(v_{i-1})$ and $t(n_2)$ has a stretched angle with respect to $t(v_i)$ (see Figure 3.13 (b)). Since v_i has at least degree 3 in T_i and precisely degree 1 in T_{i-1} , there are at least two vertices introduced in this step. Let l_1, \dots, l_k be the leaf neighbors of v_i such that $n_1, l_1, \dots, l_k, n_2$ appear consecutively on the boundary of the outer face. Subdivide the edge between $t(v_i)$ and the outer region, and connect the new point to the point between $t(v_i)$, $t(v_{i-1})$ and $t(n_2)$. Subdivide this new edge and connect the new point to s (the point between $t(v_i)$, $t(v_{i-1})$ and $t(n_1)$). The new region incident to $t(n_2)$ is $t(l_k)$, the region incident to $t(n_1)$ is $t(l_1)$ and the third region is $t(v_i)$.

For each $j = 2, \dots, k-1$ subdivide the edge between $t(l_k)$ and $t(v_i)$ and connect the new point to the point on the outer boundary, that is a corner of $t(l_2)$ and $t(n_2)$. The region that shares an edge with $t(l_{i-1})$ becomes $t(l_i)$ and the other face becomes $t(l_k)$. Note that throughout this construction, we have only used Henneberg type 2 steps.

Hence, the resulting drawing must have a GFAA by theorem 2.34. Therefore, $H(T_i)$ has a 3TTR. Consider a leaf l of T_i . If this is also a leaf in T_{i-1} then for $t(l)$ the invariant still holds. If l is not a leaf in T_{i-1} then by the construction above $t(l)$ is a 3-face unless it is $t(l_1)$ or $t(l_2)$, the first and second leaf neighbor of v_i . The 3-faces share a vertex with the outer region. From the construction it follows that $t(l_1)$ and $t(l_2)$ are 4-faces that share one edge with the outer region. Therefore the invariant still holds.

This concludes the induction step, and also the proof. \square


 Figure 3.14: The constructed 3TTR of $H(T_3)$ of Figure 3.12.

3.2.2 Biconnected Outerplanar Graphs

In this section we will characterize the biconnected outerplanar graphs that admit a cTTR. First we construct the graph H such that G is the weak dual of H . We aim for a representation whose skeleton is H . To construct H , we start with the weak dual of G . Through each

boundary edge of G , an edge that connects the inner face with a new vertex in the outer face, is added. The newly added points are cyclically connected. Contrary to the auxiliary graph in the algorithm of Chang and Yen [CY], the boundary edges are not subdivided but contracted, if possible. Every boundary edge, whose contraction does not induce a 2-face, is contracted. Throughout this section, the graph H will be called the *auxiliary graph* of G . An example is given in Figure 3.16 (a) on page 77. To find a cTTR of which H is the skeleton, we will construct an assignment of flat angles in H and then show that this assignment is good, i.e., there exists a cTTR in which precisely the prescribed angles are stretched.

By $\text{VEIN}(G)$ we denote the graph consisting of all strictly interior edges of G and their endpoints. We will make use of the relations between the components of $\text{VEIN}(G)$. In Figure 3.15 the main objects are depicted.

Definition 3.12 (Venation Graph). The vertices of the venation² graph of $\text{VEIN}(G)$ are the components of $\text{VEIN}(G)$ and the faces in G that connect two or more components of $\text{VEIN}(G)$. There is an edge between a component and a face if and only if the face has a chord of this component on its boundary. There are no other edges. The venation graph of $\text{VEIN}(G)$ is denoted by $\text{VENATION}(G)$.




Figure 3.15: A biconnected bipartite graph G , its veins, dual veins, venation graph and a valid orientation of the venation graph.

The vertices of the venation graph can be divided into five classes:

- The components without interior face, C_0 ,
- The components with precisely one interior face, C_1 ,
- The components with precisely two interior faces, C_2 ,
- The components with more than two interior faces, C_3 , and,
- The connecting faces, F .

A similar counting argument as in the proof of Theorem 3.7 shows that for biconnected outerplanar graphs that admit a cTTR, C_3 must be empty.

Lemma 3.13. *Let G be a biconnected outerplanar graph. If G has a cTTR then every component in $\text{VEIN}(G)$ has at most two interior faces.*

Proof. Suppose there exists a cTTR of the biconnected outerplanar graph G . Let H be the skeleton graph of the cTTR.

²This refers to the arrangement of the veins of a leaf.

Let c be a component of $\text{VEIN}(G)$ and G_c the subgraph of G consisting of all faces that are bounded by an edge of c . The number of interior faces of G_c , denoted by ϕ , is one more than the number of edges in c , i.e., $\phi = |Ec| + 1$. This follows from the fact that c is connected and G_c is biconnected. Let $V_{\geq 3}$ be the set of vertices of G_c that have degree at least three and $v_{\geq 3}$ the cardinality of $V_{\geq 3}$. The number of vertices of degree 2 is denoted by v_2 and the number of interior faces of c by q . Since c has q interior faces, there are q edges whose removal leave a tree, hence,

$$\phi = |Ec| + 1 = v_{\geq 3} - 1 + q + 1 = v_{\geq 3} + q.$$

On the other hand, each interior face of G_c corresponds to an interior vertex of H . Such a vertex may admit a flat angle. The number of flat angles needed, is bounded below by:

$$\sum_{v \in V_{\geq 3}} (\deg(v) - 3) = 2|EG_c| - 2v_2 - 3v_{\geq 3} = 2|EG_c| - 2|VG_c| - v_{\geq 3} = 2\phi - 2 - v_{\geq 3}.$$

In the third step we have used Euler's formula. Combining the two equations, we obtain:

$$\phi \geq 2\phi - 2 - v_{\geq 3} \quad \Rightarrow \quad v_{\geq 3} + q = \phi \leq v_{\geq 3} + 2$$

and, therefore, q is at most 2. So every component has at most two interior faces. This proves the lemma. \square

In the remainder of this chapter, we assume that $C_3 = \emptyset$, i.e., there is no component in $\text{VEIN}(G)$ that has more than two interior faces. An interior face of a biconnected outerplanar graph relates to an interior vertex in the auxiliary graph H . An interior vertex may admit a flat angle, hence, it may be assigned in a flat angle assignment. As every vertex is assigned at most once, a vertex in H admits a flat angle in at most one of the faces it is incident to. So a face in G and its representative in H , are used for the representation of at most one of the components of $\text{VEIN}(G)$. We say that an interior face f of G *belongs to* a component c of $\text{VEIN}(G)$, if the flat angle of the representative of f contributes to the representation of the subgraph G_c (see Figure 3.16).




Figure 3.16: The graph G from Figure 3.15 with its auxiliary graph is drawn in green, and a cTTR of G . The vertices of H that represent connecting faces are labeled a , b and c . The vertex a belongs to the leftmost component of $\text{VEIN}(G)$ as its assignment contributes to the representation of this component.

From Lemma 3.13 it follows that, if a component of $\text{VEIN}(G)$ has two interior faces, then all the faces incident to a chord of this component belong to this component.

Corollary 3.14. *If a component of $\text{VEIN}(G)$ has two interior faces, then in a cTTR, all the faces incident to a chord of this component belong to this component.*

Proof. From the proof of Lemma 3.13 we know that the number of flat angles needed is at least $2\phi - 2 - v_{\geq 3}$. Since $q = 2$, we have $\phi = v_{\geq 3} + 2$, and it follows that the number of flat angles needed is precisely the number of vertices available:

$$2\phi - 2 - v_{\geq 3} = 2\phi - 2 - \phi + 2 = \phi.$$

Therefore, all faces that are incident to a chord of this component, must belong to this component. \square

A similar argument shows that, of the faces that are incident to a component with one interior face, at most one can belong to another component. Of the faces that are incident to a component without interior faces, at most two can belong to another component. For the faces that are incident to more than one component, the property of ‘belonging to a component’, can be modeled as an orientation of $\text{VENATION}(G)$. The conditions on the number of faces that may belong to another component motivates the following definition.

Definition 3.15 (Valid Orientation of $\text{VENATION}(G)$). An orientation of the edges of $\text{VENATION}(G)$ is called *valid* (see Figure 3.15) if every edge is oriented and:

- Every vertex in C_2 has only incoming arcs.
- Every vertex in C_1 has at most one outgoing arc.
- Every vertex in C_0 has at most two outgoing arcs.
- Every vertex in F has at precisely one outgoing arc.

An edge in the valid orientation that is oriented $f \rightarrow c_1$, for $f \in F$ and $c_1 \in C_1$ implies that f belongs to c_1 .

Recall that every boundary edge of the auxiliary graph H whose contraction does not induce a 2-face is actually contracted. To be able to do all possible contractions with respect to the interior faces of H , we have to avoid the situation where the outer face becomes a 2-face, after the contractions (see Figure 3.17). Every degree-2 vertex has to be represented by a polygon that has 2 boundary vertices. Therefore, the situation described above only occurs, when there are at most two degree-2 vertices. The biconnected outerplanar graphs with at most two degree-2 vertices have a 3TTR, as we will show below. Afterwards, we will assume that G has at least three vertices of degree 2.

Lemma 3.16. *Let G be a biconnected outerplanar graph and H its auxiliary graph. If there is a boundary edge of H that is not contracted, only because, otherwise the outer face becomes a 2-face, then every component of $\text{VEIN}(G)$ is a tree, i.e., $C_1 = C_2 = C_3 = \emptyset$.*

Proof. Let H be the auxiliary graph of a biconnected outerplanar graph G , which has a boundary edge that is not contracted, because the outer face becomes a 2-face otherwise. Note that, G has at most two degree-2 vertices, otherwise a face that represents one of these vertices must be another reason to not contract another boundary edge. Moreover, every biconnected outerplanar graph on at least three vertices has at least two vertices of degree 2. Hence, G has precisely two vertices of degree 2. Suppose there is a component of $\text{VEIN}(G)$ that has an interior face f . Let e_1, e_2 and e_3 be three of the chords that bound f . Deleting the edges e_1, e_2 and e_3 and their endpoints, cuts the graph into at least three components. For each $i = 1, 2, 3$, let G_i be the graph that consists of the component that is connected to e_i , together with e_i and the corresponding edges, and an extra vertex x that is connected to the endpoints of e_i (see Figure 3.18). The vertex x represents ‘the rest of the graph’.

For each i , the graph G_i is biconnected, outerplanar and it has at least three vertices. It follows that it has at least two vertices of degree 2. At least one of the degree-2 vertices of each G_i , is a degree-2 vertex of G . As the graphs are disjoint, except possibly for the ends of e_i , these degree-2 vertices are distinct. Therefore, G has at least three degree-2 vertices, which contradicts the assumption. Therefore, there is no component of $\text{VEIN}(G)$ that has an interior face, so $C_1 = C_2 = C_3 = \emptyset$. \square




Figure 3.17: A graph G such that one edge of H cannot be contracted, only because, the outer face becomes a 2-face otherwise.






Figure 3.18: The graphs G_i that consist of the component that is connected to e_i , together with e_i and an extra vertex x .

If $\text{VEIN}(G)$ has no interior faces, then G has a chord-to-endpoint assignment in which all chords are assigned. Therefore, such a graph has a 3TTR which can be constructed using the construction of Fowler. For the remainder of this section we assume that G has at least one interior face and, therefore, at least three degree-2 vertices. From this assumption it follows that all boundary edges of H whose contraction does not induce an interior 2-face are contracted.

Theorem 3.17. *A biconnected outerplanar graph admits a cTTR if and only if*

- [K1] *Each component of $\text{VEIN}(G)$ has at most two interior faces.*
- [K2] *The graph $\text{VENATION}(G)$ admits a valid orientation.*

The harder part of the proof is to show that K1 and K2 are sufficient conditions. In order to prove this, we will construct assignments of the vertices that belong to a component, for each component. Then we show that using the condition of belonging to, we can construct these assignments such that they do not interfere (Lemma 3.18). Each of the assignments is obtained from a chord-to-endpoint assignment. We show that the union of the assignments is a cFAA of H (Lemma 3.19). The last step is to show that this cFAA is a good cFAA. This is done by showing that the chosen cFAA is (almost) precisely union of the cFAA's that come from Fowlers' method. Since these all induce a 3TTR, we can show that the union induces a cTTR (Lemma 3.20 and Lemma 3.21).

For a component c of $\text{VEIN}(G)$, let G_c be the subgraph of G , that consists of all faces that share an edge with c . Let V_c be the set of vertices in the auxiliary graph H which represent an interior face of G_c that does not belong to another component than c . The subgraph of H induced by V_c , $H[V_c]$, is denoted by the *dual vein* of c (see Figure 3.15 on page 76).

Lemma 3.18. *Let Σ be a contact family of pseudosegments induced by an assignment of the interior vertices of H which satisfies C_v , i.e., every vertex is assigned at most once. Then, there is no pseudosegment of Σ that has vertices of two different dual veins as interior points.*

Proof. First note that a vertex u in H , that represents a connecting face between two components of $\text{VEIN}(G)$, has at least one neighbor on the boundary of H between any two neighbors that belong to different components. From this it follows that u can only be an

interior point for a pseudosegment that starts in a boundary neighbor of u , goes through u into the dual vein to which u belongs (see Figure 3.19). As the boundary vertices are not assigned, a pseudosegment cannot have a boundary vertex of H as an interior point. From the two arguments it follows that if a pseudosegment contains vertices of two different dual veins, then one of the vertices is the connecting face between the two components represented, and this vertex is an endpoint of the pseudosegment. \square




Figure 3.19: A connecting face u , the area of B and C belong to the same component. In the auxiliary graph H , drawn in green, u has at least one neighbor on the boundary of H between any two neighbors that belong to different components. Boundary edges are drawn with dashed lines.

From Lemma 3.18, it follows that a cFAA of H , can be splitted into assignments of the vertices of each dual vein. To obtain an assignment for a dual vein, we will use the chord-to-endpoint assignment of the component of $\text{VEIN}(G)$. First, we show how to transform a chord-to-endpoint assignment into an assignment of flat angles. Later, we show how to ensure that only the vertices that belong to this component are assigned, i.e., for a face f , which belongs to a different component, the representing vertex v_f in H , is not assigned by this assignment.

Recall that a chord-to-endpoint assignment is an assignment of chords to endpoints, in such a way that at most one chord is not assigned, and, each endpoint has at most one chord assigned to it.

From chord-to-endpoint assignment to cFAA.

In order to obtain assignments of the vertices of the dual veins of G , we show how to obtain an assignment from a chord-to-endpoint assignment. For a component c of $\text{VEIN}(G)$, we have defined G_c to be the subgraph of G , that consists of all faces that share an edge with c . Let ζ be a chord-to-endpoint assignment of G_c . Let H_c be the weak dual of G_c together with a half-edge into the unbounded face for each boundary edge of G_c . The cFAA assigns the vertices of H_c to the half-faces, that is, the faces that would arise when the half-edges in the unbounded face would be cyclically connected (see Figure 3.20). As all possible contractions of the boundary edges of H are done, the half-faces are closed with one boundary vertex, unless the face represents a degree-2 vertex, then it is closed with two boundary vertices. Therefore, in the assignment there should be two not assigned vertices in every half-face that represents a vertex of degree at least 3 and one not assigned vertex in a half-face that represents a vertex of degree 2. We will now explain how to transform a chord-to-endpoint assignment into a cFAA.

Base Case. If all the chords are assigned then we start with one face of G_c . The vertex of H_c in this face is not assigned (Figure 3.21 (b)). If one chord e is not assigned, then we start with the two faces in G_c that are incident to e . Along the dual edge in H_c that represents e the two vertices of H_c are assigned to the opposite faces, any choice is fine (Figure 3.21 (a)).

Iteration. A step in the chord-to-endpoint assignment considers the following objects, the already selected part of the graph Δ , the new face f , cut off by the selected chord uv




Figure 3.20: The graphs G_c for all components c of $\text{vein}(G)$ of the graph in G , depicted on the left. The weak duals H_c are colored green (solid) and the half-edges in the unbounded faces are cyclically connected by the dashed lines.

from Δ , and, the endpoint u to which e is assigned in the chord-to-endpoint assignment. The rule is to assign the vertex of H_c that represents f , to the face of H_c that represents the vertex v of uv , i.e., the vertex to which uv is not assigned in the chord-to-endpoint assignment (Figure 3.21 (c)).

Special case. Up until now, the cFAA is precisely the cFAA one would get from the 3TTR using Fowlers' method, this will be shown in Lemma 3.20. However, if the component has no interior faces then there is one assignment too many. The problem arises since in Fowlers' method not all boundary edges that can be contracted are contracted, hence, the size of the face is different. This occurs in the face of H_c , that represents the endpoint of a chord that has no chord assigned to it. In this face an arbitrary assignment should be removed (Figure 3.21 (d)).




Figure 3.21: From chord-to-endpoint assignment to cFAA: The base case with two faces (a), the base case with one face (b), introducing a new face (c), and an example of the special case, the starting face is colored grey and the face that represents vertex a has one assignment too many (d).

Lemma 3.19. *Let G be a biconnected outerplanar graph and H its auxiliary graph. Suppose $\text{VENATION}(G)$ has a valid orientation. Then, there is a cFAA of H , that is the union of assignments which come from chord-to-endpoint assignments of the components of G .*

Proof. For a component c of $\text{vein}(G)$, V_c is the set of vertices in H that represent a face f of G such that f is bounded by an edge of c and f does not belong to a component other than c . For each component c , we will construct an assignment such that only the vertices in V_c are assigned. All the edges on the boundary of H that can be contracted without introducing a 2-face, are contracted.

Suppose c has two interior faces, then V_c consists of the representatives of all the faces in G_c . This holds since in a valid orientation the vertex that represents c has no outgoing arcs. Construct a chord-to-endpoint assignment of c , leaving out a chord that separates the two interior faces. The assignment is obtained according to the method described above.

Suppose c has one interior face, then $|V_c|$ is either equal to the number of chords in c , or equal to the number of chords plus one. This holds since in a valid orientation c has at most one outgoing arc. Construct a chord-to-endpoint assignment. In the first case, the face with the outgoing arc in the valid orientation is chosen as starting face. This ensures that the vertex that represents this face is not assigned by this assignment. In the latter case, any face can be chosen as starting face, the corresponding vertex will not be assigned.

Suppose c has no interior face. There are two special faces, we will identify them first. If c has two outgoing edges in the valid orientation then f_1 and f_2 are the two connecting faces. If c has one outgoing edge then f_1 is this connecting face and f_2 is another face that has a degree 2 vertex. If c has no outgoing edges then f_1 is one of the connecting faces³ and f_2 is another face that has a degree 2 vertex.

Set f_2 as the starting face. Construct a chord-to-endpoint assignment such that, the chord of f_1 , whose removal disconnects f_1 and f_2 , has an endpoint to which no chord is assigned. This ensures that the vertex that represents f_1 and the vertex that represents f_2 are the vertices that are not assigned by this assignment. Construct the assignment and remove the assignment of the vertex that represents f_1 .

We claim that the union of these assignments is a cFAA of H . Recall that an assignment is a cFAA if:

[C_v] Every vertex of U is assigned to at most one face, and,

[C_f] For every interior face f , precisely $|f| - 3$ vertices are assigned to f .

No boundary vertex of H is assigned. An interior vertex v of H belongs to V_c for precisely one component c . Every vertex in V_c is assigned at most once, while processing the component c . Therefore, C_v holds.

Let Δ be an interior face of H that represents vertex w of G . Let p be the number of chords incident to w . Suppose Δ includes precisely one boundary vertex of H . Then Δ is of size $p + 2$. Precisely $p - 1$ chords are not assigned to w in a chord-to-endpoint assignment, as only one chord is assigned to w . It follows, that in the flat angle assignment that comes from the chord-to-endpoint assignments including w , $p - 1$ vertices are assigned to Δ .

Suppose Δ includes two boundary vertices then w (the vertex represented by Δ) must be a degree-2 vertex. Moreover, the edges incident to w are not chords. Therefore, there is no assignment in Δ and Δ is a 3-face. We conclude that C_f must hold and the assignment is a cFAA of H . \square

In order to use that this construction of flat angle assignments will turn out to be a good flat angle assignment, we will first show that the local assignments are good. For a biconnected outerplanar graph G , and a component c of $\text{VEIN}(G)$, we have defined G_c to be the subgraph of G , that consists of all faces that share an edge with c . To show that the cFAA obtained from a chord-to-endpoint assignment is locally good, we consider the auxiliary graph of G_c , denoted by \overline{H}_c . We assume that there is at least one face that is a connecting face. This face has at least two degree-2 vertices, and there must be at least one more degree-2 vertex. Therefore, G_c at least three degree-2 vertices.

Lemma 3.20. *Let ψ be a flat angle assignment of \overline{H}_c , obtained from a chord-to-endpoint assignment ζ of c . Then every simple outline cycle has at least three combinatorial convex corners.*

³If there is no connecting face then c is the only component and the result of Fowler shows that there is a 3TTR, therefore, we may assume that c is not the only component

Proof. First we will show that ψ is exactly the assignment that appears in the representation constructed with the method of Fowler in almost all cases. Then we discuss the special case. When ψ is exactly the assignment that appears in some representation, then it must be a good FAA and, hence, every simple outline cycle has at least three combinatorial convex corners.

To show that the two assignments are equivalent we consider the base case and the iteration steps. By ψ we denote the assignment obtained by the algorithm described above and by ϕ_{Fowler} we denote the assignment extracted from a representation that is obtained by the method of Fowler.

In the base case there are two options. When c has at most one interior face, the vertex in \overline{H}_c that represents the starting face, is not assigned in ϕ_{Fowler} nor in ψ . When c has two interior faces, the two vertices that are endpoints of the starting chord, are 4-faces in the representation of Fowler. The two different assignments, yield the two different drawings in Figure 3.22 (a). In the method of Fowler the appropriate one can be chosen, such that the for these vertices the assignment in ϕ_{Fowler} is equal to the assignment in ψ .

Recall that in a step, one face is added and all its vertices are introduced in the triangle that represents the endpoint to which the connecting chord is assigned (see Figure 3.22 (b)). Let u and v be the endpoints of the chord, and let the chord be assigned to u . The size of the polygon that represents v , increases by one. This new vertex is placed on a straight-line segment, hence, the new point is assigned to the polygon that represents v . This is exactly how the assignment is introduced in the algorithm with which ψ is obtained. It follows that, after a step, ϕ_{Fowler} and ψ are still the same.

Suppose every vertex of degree at least 3 has a chord assigned to it. Then in the method of Fowler, the triangles that have two boundary vertices in the representation, are precisely the degree 2 vertices. Hence, \overline{H}_c is the skeleton of the representation, and the assignment extracted from the chord-to-endpoint assignment must be good (it can even be stretched to a 3TTR!).




Figure 3.22: The two base cases with two faces (a). The extension step of Fowler implies the same assignment as is extracted from the chord-to-endpoint assignment (b).

We proceed with the special case, in which there is one assignment removed. There is one vertex of degree at least three that has no chord assigned to it, this vertex is denoted by v and the polygon that represents v by Δ_v . In this case, c has no interior faces and, in the end, one assignment in Δ_v is removed of ψ . In the method of Fowler, the polygon that represents v has two vertices on the boundary of the representation. In \overline{H}_c , it has only one boundary vertex and therefore, in ψ , one assignment is removed in Δ_v . We will use the method of Fowler, with a slight change, to obtain a 3TTR of which \overline{H}_c is the skeleton.

Recall that we have removed the assignment of the representative of f_1 , the face f_1 is a connecting face which possibly belongs to another component. Moreover, the starting face is known, it is denoted by f_2 and the chord-to-endpoint is such that v is incident to the chord of f_1 that disconnects f_1 and f_2 .

It follows that in the method of Fowler, f_1 is introduced after v . The second valuable property is that f_1 has size at least 4. We change the introduction of f_1 in the method of Fowler. The invariant that is maintained along the construction, says that the triangle of a vertex that is not processed, shares an edge with the boundary. To obtain a representation which has $\overline{H_c}$ as a skeleton, we want to pretend that v is processed. If v would be processed, Δ_v would only share one vertex with the boundary, precisely as in $\overline{H_c}$. Let uv be the chord over which f_1 is introduced, uv is assigned to u . As $|f_1| \geq 4$, when f_1 is introduced we can choose to put the neighbor of v into Δ_v and all the other vertices into Δ_u , the representative of u . The procedure is depicted in Figure 3.23.




Figure 3.23: Vertex l_1 is introduced in v instead of in u .

The rest of the construction is as in the method of Fowler. By pretending that v is processed we have obtained a representation of which $\overline{H_c}$ is the skeleton. With this special construction step, the property that ψ and ϕ_{Fowler} are equal is maintained. We conclude that the assignment that comes from a chord-to-endpoint assignment, induces a 3TTR and therefore, such an assignment is locally good. Hence, every simple outline cycle in $\overline{H_c}$ has at least three combinatorial convex corners under ψ . \square

We proceed by showing that the cFAA, which is the union of the assignments that come from chord-to-endpoint assignments, is stretchable to a cTTR. We call a straight-line segment in a cTTR that connects two boundary vertices a *diagonal*. In a cFAA we can also speak about a diagonal, this is a path which becomes a diagonal in a cTTR that belongs to this cFAA.

Lemma 3.21. *Let G be a biconnected outerplanar graph and H its auxiliary graph. Suppose $\text{VENATION}(G)$ has a valid orientation. Let ψ be a cFAA of H , that comes from chord-to-endpoint assignments of the components of $\text{VEIN}(G)$. Then, every simple outline cycle in H has at least three combinatorial convex corners.*

Proof. Let γ be a simple outline cycle in H . Recall that $\text{int}(\gamma)$ is the outline cycle γ together with all vertices, edges and faces that are enclosed by γ .

Suppose that $\text{int}(\gamma)$ contains vertices of at most one dual vein, say of $H[V_c]$. Note that the vertices of $\text{int}(\gamma)$ are not necessarily all included in V_c , as the boundary vertices of H do not belong to dual veins. Since the boundary vertices are not assigned, it follows from Lemma 3.20 that γ has at least three combinatorial convex corners.

Suppose there is a simple outline cycle γ that has at most two combinatorial convex corners. By the above, $\text{int}(\gamma)$ must contain vertices of at least two dual veins. But then γ can be splitted into smaller cycles, each belonging to one dual vein (see Figure 3.24). Each of these has at least three combinatorial convex corners. Two smaller cycles are glued together at two points. If the path between the two points is an edge (e.g., p_1, p_2 in Figure 3.24), then

one of p_1, p_2 must be a boundary point and, therefore, a combinatorial convex corner for γ in H . Furthermore, both smaller cycles contribute at least one more corner to their union. If the path between the two gluing points consists of two edges (e.g., p_1, p_3 in Figure 3.24), then both p_1 and p_3 must be boundary vertices and, therefore, combinatorial convex corners for γ in H .




Figure 3.24: Splitting an outline cycle into outline cycles each of which belongs to one dual vein.

Suppose the two gluing points are the only corners after gluing. Without loss of generality, we may assume that on one side of the “gluing path” we add (part of) only one dual vein, as we consecutively glue on parts. Let X be the subgraph of H consisting of this part. By assumption, X and the part on the other side of the gluing path, both have only three combinatorial convex corners and they are all on the gluing path. The middle vertex of the gluing path (p_2 in Figure 3.24) is a combinatorial convex corner for both of the cycles that we are gluing together. Therefore, this vertex is not assigned by either of its neighboring components. From this it follows that both these components have at most one interior face in their vein (they are C_0 or C_1 components). Hence, X belongs to a C_0 or a C_1 component.

The cycle bounding X has no combinatorial convex corners on the path between p_1 and p_3 that does not go through p_2 (see Figure 3.25 (a)). Suppose p_1 and p_3 are neighbors (see Figure 3.25 (b)). Then we consider another gluing sequence, starting with X . Note that we cannot end up in the same situation since then H has only two boundary vertices (see Figure 3.25 (c)).




Figure 3.25: Gluing over the path p_1, p_2 and p_3 : there must be a combinatorial convex corner on at least one side, or the original graph has only two degree-2 vertices.

Hence, the path between p_1 and p_3 has at least one vertex and all the vertices are assigned inside X or outside X with the restriction that this vertex does not have a neighbor outside X . We finish the argument by proving the following three claims, which together yield that there cannot be such a path without a combinatorial convex corner for X .

Claim 1. In a cFAA of a C_0 or a C_1 component which comes from a chord-to-endpoint assignment there is no diagonal.

This immediately follows from the result of Fowler. For a C_1 component the obtained cFAA belongs to a cTTR which interiorly is precisely the 3TTR as obtained by Fowlers' method. For a C_0 component the cFAA belongs to a cTTR which is a 3TTR obtained by Fowlers' method with possibly one fake-processed vertex. The 3TTR obtained by Fowlers' method does not have a diagonal unless there is a diagonal in the base case. The base case has a diagonal only if the component is of type C_2 . \triangle

Claim 2. The path between p_1 and p_3 on the boundary of X that does not go through p_2 , cannot have all its interior vertices assigned inside X .

Suppose all vertices on this path are assigned inside X . We count the number of assignments needed and compare it to the number of assigned vertices if all vertices on this path are assigned inside X . We consider a graph Z as in Figure 3.26 (a), which consists of X together with a vertex z which is connected to p_1, p_2 and p_3 .




Figure 3.26: The graph Z consists of X and an extra vertex z which is connected to p_1, p_2 and p_3 . A vertex y on the boundary of X such that y has no neighbors outside of X .

Let v_Z be the number of interior vertices of Z , let $b \geq 4$ be the number of boundary vertices, e_Z the number of edges and F_Z the set of interior faces. The underlying outerplanar graph is denoted by G . The number of vertices of degree at least 3 in G is denoted by $v_{\geq 3}$ and the number of degree-2 vertices in G is denoted by $v_2 = b$. Suppose p_2 is the only not assigned vertex interior to Z . If all vertices on the boundary, except for p_1, z and p_3 , are assigned to a face inside Z then we need $v_Z - 1 + b - 3$ assigned vertices. On the other hand, the number of assignments is given by:

$$\sum_{f \in F_Z} (|f| - 3) = 2e_Z - b - 3|F_Z| = 2(v_Z + b) - 2 - b - |F_Z| = 2v_Z + b - 2 - v_{\geq 3} - v_2 = 2v_Z - v_{\geq 3} - 2.$$

Recall that the interior vertices of Z each represent a face of G . If Z belongs to a C_0 component then $v_Z = v_{\geq 3}$. If Z belongs to a C_1 component then $v_Z = v_{\geq 3} + 1$. From identifying the two expressions for the number of assigned vertices we obtain:

$$2v_Z - v_{\geq 3} - 2 = v_Z - 1 + b - 3 \Rightarrow b = v_Z - v_{\geq 3} + 2.$$

This implies that $b = 2$ if Z belongs to a C_0 component and $b = 3$ if Z belongs to a C_1 component. Both contradict $b \geq 4$. Note that we started with the assumption that p_2 is the only vertex that is not assigned. When X belongs to a C_0 component there may be another such vertex. In this case we obtain:

$$2v_Z - v_{\geq 3} - 2 = v_Z - 2 + b - 3 \Rightarrow b = v_Z - v_{\geq 3} + 3$$

and this implies that $b = 3$. We again obtain a contradiction to $b \geq 4$ and conclude that the claim must hold. \triangle

Claim 3. The path between p_1 and p_3 on the boundary of X that does not go through p_2 has at least one combinatorial convex corner of X in its interior.

First we will show that if X contains the starting face then a vertex on the boundary of X that has no neighbors outside X cannot be assigned outside of X . This proves the claim when p_2 represents the starting face.

Consider the vertex y as in Figure 3.26 (b). In the method of Fowler every face is introduced by gluing it to the current graph on some chord. Since X contains the starting face, the face that y represents must have been introduced over one of the solid edges in Figure 3.26 (b). It cannot be introduced over the dashed edges since one of the vertices of these edges lies strictly outside X . The chord over which a face is introduced in Fowler's method ensures that the vertex that represents this face is assigned in the representation of the endpoint to which this chord is assigned. From this it follows that y must be assigned to the angle labeled 1, 2, or 3, or it is not assigned.

If p_2 is not the starting face, then X belongs to a C_0 component and there are two vertices not assigned. The second vertex that is not assigned must be outside of X as otherwise, the starting face belongs to X and the previous argument suffices. From Claim 2 we know that there must be some vertices on the boundary of X that are assigned outside X but have no neighbors outside of X . Suppose there are k such vertices. Again using the expressions obtained before we get:

$$2v_z - v_{\geq 3} - 2 = v_z - 1 + b - 3 - k \quad \Rightarrow \quad b = 2 + k.$$

But p_1, z and p_3 count for b and not for k , therefore, $b \geq 3 + k$, contradiction.

It follows that not all internal vertices of the path between p_1 and p_3 can be non-corners, hence, at least one of them must be a combinatorial convex corner for X . \triangle

Therefore, there is no path that connects p_1 and p_3 that is assigned as in Figure 3.25 (a).

Note that, since the cycles belong to different dual veins, there cannot be a strictly interior edge in the gluing path. Iteratively gluing together the smaller cycles results in the original outline cycle γ . In each step the glued cycle has at least three combinatorial convex corners and, therefore, γ has at least three combinatorial convex corners.

Hence, every simple outline cycle of H has at least three combinatorial convex corners. \square

With Lemma 3.21 it has become easy to show that the conditions K1 and K2 of Theorem 3.17 are sufficient.

Theorem 3.17. *A biconnected outerplanar graph admits a cTTR if and only if*

[K1] *Each component of $\text{VEIN}(G)$ has at most two interior faces.*

[K2] *The graph $\text{VENATION}(G)$ admits a valid orientation.*

Proof. “ \Leftarrow ”: Let G be a biconnected outerplanar graph. Suppose that $\text{VEIN}G$ has no interior face, then G has a 3TTR which can be transformed into a cTTR (by slightly moving the vertices that have an angle of size π in the outer face). Therefore, we assume that $\text{VEIN}G$ has at least one interior face and at least three degree two vertices.

The graph H is obtained by taking the weak dual of G , adding an edge into the outer face for each boundary edge of G . The new vertices are cyclically connected and every boundary edge of H , whose contraction does not induce a 2-face, is contracted. From Lemma 3.19,

and conditions K1 and K2, it follows, that H has a cFAA. From Lemma 3.21, it follows that the cFAA that comes from the chord-to-endpoint assignments induces a cTTR.

“ \Rightarrow ”: Suppose there exists a cTTR of a biconnected outerplanar graph G . Let H be the underlying graph of the cTTR. From Lemma 3.13 it follows that K1 holds, for convenience we repeat the proof here. Let c be a component of $\text{VEIN}(G)$ and G_c the subgraph of G consisting of all faces that are bounded by an edge of c . The number of interior faces of G_c , denoted by ϕ_c , is equal to the number of chords, $|Ec|$, plus one, this follows from the fact that c is connected and G_c is biconnected. Let $V_{\geq 3}$ be the set of vertices of G_c that have degree at least three, $v_{\geq 3}$ the cardinality of $V_{\geq 3}$. The number of vertices of degree 2 is denoted by v_2 and q the number of interior faces of c . First we express the number of interior faces of G_c (which is the number of possible flat angles in H_c) in $v_{\geq 3}$ and q .

$$\phi_c = |Ec| + 1 = v_{\geq 3} - 1 + q + 1 = v_{\geq 3} + q$$

On the other hand, the number of flat angles needed is bounded below by (in the third step Euler’s formula is used):

$$\sum_{v \in V_{\geq 3}} (\deg(v) - 3) = 2|EG_c| - 2v_2 - 3v_{\geq 3} = 2|EG_c| - 2v - v_{\geq 3} = 2\phi_c - 2 - v_{\geq 3}.$$

Hence, we obtain

$$\phi_c \geq 2\phi_c - 2 - v_{\geq 3} \quad \Rightarrow \quad \phi_c \leq v_{\geq 3} + 2$$

and therefore q is at most 2, i.e., c has at most two interior faces and K1 holds.

For each component c of $\text{VEIN}(G)$, let V_c be the set of vertices of H that are assigned and belong to c . In other words, a vertex in V_c is assigned to a face Δ and Δ represents a vertex of c . It is obvious that $|V_c| \leq \phi_c$. Using the relations obtained before, it follows that the following implications must hold.

$$q = 2 \quad \Rightarrow \quad |V_c| \geq 2\phi_c - 2 - v_{\geq 3} = v_{\geq 3} + 2 = \phi_c \quad (3.1)$$

$$q = 1 \quad \Rightarrow \quad |V_c| \geq 2\phi_c - 2 - v_{\geq 3} = v_{\geq 3} = \phi_c - 1 \quad (3.2)$$

$$q = 0 \quad \Rightarrow \quad |V_c| \geq 2\phi_c - 2 - v_{\geq 3} = v_{\geq 3} - 2 = \phi_c - 2 \quad (3.3)$$

An orientation of $\text{VENATION}(G)$ can be obtained from the assignments of the vertices that represent a face that connects two or more components. If a vertex v that represents such a face belongs to component c , the edge in $\text{VENATION}(G)$ is oriented from v to c . As v is assigned at most once, the other edges of v are incoming. If a vertex is not assigned, then the outgoing arc is chosen arbitrarily.

Every edge incident to a component with two interior faces, is oriented towards this component, due to 3.1. At most one edge incident to a component with one interior face, is outgoing due to 3.2, and at most two edges of component without interior faces, are oriented outwards due to 3.3. As all the edges are oriented and each vertex in F has at precisely one outgoing arc, we have obtained a valid orientation, as desired by K2. \square

Unfortunately, this does not classify the biconnected outerplanar graphs that admit a 3TTR. The question “given an integer k , can a particular cTTR be transformed into a k TTR?” remains open. The transformation consists of selecting the k outer face vertices and assigning all other vertices to the outer face. However, this transformation is not always possible (see Figure 3.27).




Figure 3.27: Selecting four vertices to transform the cTTR into a 4TTR: the four white vertices in (a) are not a good choice since the grey area will be forced onto a line-segment (b) and a 4TTR of the same graph using different vertices as corners (c).

Consider the contact family of pseudosegments which arises from the cFAA that is stretchable. Recall that a diagonal is a straight-line segment that connects two boundary vertices in a cTTR. If there is no diagonal, then every tuple of three vertices can be chosen as corners of the outer face. However, between two boundary vertices that are the endpoints of a diagonal, there must be a suspension (on both sides). In a cFAA that belongs to an outerplanar graph a diagonal occurs (unavoidably) in a component of $\text{VEIN}(G)$ that has two interior faces.

Lemma 3.22. *In any representation of a component of $\text{VEIN}(G)$ that has two interior faces there is a diagonal.*

Proof. An interior vertex of the dual vein of degree k is an endpoint for $k - 2$ interior pseudosegments. Since every interior vertex is assigned the vertex will also be an interior point for one pseudosegment. Let V be the set of vertices of the dual vein and e the number of edges of the dual vein, these are the strictly interior vertices and edges in the representation. Let G_c be the subgraph of G consisting of all the faces that are bounded by an edge of c . Let x be the number of vertices in G_c , which is equal to the number of edges in H that have one endpoint on the boundary due to outerplanarity. Note that $e + 1 = |V|$ since the dual vein is a tree. In total the number of interior endpoints is given by

$$\sum_{v \in V} (\deg(v) - 2) = 2e + x - 2|V| = e + x - |V| - 1 + (e + 1 - |V|) = e + x - |V| - 1.$$

The number of interior pseudosegments is $e + x - |V|$, i.e., the number of interior edges minus the number of interior vertices since every vertex is assigned. There are more interior pseudosegments than there are interior endpoints, therefore, there must be at least one pseudosegment with two ends on the boundary, this is a diagonal. \square

An obvious upper bound on the number of corners needed on the outer face is twice the number of components in C_2 , if $|C_2| > 1$. However, when there are many components in C_2 all of which are “parallel”, this bound is far from optimal (see Figure 3.28). Let $c, d \in C_2$ be two components and the endpoints of the diagonals in the representation of these components are denoted by x_c, y_c and x_d, y_d , respectively. The components c and d are *parallel* if there is a k TTR in which x_c, x_d and y_c, y_d pairwise belong to the same boundary segment. The three components of C_2 in the graph in Figure 3.28 all are pairwise parallel.

To obtain a characterization the property of being parallel should be changed into a combinatorial description that fits the original graph G . Intuitively, two components are parallel




Figure 3.28: A 3TTR of a graph with three components with two interior faces, the diagonals are drawn in red, and they all end on the same two boundary segments.

if the boundary vertices of G can be covered by two paths on the boundary in such a way that for each of the components the interior faces have vertices of different color.

Definition 3.23 (Boundary Path Coloring). Let R be a cTTR and $k \geq 3$ an integer. Let $\mathcal{P} = P_1 \cup P_2 \cup \dots \cup P_k$ a covering of the boundary vertices of the cTTR such that:

- For $i = 1, \dots, k$: P_i is a path.
- For $i = 1, \dots, k - 1$: $|P_i \cap P_{i+1}| = 1$ and $|P_k \cap P_1| = 1$.

\mathcal{P} is a boundary path coloring of R if:

- [D1] For every diagonal s that ends in two boundary vertices x, y of R , there does not exist an i such that $x \in P_i$ and $y \in P_i$.

We call k the number of colors of the boundary path coloring.

Theorem 3.24. Let $k > 2$ be an integer. A biconnected outerplanar graph admits a k TTR if and only if it has a cTTR that admits a boundary path coloring \mathcal{P} with at most k colors.

Remark. The cTTR considered is a cTTR as obtained by the method described before; we rely on the properties of such a cTTR. We assume that there are at least two components in $\text{VEIN}(G)$ that have two interior faces. Otherwise there is only one diagonal in a cTTR obtained by the method described above and it follows that the graph has a 3TTR.

Proof. Surely, if a biconnected outerplanar graph admits a k TTR then this can be transformed into a cTTR with a boundary path coloring with at most k colors. The k boundary segments of the k TTR are the paths of the boundary path coloring. By perturbing the vertices that are assigned to the outer face in the k TTR slightly into the outer face, we obtain a cTTR.

To prove the converse we have to show that the boundary path coloring \mathcal{P} induces an assignment of the vertices in the outer face which together with a cTTR yields a k TTR. First we show that there is a cTTR R for which \mathcal{P} is a boundary path coloring and for which the following property holds. This property is important, since D1 only ensures that diagonals do not induce a flattened part. The following property ensures that also “concave” paths connecting two boundary vertices do not induce a flattened part (see Figure 3.29). Such a “concave” path arises when the valid orientation is chosen such that a part with only C_0 and C_1 components has the next connecting face belonging to one of its components. Changing the orientation such that this face does not belong to this part resolves the problem.




Figure 3.29: Part of a biconnected outerplanar graph with only components of C_0 and C_1 (black), together with its auxiliary graph (green). A chord-to-endpoint assignment is given by the black arrows on the chords and the assignment with red arrows. The valid orientation of the venation graph is given on the right. The vertex w is assigned. Suppose the whole boundary of this part belongs to one color class, then, due to the two edges from w to the boundary and the assignment of w , this part would be flattened onto its boundary. Therefore, we want to change the assignment such that w is not assigned.

Property 3.25. Let r be a path that connects two boundary vertices in R , r is not a straight-line segment, and, r contains precisely two boundary vertices. Let R_1 and R_2 be the two parts of R that are separated by r . Then the following properties must hold.

- There is a convex corner of $R_1 \cup r$ interior to r or $R_1 \cup r$ has (part of) a diagonal inside.
- There is a convex corner of $R_2 \cup r$ interior to r or $R_2 \cup r$ has (part of) a diagonal inside.

Suppose a path r has only one interior vertex, denoted by w . Suppose w represents a connecting face of $\text{VEIN}(G)$. If w is not assigned then both sides have a convex corner that is an interior vertex of r and the properties hold. Suppose w is assigned and let R_1 be the side to which w is assigned. Suppose there is no diagonal in R_1 . Then there are only components of C_0 and C_1 in R_1 . Consider this part of the venation graph as a rooted tree, rooted in the face that is represented by w . We change the valid orientation such that every element of the tree has an outgoing edge to its parent. This is an orientation in which w has no outgoing edge. We orient an edge of w towards a component of C_2 , if there is no such component then the graph has a 3TTR by Fowlers theorem. This again is a valid orientation and in the new cTTR the property holds for this path.

Suppose that r has more than one interior vertex and at least one of them represents a connecting face, then the same reasoning works. Suppose that none of the interior vertices of w is a connecting face. But then r belongs to the interior of one component and the property must be satisfied.

By changing the orientation of the venation graph only assignments of connecting faces are changed. A vertex that represents a connecting face may only be interior to a diagonal if it belongs to a C_2 component. In such a situation, the assignment is not changed as the property is already satisfied. It follows that the boundary path coloring is still valid.

Let ψ be the assignment of the (new) cTTR together with the assignment of the interior vertices of the paths in the boundary path coloring to the outer face. Suppose that ψ does not belong to a k TTR. Then there must be a simple outline cycle, which has at most two combinatorial convex corners. This implies that all the combinatorial convex corners of this

simple outline cycle in the cTTR must be on the boundary. Then, the boundary of this outline cycle that is not on the boundary of the cTTR, must be a straight-line segment, as it cannot be concave due to Property 3.25. But then, its endpoints belong to strictly different sets P_i and P_j . It follows, that the endpoints are combinatorial convex corners, as well as the vertex that belongs to P_i and $P_{i\pm 1}$. Hence, every simple outline cycle has at least three combinatorial convex corners and the graph has a k TTR. \square

We believe, that a boundary path coloring of a cTTR, can be transformed into a boundary path coloring of an auxiliary graph H , and then into a boundary path coloring of the original graph G . The property D1, which holds for a boundary path coloring of a cTTR is transformed to a condition of H . For every component $c \in C_2$, let W_c be the set of pairs of boundary vertices of H such that the diagonal in a representation of this component must end in one of the pairs. In other words, W_c is the set of all possible endpoints of the diagonal of c .

[D2] For every $c \in C_2$, there must be a pair $(x, y) \in W_c$ such that there does not exist an i such that both $x \in P_i$ and $y \in P_i$.

To show that this condition is still sufficient, one needs to show that there exists a cTTR in which all the diagonals end in prescribed pairs of W . We believe that this can be done using a particular choice of starting faces for each component of C_2 . Then D2 can be translated into a property of a boundary path coloring of G . For every component $c \in C_2$, let V_c^1 and V_c^2 be two sets of degree-2 vertices. A vertex v is an element of V_c^1 if the interior face of G that is represented by v , is bounded by a chord of the first interior face of c . Similarly, a vertex v is an element of V_c^2 if the interior face of G that is represented by v , is bounded by a chord of the second interior face of c .

[D3] For every $c \in C_2$, there must exist a pair (x, y) , $x \in V_c^1$ and $y \in V_c^2$, such that there does not exist an i for which both $x \in P_i$ and $y \in P_i$.

To show that this condition is again sufficient, one needs to take care of the difference between the boundary vertices of H and the boundary vertices of G . The degree-2 vertices of G relate to two boundary vertices of H . Therefore, we believe that the condition on G implies the condition on H . Secondly, one needs to show that this condition is not stronger than the condition on H . We also believe that it can be shown that if a graph has a k TTR then the condition on G is satisfied. We conjecture the following. An example with a boundary path coloring in G is given on page 94.

Conjecture 3.26. *A biconnected outerplanar graph admits a k TTR if and only if the following three statements hold.*

[K1] *Each component of $\text{VEIN}(G)$ has at most two interior faces.*

[K2] *The graph $\text{VENATION}(G)$ admits a valid orientation.*

[K3] *There is a boundary path coloring of G that satisfies D3.*

3.3 Conclusion

We have shown that Halin graphs admit a 3TTR and we have characterized the biconnected outerplanar graphs that admit a cTTR. The conclusion of this chapter is that there is still a lot of work to do in this area.

For example, can the characterization of biconnected outerplanar graphs be used to characterize:

- Biconnected outerplanar graphs that admit a k TTR?
- Biconnected internally cubic planar graphs that admit a cTTR?
- 2-outerplanar graphs that admit a cTTR?

And for which other graph classes can we find a characterization?

What happens if triangles and quadrangles are allowed?

Which of the representations are area-universal, i.e., realizable with a prescribed size for each triangle.




Figure 3.30: Top-left: A graph G in black, with its chords highlighted in green and a boundary path 3-coloring. Top-right: A valid orientation of $\text{VENATION}(G)$. Center: G and chord-to-endpoint assignments in green, the graph H where the contracted edges are dashed, in black, and an FAA given by the red arrows that comes from the chord-to-endpoint assignments. Bottom: A 3TTR of G .

"It made her think that it was curious how much nicer a person looked when he smiled. She had not thought of it before."

Frances Hodgson Burnett, *The Secret Garden*

4

Grid-Paths Contact Representations

In this chapter we study planar graphs that admit a VCPG, i.e. graphs admitting a contact representation of paths on a grid. The abbreviation VCPG comes from Vertex Contact representation of Paths on a Grid and is inspired by VPG (Vertex intersection graph of Paths on a Grid) and EPG (Edge intersection graph of Paths on a Grid), which were introduced by Asinowski et al. [ACG⁺12] and Golumbic et al. [GLS09], respectively. In such a representation the vertices of G are represented by a family of internally disjoint grid-paths. Adjacencies are represented by contacts between an *endpoint* of one grid-path and an *interior point* of another grid-path. When the number of bends of each path is at most k , we denote the representation by B_k -VCPG and when every path has precisely k bends, we speak about *strict* B_k -VCPG.

The graphs that can be represented by grid-paths without bends are segment graphs where the segments are placed in only two directions. Intersection graphs of segments in two directions are also known as *Grid Intersection Graphs* (GIG). This class is well-studied. It is clear that no such graph can have an odd cycle, therefore this is a subclass of the class of bipartite graphs. The restriction to contact representations requires the graphs to be planar. It has been shown that B_0 -VCPG (or contact 2DIR) graphs are precisely bipartite planar graphs [HNZ91]. An algorithm to construct such a drawing for the vertex-edge incidence graph of a planar graph is given in [RT86], an algorithm for bipartite graphs is given in Chapter 1 (page 6 and further). It has been shown that the 2-orientations of a maximal bipartite planar graph are in bijection with the separating decompositions of this graph (e.g. [dFdM01]). In turn a separating decomposition induces a segment contact representation by segments in two directions (cf. [HNZ91, dFdMP95, Fel13]). A simple algorithmic proof of existence of a segment contact representation of a bipartite graph is given by Czyzowicz, Kranakis and Urrutia [CKU98].

Intersection graphs of grid-paths are denoted by VPG graphs. VPG graphs were introduced by Asinowski et al. [ACG⁺12]. From the result that every planar graph has a contact representation by T-shapes ([dFdMR94], see page 66), it follows that every planar graph has an intersection representation by grid-paths where each path has at most three bends [ACG⁺12]. To obtain the B_3 -VPG every T-shape is replaced by a grid-path that follows the T-shape in such a way that every point is covered on at least one side (see Figure 4.1). Asinowski et al. conjectured that this bound was tight, i.e., there exists a planar graph for which three bends are necessary. Chaplick and Ueckerdt disproved this by showing that every planar graph admits a B_2 -VPG [CU13]. To the best of our knowledge, it is still open to decide whether this is tight. It could be that every planar graph admits a B_1 -VPG.




Figure 4.1: A contact representation by T's and a B_3 -VPG obtained from it.

In a contact representation by grid-paths (as well as segments, pseudosegments, etc.) the vertices are represented by a family of internally disjoint objects. An *endpoint* of one object coincides with an *interior point* of another object if and only if the two represented vertices are adjacent. In an intersection representation obtained from the contact representation

by T-shapes (see Figure 4.1), there are grid-paths that cross twice. To represent this as a contact, the two paths would touch at an interior point of both paths, which is not allowed in our definition of contact representation.

Graphs that admit a VCPG must be planar and $(2,0)$ -sparse (Proposition 4.4). Therefore in this chapter we consider planar $(2,l)$ -tight graphs for varying l . Unless stated otherwise, the graphs we consider are simple, i.e., two vertices are connected by at most one edge and there are no loops.

Kobourov, Ueckerdt and Verbeek [KUV13] have shown that all planar $(2,3)$ -tight graphs admit an L-contact representation, i.e. a strict B_1 -VCPG (a sketch of their algorithm is given in Section 4.2.1). It is immediate that every planar $(2,3)$ -sparse graph also admits a strict B_1 -VCPG. There are graphs that are not $(2,3)$ -sparse but admit a strict B_1 -VCPG, e.g. K_4 . In [KUV13] the question was posed which conditions are necessary and sufficient for a graph to have such a representation.

Chaplick and Ueckerdt show that triangle-free planar graphs, i.e., $(2,4)$ -sparse planar graphs, admit a contact representation of $\{L, \Gamma, |, -\}$ -shapes [CU13]. In other words, triangle-free planar graphs admit a B_1 -VCPG with only two possible orientations of the one-bend paths.

Another related area of research is *orthogonal graph drawing*. Here one wants to draw a graph in the classical setting, i.e., the vertices are points in the plane and the edges are grid-paths. In orthogonal graph drawing there have been many results on minimizing the number of bends. Note that in this setting vertices have at most degree four, or as a workaround, the vertices can be represented as boxes. Tamassia [Tam87] gives an algorithm that obtains, for an embedded graph, an orthogonal drawing with minimal bend number. Optimizing the bend number locally (for each path) has received considerable attention, too. Sch  ffter gives an algorithm that draws 4-regular graphs on a grid with at most two bends per edge (which is tight when not restricted to planar graphs) [Sch95]. For orthogonal drawings without degree restriction, F  lmeier, Kant and Kaufmann have shown that every plane graph has an orthogonal drawing with at most one bend per edge [FKK96].

In Section 4.1 we give a combinatorial description of a VCPG in terms of an orientation of the edges and a flow in the angle graph. A pair, an orientation and a flow, is called *realizable* if there is a VCPG represented by this pair. We give a necessary and sufficient condition for a pair to be realizable. Our proof is constructive. Realizable pairs are then used to obtain bounds on the number of bends of a VCPG.

In Section 4.2 the problem of minimizing the number of bends is addressed. Tight bounds are obtained for planar $(2,2)$ -tight graphs. For planar $(2,1)$ -tight graphs an upper bound is obtained, however, we do not believe that this bound is tight. For simple planar $(2,0)$ -tight graphs, we have not been able to obtain a tight bound either; the bound that we give also holds for planar $(2,0)$ -tight graphs that have loops and double edges. Even in the latter case, we believe the bound is not tight.

4.1 A Combinatorial Characterization of VCPGs

In a similar manner as for SLTRs in Chapter 2, we deduce necessary conditions for the existence of a representation, from a given representation. Then we show that these necessary conditions are also sufficient, and thereby, obtain a combinatorial characterization of VCPGs.

Let $G = (V, E)$ be a planar graph and R a VCPG of G . When we speak about the embedding

of G , we mean the embedding that agrees with R . In R , the vertices of G are represented by internally disjoint grid-paths. We denote the grid-path of $v \in V$ by p_v . Each grid-path has two endpoints. When two grid-paths intersect, the point is an endpoint for precisely one of the grid-paths. An endpoint of p_v coincides with an interior point of p_u precisely then when u and v are adjacent. If an endpoint does not coincide with another grid-path, we call it a *free end*. The number of free ends is denoted by l . For simplicity we consider R to only have free ends in the outer face. A free end in an interior face can be extended to represent an edge, which is considered as dummy edge. Therefore interior free ends can be ignored. Free ends in the outer face reduce the number of bends needed to close the outer face, therefore the simplification is attractive.

We will show that every planar $(2, l)$ -sparse graph is a subgraph of some planar $(2, l)$ -tight graph. From this it follows that the assumption above is not a restriction. Even though this is a very natural problem, we have not been able to find it in the literature¹.

Lemma 4.1. *Every planar $(2, 2)$ -sparse graph is a subgraph of a planar $(2, 2)$ -tight graph.*

Proof. By the theory of Nash-Williams [NW64], every $(2, 2)$ -sparse graph has a decomposition into two edge-disjoint forests, a blue and a red forest. We show that the graph can be extended by induction on the number of trees in a forest. Let the blue forest have at least two trees. Let T_b be a blue tree, it is not spanning. Therefore, there must be a vertex v of T_b that lies in a face f where there is a vertex u which does not belong to T_b . If adding the edge uv does not violate simplicity of the graph, then the edge is added and the blue forest has one tree less. If u and v are neighbors, then a new vertex is added in the face f , the vertex is connected to u and v , those edges belong to the blue forest. The new vertex is connected to a third vertex in f , not u or v , and this edge is added to the red forest. The forests are still spanning and the blue forest has one tree less. By induction, this procedure gives an extension of the graph which has a decomposition into two spanning trees. \square

For $l = 1$ and $l = 0$ the spanning forest decomposition does not exist, as there are too many edges. Instead, we will build a (much) larger graph, of which the sparse graph is an induced subgraph. This is not a very efficient construction, however, it suits the purpose. We consider each vertex to have degree at least 2. Degree 1 vertices can be removed, then the construction is applied. Afterwards the degree 1 vertices are added with Henneberg type 1 steps, they are degree two vertices in the extended graph. Henneberg steps leave tightness intact, hence, the result follows. We also assume the graph to be connected, if it is not, the components can be connected by Henneberg type 1 steps before applying the construction.

Lemma 4.2. *Every planar $(2, 1)$ -sparse graph is a subgraph of a planar $(2, 1)$ -tight graph.*

Proof. Let $G = (V, E)$ be a planar $(2, 1)$ -sparse graph with $|E| = 2|V| - 1 - m$. For a vertex x , let $\pi(x)$ be the smallest integer such that for every set $U \subseteq V$, $x \in U$, the number of edges induced by U is at most $2|U| - 1 - \pi(x)$. In words, π denotes the number of edges that can be added to a set containing x without violating the sparseness condition for a set that contains x . Let X be the set that contains $\pi(x)$ copies of x for all $x \in V$.

Claim. When a graph $G = (V, E)$ is not $(2, 1)$ -tight, then $X \neq \emptyset$.

¹After the final version of this thesis appeared, Alam et al. published their results on contact representations of sparse graph. In this article they also show that every $(2, l)$ -sparse graph can be augmented to a $(2, l)$ -tight graph with the same number of vertices [AEK⁺15].

Suppose the claim does not hold. Let $U \subset V$ be a vertexwise maximal critical set, by assumption $U \neq V$. Let $w \notin U$. Since $X = \emptyset$, there exists a critical set W , such that $w \in W$. Note that there cannot be a vertex that has degree 1 in a critical set, otherwise the set that arises from deleting this vertex violates the sparsity condition.

- If $|W \cap U| \geq 2$, then $E[U \cup W] = |E[W]| + |E[U]| - |E[W \cap U]| \geq 2|W| + 2|U| - 2|W \cap U| - 1$, hence, $U \cup W$ is a larger critical set, contradiction.
- If $|W \cap U| = 1$, then $E[U \cup W] = 2|W| - 1 + 2|U| - 1 = 2|W \cup U| - 1$, hence, $U \cup W$ is a larger critical set, contradiction.
- If $|W \cap U| = 0$, then there cannot be an edge between U and W , as then again $U \cup W$ is a larger critical set. Therefore, there must be a vertex $y \notin W, U$.

If there is such a vertex y , by the same reasoning, there is a critical set that contains y but is disjoint from U and W . It follows now from the fact that G is finite, that at some point there is a vertex which does not belong to any critical set, hence, $X \neq \emptyset$. \triangle

The complete graph on five vertices of which one edge is deleted, is denoted by $K_5 - e$. This graph is (2,1)-tight. Let G' be the graph that is build by consecutively adding a $K_5 - e$ to G . The two components are connected by connecting a representative of an element in X to two vertices of the $K_5 - e$. This step is repeated until the resulting graph is (2,1)-tight, by then m copies of $K_5 - e$ are added. We have obtained a graph with $|V| + 5m$ vertices and the number of edges is

$$|EG'| = 2|V| - 1 - m + 9m + 2m = 2(|V| + 5m) - 1.$$

By construction, there is no subset of vertices U which induces more than $2|U| - 1$ edges, hence, G' is (2,1)-tight. Furthermore, G' is planar, every $K_5 - e$ can be drawn in a face to which its connecting vertex belongs, and G is an induced subgraph of G' . \square

Lemma 4.3. *Every planar (2,0)-sparse graph is a subgraph of a planar (2,0)-tight graph.*

Proof. Let $G = (V, E)$ be a planar (2,0)-sparse graph with $|E| = 2|V| - m$. For a vertex x , let $\pi(x)$ be the smallest integer such that for every set $U \subseteq V$, $x \in U$, the number of edges induced by U is at most $2|U| - \pi(x)$. In words, π denotes the number of edges that can be added to a set containing x without violating the sparseness condition for a set that contains x . Let X be the set that contains $\pi(x)$ copies of x for all $x \in V$.

Claim. When a graph $G = (V, E)$ is not (2,0)-tight, then $X \neq \emptyset$.

Suppose the claim does not hold. Let $U \subset V$ be a vertexwise maximal critical set, by assumption $U \neq V$. Let $w \notin U$. Since $X = \emptyset$, there exists a critical set W , such that $w \in W$. Note that there cannot be a vertex that has degree 1 in a critical set, otherwise the set that arises from deleting this vertex violates the sparsity condition.

- If $|W \cap U| \geq 2$, then $E[U \cup W] = |E[W]| + |E[U]| - |E[W \cap U]| \geq 2|W| + 2|U| - 2|W \cap U| = 2|W \cup U|$, hence, $U \cup W$ is a larger critical set, contradiction.
- If $|W \cap U| = 1$, then $E[U \cup W] = 2|W| + 2|U| = 2|W \cup U| + 1$, this contradicts sparsity.
- If $|W \cap U| = 0$, then $U \cup W$ is a larger critical set, contradiction.

Hence, there must be a vertex that does not belong to a critical set, hence, $X \neq \emptyset$. \triangle

The octahedron is denoted by T_6 , as it is a triangulation on six vertices. This graph is (2,0)-tight. Let G' be the graph that is build by consecutively adding copies of T_6 . The two

components are connected by connecting a representative of an element in X to a vertex of the T_6 . The step is repeated until the graph is $(2,0)$ -tight, by then, m copies of T_6 are added. We have obtained a graph with $|V| + 6m$ vertices and the number of edges is

$$|EG'| = 2|V| - m + 12m + m = 2(|V| + 6m).$$

By construction, there is no subset of vertices U which induces more than $2|U| - 1$ edges, hence, G' is $(2,0)$ -tight. Furthermore, G' is planar and G is an induced subgraph of G' . \square

Remark. An easier proof is obtained from the VCPG representation later on.

4.1.1 ≤ 2 -Orientation representing Edges

We start with an easy proposition which gives an upper bound on the number of edges in a VCPG.

Proposition 4.4. *If $G = (V, E)$ admits a VCPG then:*

$$\forall W \subseteq V : |E_W| \leq 2|W|, \quad (4.1)$$

where E_W is the set of edges induced by W .

Proof. Each edge is represented as a proper contact between two grid-paths representing two vertices. For one of the vertices this contact point is an endpoint of its grid-path. This vertex is the representative for the edge. Each vertex can be a representative for at most two edges. Therefore the number of edges induced by a set $W \subseteq V$ is at most twice the number of vertices. \square

Defining $u \rightarrow v$ if p_u ends on p_v , we obtain an orientation on G from a VCPG. Every vertex has outdegree at most two in this orientation. The vertices with outdegree less than two, are precisely those for which the grid-path has one or two free ends.

A VCPG is not completely described by a plane graph G and a ≤ 2 -orientation. Consider the VCPGs in Figure 4.2, the VCPGs induce the same orientation, but the number of bends in the representations is different.




Figure 4.2: Two VCPGs of the octahedron, which induce the same 2-orientation (shown in the middle) but the number of bends differs.

For simplicity we write 2-orientation instead of ≤ 2 -orientation, it may be clear that for a $(2,l)$ -tight graph with $l > 0$ not every vertex has outdegree precisely two. Later, we also

assume that the vertices that do not have outdegree 2, are on the outer face. With the next lemma we show that every $(2, l)$ -tight graph has such an orientation.

Lemma 4.5. *Every planar $(2, l)$ -tight graph has a 2-orientation. If $l > 0$, then for every embedding, there exists a 2-orientation, such that, all vertices with outdegree less than 2, are on the boundary of the outer face.*

Proof. Let $G = (V, E)$ be a planar $(2, l)$ -tight graph. Suppose there is a subset W of the vertices of G that has less than $2|W|$ incident edges. Then $G[V - W]$ must induce at least:

$$2|V| - l - (2|W| - 1) = 2|V - W| - l + 1$$

edges, which contradicts $(2, l)$ -tightness. Hence, every subset W of the vertices of G has at least $2|W|$ incident edges. Now we construct a bipartite graph B . The first vertex class, V_1 , consists of two copies of all but l of the vertices of G . The remaining l vertices are only added once. These l vertices can be chosen arbitrarily. The second class, V_2 , consists of the edges of G . The edge set of B is defined by the incidences in G : two vertices are connected if the corresponding vertex of G is an endpoint of the corresponding edge of G . We will show that $(2, l)$ -tightness of G implies that this bipartite graph has a perfect matching. A perfect matching defines a 2-orientation of G .

Let $U \subseteq V_1$ that consists of n vertices with their copies, and m vertices without a copy. So $|U| = 2n + m$. Let $N_B(U)$ denote the set of neighbors of U in B . The neighbors are all edges in G , except for the edges between two vertices that are not in U . There are at most $2(|V| - n - m) + l$ such edges. Therefore, the number of neighbors of U is at least:

$$|N_B(U)| \geq |E| - 2(|V| - n - m) + l \geq 2|V| - l - 2|V| + 2n + 2m + l = 2(n + m) \geq |U|.$$

Hence, Hall's marriage condition is satisfied for each subset of vertices of V_1 . Let $A \subseteq V_2$ be a subset of the edges of G and $V_A = V \cap A$ is the set of endpoints of the edges in A .

$$|A| \leq 2|V_A| - l \leq |N_B(A)|$$

Hence, Hall's marriage condition is also satisfied for each subset of vertices of V_2 .

To prove the second part of the lemma we fix an embedding of G . The vertices that are not added twice to V_1 are chosen arbitrarily from the boundary of the outer face. The counting argument still holds, therefore we obtain a perfect matching and thus the desired 2-orientation. \square

Remark. Lemma 4.5 holds for all $(2, l)$ -tight graphs, not only simple and planar $(2, l)$ -tight graphs. In the proof simplicity nor planarity is used, except for the selection of the vertices that have outdegree less than two. In a non-planar graph the vertices with outdegree less than two should be selected differently than "being on the outer face".

4.1.2 A Feasible Flow representing Bends

From here on, we consider the graph to be 2-connected. Note that any $(2, l)$ -tight graph can easily be extended to a 2-connected $(2, l)$ -tight graph by adding an appropriate number of degree two vertices.

To describe the behavior of the bends of a grid-path, we introduce a flow network similar to the flow network introduced by Tamassia [Tam87]. A flow network \mathcal{N} is a graph that has

two distinct sets of vertices that are the sources and the sinks of the network. The sources and the sinks in our case, as well as in the network designed by Tamassia, are the faces of the graph (see Figure 4.3). For a face of a graph G , needing a convex angle of a bend is presented as having an excess, i.e., it is a source. Similarly, needing a concave angle of a bend is presented as having a demand, i.e., it is a sink. Whenever we speak about excess, the face is a source and when we speak about demand, the face is a sink.




Figure 4.3: Examples of faces drawn by grid-paths, the number in the face denotes the number of convex angles needed.

A path going from a source to sink must go over vertices, as the vertices will be represented as grid-paths with bends. Therefore, the network $\mathcal{N} = (A(G), S, T)$ is the *angle graph* together with a subset of the vertices of the angle graph that represents the sources and a subset that represents the sinks. The angle graph $A(G)$ is a plane bipartite graph in which each edge represents an angle of G . Formally, the angle graph arises from G by setting the union of the vertices and faces of G as the vertices of $A(G)$ and the edges of $A(G)$ are the pairs vf , $v \in V(G)$, $f \in F(G)$, such that v is a vertex on f in G (see Figure 4.5). The angle graph of a 2-connected plane graph is a maximal planar bipartite graph. The outer face has different properties in a VCPG, therefore we consider graphs with a fixed embedding.

A path in the network, going from a source to a sink, passing through a vertex v , represents a bend in the grid-path of the vertex v . Let f_1 be the previous node with respect to the direction of the flow, and f_2 be the next node, with respect to v . Then this bend of v is such that, the convex angle lies in f_1 and the concave angle lies in f_2 . The difference between the network needed for VCPGs and the network defined by Tamassia for orthogonal drawings is that in the case of VCPGs the flow goes through a vertex and in the case of orthogonal drawings the flow goes through an edge. In other words, the underlying graph of the network defined by Tamassia is the dual graph and not the angle graph (see Figure 4.4). We will now proceed with the formal introduction of the encoding of the bends of a VCPG.




Figure 4.4: An example of a feasible flow for K_4 with as underlying graph the dual graph (left) or the angle graph (right).

A flow ψ is a weighted directed graph, with the network \mathcal{N} as underlying graph. The weight of an edge, denotes the amount of flow that is routed over this edge. On the right side of

Figure 4.5 a flow ψ is depicted. The edges that get weight zero are left out. The edges that do not have weight zero are labeled with their weight. The face-vertices of $A(G)$ can be a source or a sink, depending on the size of the face in G . The vertex-vertices of $A(G)$ are neither sources nor sinks. The capacity of the edges is unbounded. From a VCPG R we construct a flow ψ as follows. For each bend of p_v , such that the convex angle of this bend lies in f_1 and the concave angle lies in f_2 , add a unit of flow $f_1 \rightarrow v \rightarrow f_2$.




Figure 4.5: The angle graph (black) of the octahedron (grey) and the angle graph with the feasible flow induced by the VCPG on the left of Figure 4.2.

Proposition 4.6. *Let ψ be the flow obtained from a VCPG of a $(2, l)$ -tight graph G . Then the following holds. Every interior 3-face has excess 1. Every interior k -face f , for $k > 3$, has demand $|f| - 4$. The outer face f_∞ has demand $(4 - 2l) + |f_\infty|$.*

Proof. Following the boundary of an interior region of a VCPG and adding the changes in direction one should obtain 2π . Each edge is represented as a proper contact and therefore changes the direction with $\pi/2$. A convex angle at a vertex changes the direction with $\pi/2$ as well and a concave angle at a vertex changes the direction with $-\pi/2$. For a face f , let $c(f)$ be the number of concave angles at vertices minus the number of convex angles at vertices. For an interior face f we obtain:

$$2\pi = \pi/2 \cdot |f| - \pi/2 \cdot c(f). \quad (4.2)$$

The value $c(f)$ also counts the incoming minus the outgoing flow in ψ . Therefore, by rearranging 4.2 to

$$c(f) = 4 - |f|$$

we obtain the desired result.

For the outer face f_∞ , the bends have to be counted differently. For the outer face, an edge gives a change in direction of $-\pi/2$. A free end gives a change in direction of π . For the angles at vertices there is no difference, a convex angle at a vertex changes the direction with $\pi/2$ and a concave angle at a vertex changes the direction with $-\pi/2$. Similarly as for the interior faces we obtain

$$c(f_\infty) = (2l - 4) - |f_\infty|$$

where l is the number of free ends. As $c(f_\infty) < 0$, the outer face is a sink with demand $(4 - 2l) + |f_\infty|$ in ψ . \square

A valid flow for the network \mathcal{N} , satisfies the flow conservation law at each vertex that is not a source or a sink, i.e., the net flow is zero at each vertex that is not a source or a sink. At a

sink, the sum of the incoming flow minus the sum of the outgoing flow is at most the demand of the sink. At a source, the sum of the outgoing flow minus the sum of the incoming flow is at most the size of the excess. A valid flow ψ is a feasible flow for the network \mathcal{N} if the demand of every sink is satisfied. As the demands of the sinks add up to the sum of the excesses of the sources, in a feasible flow the excess of each source is also completely used. The total value of a feasible flow is equal to the number of interior 3-faces (the number of sources), which in turn is equal to the sum of the demands of all sinks.

$$\sum_{f \in F_{\text{int}}} (4 - |f|) + (2l - 4) - |f_{\infty}| = 4|F| - \sum_{f \in F} |f| - 2l - 8 = 2(2|V| - |E| - l) = 0$$

The number of bends prescribed by the flow is

$$\sum_{v \in V} \psi(v),$$

where $\psi(v)$ denotes the amount of flow that goes through the vertex v . In order to relate to a VCPG it is necessary that ψ is an *integral* feasible flow for the network \mathcal{N} . Note that since the demands and excesses are integral the existence of a (minimum cost) integral feasible flow is guaranteed, for example by the result of Tardos [Tar85]. In the sequel we will omit the word ‘integral’ as we only consider integral feasible flows.

A flow does not uniquely define a VCPG either. The two VCPGs in Figure 4.6 induce the same feasible flow, but the edges are represented in a different way. The obvious question is whether every feasible integral flow in $A(G)$ belongs to a VCPG. Unfortunately this is not the case (see Figure 4.7). However, using both the orientation and the feasible flow, we will obtain a characterization of VCPGs.




Figure 4.6: A feasible flow does not uniquely define a VCPG. On the left two VCPGs and on the right the feasible flow induced by both of the VCPGs. The difference is the orientation of the cycle e, a, b, f (in red).

4.1.3 Realizable pairs

Starting with a VCPG, a 2-orientation α and a feasible flow ψ can be obtained, as shown in the previous sections. In this section we identify a property of such a pair, (α, ψ) that comes from a VCPG. This ensures that the property is necessary. Not every pair (α, ψ) on G induces a VCPG of G . We call a pair (α, ψ) *realizable* when it does. We will prove that the necessary property is also sufficient, hence, realizable pairs are in bijection to VCPGs. Our proof method is algorithmic, it shows how one can construct a VCPG (the geometric setting) from a realizable pair (the combinatorial setting).




Figure 4.7: A feasible flow ψ that does not belong to a VCPG, i.e. there exists no 2-orientation α such that the pair (α, ψ) is realizable.

A Property of Realizable Pairs

The property depends on α and on ψ . First we define the property for a vertex v (see Figure 4.8 (c)). Let $A[N_{A(G)}[v]]$ denote the angle graph induced by the closed neighborhood of a vertex v , i.e. induced by v and all its neighbors in $A(G)$. Let n_1, n_2 be the neighbors of v along the outgoing edges of v in α . If v has outdegree 0, then n_1, n_2 are its neighbors on the outer face. If v has outdegree 1, then n_1 is its neighbor along the outgoing edge. The vertex n_2 is the neighbor of v on the outer face, chosen such that the units of flow are equally distributed on the clockwise and counterclockwise side of the path n_1, v, n_2 . Informally, a unit of flow through a vertex v represents a bend of the grid-path of v . Following the grid-path from n_1 to n_2 , looking left and right, the bends are met at the same time. This implies that the flow through a vertex must be laminar, i.e., non-crossing.

Definition 4.7 (Realizability Condition). The pair (α, ψ) satisfies the realizability condition at vertex v if and only if, given $A[N_{A(G)}[v]]$ and the flow in this subgraph, (see Figure 4.8 (b))

- There is a decomposition of the flow into non-crossing paths, and,
- Every path of such a decomposition crosses the path n_1, v, n_2 .

When the pair (α, ψ) satisfies the realizability condition at each vertex we say that the pair is realizable.

Consider the feasible flow in Figure 4.7, call this flow ψ . The graph is (2,0)-tight and therefore all vertices have outdegree 2 in an orientation that comes from a VCPG. A 2-orientation α such that (α, ψ) is a realizable pair, demands that $u \rightarrow v$ and $w \rightarrow v$. Then v has outdegree at most 1 and there has to be a vertex with outdegree 3 in α . Therefore, α cannot come from a VCPG. We conclude that there does not exist a 2-orientation α such that (α, ψ) is realizable and therefore ψ does not come from a VCPG.

First we show that a pair that comes from a VCPG satisfies the realizability condition.

Lemma 4.8. *A pair (α, ψ) that comes from a VCPG is realizable.*

Proof. First note that a VCPG of G describes an embedding of G . If there is a grid-path with one free end, then before proceeding we reduce all unnecessary bends, i.e. if a grid-path has




Figure 4.8: (a) The vertex v in G and the local flow. (b) The expansion of v : The flow through v is decomposed into disjoint paths and each of these paths cuts the path consisting of the outgoing edges of v . (c) A close look on the grid-path of a vertex v in a VCPG.

bends between its last neighbor and its free end, these bends are removed. A 2-orientation can be constructed from a VCPG by orienting an edge $u \rightarrow v$ if and only if the grid-path of u ends on the grid-path of v . Consider the grid-path that represents a vertex v . If this path has no bends, the realizability condition is satisfied at this vertex. Suppose the path has k bends. Draw an arrow from the face containing a convex corner to the face in which the associated concave corner lies. Now the set of arrows represents the flow $\psi(v)$. This flow is non-crossing through v and every unit of flow is cut by the the grid-path of v . When these arrows are introduced for all bends of all grid-paths, the flow given by these arrows satisfies the demand of each face. Contract the strictly interior steps of the grid-path to a vertex and every unit of the non-crossing flow through v is now cut by the outermost two segments of the grid-path, which correspond to the outgoing edges of v , or to the outgoing edge and the location of the last incoming edge before the free end of the grid-path. Hence the realizability condition is satisfied at each vertex, therefore the pair (α, ψ) obtained from the VCPG is realizable. \square

4.1.4 Realizable Pairs are in Bijection with VCPGs

Now we are ready for the main result.

Theorem 4.9. *The realizable pairs are in bijection with VCPGs.*

The remainder of this section is dedicated to showing that given a realizable pair (α, ψ) , there exists a VCPG, with the same vertices on the outer face as in the chosen embedding of G , and such that:

- (a) The grid-path of u ends on the grid-path of v if and only if the edge uv is oriented from u to v in α , and,
- (b) The grid-path of v has precisely $\psi(v)$ bends.

We will show how to construct a VCPG given a realizable pair. Note that an embedding can be derived from $A(G)$ (in which the flow ψ is defined). Consider a realizable pair (α, ψ) . The construction consists of four steps, which we first outline here.

Step 1: First we expand the vertices that have k units of flow going through them, to a path² of length k . We obtain a bipartite graph.

²The length of a path counts the number of edges.

Step 2: We introduce help-edges and vertices in the bipartite graph to construct a quadrangulation (see Figure 4.9 (b)). The orientation α is extended to a 2-orientation of the quadrangulation.

Step 3: We then find a segment contact representation of the quadrangulation. It has been shown that the 2-orientations of maximal bipartite planar graphs are in bijection with separating decompositions of this graph (e.g. [dFdM01]). In turn a separating decomposition induces a segment contact representation (cf. [HNZ91, dFdMP95, Fel13]). Hence we can construct a segment contact representation where the representation of the edges is in bijection with the given 2-orientation. An example is shown in Figure 4.9 (c).

Step 4: Last, we will show that the extra edges that have been introduced to make a quadrangulation of the bipartite graph can be deleted in order to obtain a VCPG of G (see Figure 4.9 (d)).




Figure 4.9: From a realizable pair to a VCPG: (a) a plane (2,0)-tight graph with a realizable pair (ψ in red); (b) expanding the vertices according to the flow (in red) and extending the bipartite graph to a quadrangulation (in green); (c) a segment contact representation of the quadrangulation with the segments that belong to the original graph highlighted; (d) a VCPG.

Let us describe the steps in more detail.

Step 1. Let (α, ψ) be a realizable pair for G . We expand all vertices with non-zero flow. The plane graph we obtain is denoted by \tilde{G} . For every vertex v for which $\psi(v) \neq 0$, expanding v

consists of the following steps (see Figure 4.8):

1. Expand v to a circle, we will denote this the *bag* of v .
2. Inside the circle, add a path with $\psi(v) + 1$ vertices between the two outgoing edges of v . If v has outdegree 1 or 0, then v is on the outer face. The path is added between the outer face and the outgoing edge, or between the two edges on the outer face, respectively. The edges of the path are denoted by *path-edges*, the inner vertices of the path by *path-vertices*, and all the vertices that belong to a bag are denoted by *bag-vertices*.
3. Connect the edges that end on the circle to the path-vertex in such a way that the flow between two faces only crosses an edge of the path.

Step 2. After all the expansions have been done we obtain a graph where all faces have even length. Each face gets $|4 - |f|| + 2k$ extra vertices due to the expanding step, where k is the amount of flow proceeding through the face. The resulting faces in \tilde{G} have size $|f| + |4 - |f|| + 2k$, for $|f| = 3$ this gives $4 + 2k$ and for $|f| > 3$ this gives $2|f| - 4 + 2k$, both even. So in \tilde{G} all faces have even length and therefore \tilde{G} is a bipartite graph. Now we add the help-edges to extend \tilde{G} to a quadrangulation. We denote the quadrangulation by G_Q . We will also orient the new edges to obtain a 2-orientation of G_Q . In order to explain how the help-edges are added, we need the following lemma.

Lemma 4.10. *Every interior face \tilde{f} of \tilde{G} has $(|\tilde{f}| - 4)/2$ units of incoming flow.*

Proof. Let $\psi^+(f)$ (respectively $\psi^-(f)$) denote the incoming (respectively outgoing) flow in face f . Let \tilde{f} be the equivalent face of \tilde{f} in the original graph G .

As ψ is a feasible flow, the demand $c(f)$ of f , is the difference between incoming and outgoing flow:

$$\psi^-(f) - \psi^+(f) = c(f) = 4 - |f|.$$

Now we use that the size of the extended face \tilde{f} is the size of f plus the incoming and the outgoing flow.

$$|\tilde{f}| = |f| + \psi^+(f) + \psi^-(f) = |f| + \psi^+(f) + \psi^-(f) - |f| + 4 = 2\psi^+(f) + 4$$

Hence we find

$$\psi^+(f) = \frac{|\tilde{f}| - 4}{2}. \quad (4.3)$$

□

Using $(|\tilde{f}| - 4)/2$ edges, we can quadrangulate \tilde{f} (see Figure 4.10). The help-edges should be added in such a way that every bag (vertex expansion) gets as many help-edges as it has flow going into \tilde{f} . Informally, a concave corner arises from two segments that both end in one point. In the theory of segment contact representations that we will use, there are only proper contacts or free ends. Each help-edge represents a segment of a concave corner that proceeds into the face, and, this part will later be removed. Such a help-edge will be oriented outgoing from the bag.

Later, we construct the segment contact representation. For this, it is necessary that every interior vertex has outdegree precisely two. Therefore the help-edges must be added in such

a way that this is possible for all vertices. Each interior bag-vertex should gain precisely two outgoing arcs (which are not edges in the original graph). The vertices on the boundary of the bag, i.e., the vertices incident to the outgoing edges of the vertex in the original graph, gain precisely one outgoing arc that is not in the original graph. The help-edges are added along the flow from a vertex into a face, this will give the correct amount of new edges for every bag.

Lemma 4.11. *Each inner face \tilde{f} of \tilde{G} can be quadrangulated in such a way that each bag through which k units of flow enter \tilde{f} gets k new outgoing arcs. The outer face of \tilde{G} can be quadrangulated using four help-vertices (v_t, v_r, v_b, v_l) , in such a way that each bag through which k units of flow enter the outer face gets k new outgoing arcs.*

An example of quadrangulating an interior face is depicted in Figure 4.10. The flow ψ is given by the red arrows. First half-arcs are added, the green solid arcs. Then these half-arcs are subsequently connected in such a way that they close one 4-face (green dashed lines). The quadrangulation of the outer face is based on the same idea.

Proof. Assign a half-arc into \tilde{f} from each vertex that comes clockwise after a unit of incoming flow, see Figure 4.10. There are $(|\tilde{f}|-4)/2$ vertices who get a half-arc, thus $(|\tilde{f}|+4)/2$ without a half-arc. Hence there exists a vertex with a half arc which is followed by two consecutive vertices u, v without a half-arc. Consider the half-arc clockwise before u, v and connect it to the vertex after u, v . We have constructed a 4-face and completed one half-arc. This step can be repeated considering the resulting face and its half-arcs. From counting it follows that all half-arcs can be extended into arcs with this method and that the result is a quadrangulation of \tilde{f} in such a way that each bag through which k units of flow enter \tilde{f} gets k new outgoing arcs.



Figure 4.10: Adding help-edges in a face. The flow ψ is depicted by red arcs. The green half-arcs together with the dashed extensions represent the help-edges.

To take care of the outer face we add a quadrilateral around the graph and quadrangulate the new inner face between the graph and the quadrangle.

We distinguish four cases:

- (a) There exists a vertex s in the original graph G which has no outgoing arcs under α .
- (b) There exist two vertices s, t in the original graph G which both have precisely one outgoing arc under α .
- (c) There exists precisely one vertex s which has precisely one outgoing arc under α .
- (d) All vertices have outdegree 2 under α .

We add a quadrilateral around the graph and construct an inner face, of size $2k + 4$ and such that the amount of incoming flow is k (see Figure 4.11). Then we can use the same method as for the interior faces.




Figure 4.11: Adding a quadrilateral around the graph.

Recall that we consider $(2, l)$ -tight graphs. Hence the graph has $2|V| - l$ edges and there are precisely l free ends in the outer face of the representation.

(a) Note that $l = 2$. Add a quadrangle around the graph with vertices v_l, v_t, v_r, v_b in clockwise order. If s is expanded, we label the vertices in the expansion in counterclockwise order, following the boundary of \tilde{f}_∞ , by s_1, \dots, s_k . If s is not expanded then we label $s = s_1$. Add the arcs (s_1, v_t) and (s_1, v_b) . Now the bounded face f^* containing s_1, v_t, v_r, v_b on its boundary has the following properties: $|f^*| = |\tilde{f}_\infty| + 4$, it has $\frac{1}{2}|\tilde{f}_\infty|$ incoming flow. The same method as for an inner face can be used: first half-arcs are added and then these are extended to arcs by consecutively closing 4-faces.

(b) Note that $l = 2$. Add a quadrangle around the graph with vertices v_l, v_t, v_r, v_b in clockwise order. If s and t are expanded, we label the expansion vertices in the respective bags such that s_1 and t_1 have no outgoing arc under α and they are end vertices of the extension path. If s respectively t is not expanded, we label $s = s_1$ respectively $t = t_1$. Add the arc (s_1, v_t) and let f^* be the face between the quadrangle and \tilde{f}_∞ . Let ψ_{t_1, s_1} denote the incoming flow to f^* between t_1 and s_1 clockwise around \tilde{f}_∞ . Assign the label q to the vertex at distance $2\psi_{t_1, s_1} + 3$ from t_1 walking counterclockwise around f^* . Add the arc (t_1, q) . Now we have obtained two faces f_u, f_d , for which the incoming flow $\psi^+(f_u), \psi^+(f_d)$ is equal to $|f_u|/2 - 2, |f_d|/2 - 2$. The same method as for an inner face can be used: first half-arcs are added and then these are extended to arcs by consecutively closing 4-faces.

(c) Note that $l = 1$. Add a quadrangle around the graph with vertices v_l, v_t, v_r, v_b in clockwise order. The vertex s has outdegree precisely one under α . Label the vertex in the bag of s that has no outgoing edge and is an end vertex of the extension path s_1 or if s is not expanded we label $s = s_1$. Add the arc (s_1, v_t) . We obtain a face f^* between the quadrangle and \tilde{f}_∞ for which the incoming flow is of size $\frac{1}{2}|\tilde{f}_\infty| + 1$ and $|f^*| = |\tilde{f}_\infty| + 6$. The same method as for an inner face can be used: first half-arcs are added and then these are extended to arcs by consecutively closing 4-faces.

(d) Note that $l = 0$. Add a quadrangle around the graph with vertices v_l, v_t, v_r, v_b in clockwise order and let f^* be the face between the quadrangle and \tilde{f}_∞ . We will use one unit of flow to connect \tilde{f}_∞ to the quadrangle. First add a half-arc into f^* from each vertex clockwise after a unit of incoming flow. Choose any half-arc and connect it to v_t . We obtain a face f with $\frac{1}{2}|\tilde{f}_\infty| + 1$ incoming flow and $|f| = |\tilde{f}_\infty| + 6$. The same method as for an inner face can be used: first half-arcs are added and then these are extended to arcs by consecutively closing 4-faces. \square

To obtain a 2-orientation of the quadrangulation G_Q , the edges that are strictly inside the

bags, and the four boundary edges need to be oriented. The orientation of the original edges is inherited from α , the help-edges are already oriented. Each bag b_v contains $|b_v| - 1 = \psi(v)$ edges which are not yet oriented, all others are oriented and such that, in total, b_v has outdegree $|b_v| + 1$.

Lemma 4.12. *Each bag b_v in G_Q has precisely outdegree $|b_v| + 1$ and each vertex in b_v has outdegree at most two.*

Proof. The bag b_v of a vertex v has $\psi(v) + 1$ vertices. According to the flow, $\psi(v)$ outgoing arcs are added at the introduction of the help-edges.

If a bag b_v comes from a vertex v which has outdegree 2, the bag has outdegree $|b_v| + 1$. If a bag b_v comes from a vertex v which has outdegree 1 in α , then the quadrangulation step has assigned another outgoing arc to this bag. Therefore, the bag has outdegree $|b_v| + 1$. If a bag b_v comes from a vertex v which has outdegree 0 in α , then the quadrangulation step has assigned two outgoing arcs to this bag. Therefore, the bag has outdegree $|b_v| + 1$.

Suppose one of the bag-vertices has outdegree three or more (see Figure 4.12). At most one of the arcs is an edge of the original graph, as otherwise the vertex would not have been expanded and no new outgoing arc is added. When an outgoing arc is added to a vertex, the vertex is clockwise after a path-edge³ which is crossed by a unit of flow routed into the adjacent face (for which the vertex is clockwise after the edge). If a vertex has two added arcs, it must be incident to two faces, one on each side of the path and such that the vertex is clockwise next of the path-edge with respect to this face. Hence, this vertex is incident to two path-edges and is not an end vertex of the path. Therefore, it cannot have an outgoing arc that is an edge of the original graph. It follows that every vertex has outdegree at most 2. \square




Figure 4.12: There is no outdegree-3 vertex.





Figure 4.13: Greedily orienting edges of the path: a vertex of degree two at the end of the path, or internal to the path.

Orient $v_l v_t$, $v_r v_t$, $v_l v_b$ and $v_r v_b$ towards v_t and v_b , respectively. The vertices v_t and v_b are the two poles of the 2-orientation. The orientation can be completed in a greedy way.

Lemma 4.13. *The path-edges can be oriented (greedily) such that the resulting orientation is a 2-orientation of G_Q .*

Proof. We subsequently orient the edges of a path towards a vertex which has outdegree 2 and show that in each step, all neighbors of a vertex with outdegree 2, have outdegree at most 1. Hence, we can continue orienting edges until all the edges are oriented (see Figure 4.13).

³Recall that by path-edge we denote the edges that are added during the expansion of the vertex.

A neighbor of a path-vertex with outdegree 2 cannot also have outdegree 2. This follows from the fact that through every path-edge at most one unit of flow proceeds and, hence, a neighboring pair in the path has together outdegree at most 3. Consider a vertex v which has outdegree 2. Let u be a neighbor of v on the path, v has at most outdegree 1. We need to show that if it has outdegree 1, its next neighbor has at most outdegree 1 or it is an end vertex of the path. Because then, the path-edge between u and v can be oriented towards v . Suppose u has outdegree 1 and it is not an end vertex of the path. Let w be the next vertex on the path. Vertex v is the clockwise subsequent vertex for the flow through uv , so the outgoing arc of u must come from the flow through uw . Hence for w the only possibility is an arc that comes from the flow through the next edge of the path, or, w is an end vertex and has an outgoing edge in α . Either way, w has outdegree at most 1. Hence, the path-edge uv is oriented from u to v . We consider another vertex with outdegree 2.

As the number of path-edges is precisely the number of outarcs needed for every bag, this process will end and each bag-vertex has outdegree precisely 2. The vertices on the outer face, have outdegree 2 or are a pole. All other vertices have outdegree 2 in α and did not get new outgoing arcs. Hence, we have a 2-orientation $\hat{\alpha}$ of G_Q . \square

Step 3. From G and the realizable pair (α, ψ) we constructed the quadrangulation G_Q with a 2-orientation $\hat{\alpha}$ (v_t and v_b are the only two vertices with outdegree 0 instead of outdegree 2). We construct a segment contact representation, the vertices of the two color classes of G_Q become horizontal and vertical segments and the edges are proper contacts between the segments satisfying $\hat{\alpha}$. A method to obtain such a representation from a 2-orientation is given in Chapter 1, on page 6 and further.

Step 4. Last to show is that this segment representation of G_Q is equivalent to a VCPG of G where the path of a vertex v is given by its outgoing arcs in α and $\psi(v)$ denotes the number of bends of the path of v . The segment contact representation of G_Q does not necessarily contain the VCPG. That is to say, it might be such that a grid-path that is supposed to end on a particular part of another grid-path, does not end there. In the example in Figure 4.15 (c) on page 115, the path of 1 does not end on the path of 10 that will be part of the grid-path of vertex b . The following lemma shows that the sets of segments ending on different sides of a segment, can be moved independently. Hence, the segment contact representation can be changed such that all contacts are as we want them to be.

Lemma 4.14. *Given a vertical segment in a segment contact representation, then its (horizontal) left neighbors can be shifted independently from its (horizontal) right neighbors. The same holds for a horizontal segment and its top respectively bottom incoming neighbors.*

Proof. Suppose the horizontal segment v_b ends on the left and v_t ends on the right of the vertical segment h . The segment v_b is below the segment v_t and we want it to be above of v_t . In Figure 4.14 an example is shown where $h = 10$, $v_b = 9$ and $v_t = 1$.

Consider a cutline as follows: from h just below of v_b to the left of the drawing, such that, it intersects only horizontal segments; from h just to above v_t to the right, such that it only intersects horizontal segments; and the part on h between the these two horizontal cutlines. We cut the graph along the cutline. We move the half containing v_b to the top, such that v_b now is above v_t . Extend all the vertical lines that are cut. We have obtained a new, equivalent, segment contact representation where v_b is above v_t . \square

Theorem 4.15. *The segment representation of G_Q obtained from a realizable pair (α, ψ) induces a VCPG of G .*




Figure 4.14: Cutting open and shifting shows that neighbors on each side can be moved independently: (a) a segment contact representation where the highlighted part of 1 does not end on the highlighted part of 10, and a cutting line (dashed); (b) the segment contact representation is cut and the top is pulled upwards extending the vertical segments that are cut.

Proof. We apply identification, shifting and deletion to the segment representation of G_Q and will then show that the result is a VCPG of G .

Identification. For each vertex v that is not expanded we color the segment with color v . For each bag b_v , select the segments representing the bag-vertices v_1, \dots, v_k . For each bag-vertex v_i color the part of the segment between v_{i-1} and v_{i+1} with color v . For the end vertices of the bag, v_1 (respectively v_k) color the part of the segment between v_2 (respectively v_{k-1}) and the outgoing neighbor of v_1 (respectively v_k) that is not in the bag. Now the grid-path of vertex v is highlighted among the selected segments. When $l = 1, 2$ we have added arcs while quadrangulating the outer face, that do not correspond to the flow nor to original edges. These edges represent the free ends of grid-paths. Hence we extend these to be just further than the last neighbor ending on this segment.

Shifting. It may occur that for an arc of G , say (u, v) , the endpoint from u on v appears on a non-highlighted part of v , in this case we need to shift. Let v_i be the vertex represented by the segment on which u ends, and v_{i-1}, v_{i+1} the neighbors of v_i in the bag. We need that around v_i , there are consecutively one outgoing edge, at most one incoming edge from a bag neighbor, say v_{i-1} , incoming other edges, one outgoing edge, at most one incoming edge from a bag-neighbor, say v_{i+1} , incoming other edges. In other words, in either clockwise or counter clockwise order, there is no incoming edge of the original graph between the outgoing edge and the incoming bag edge. This ensures (by Lemma 4.14) that there is a segment representation in which the contacts of edges of the original graph to v_i lie between the contacts of v_{i-1} and v_{i+1} with v_i . The statement follows from the construction of G_Q . If the outgoing edges of v_i are to v_{i+1} and v_{i-1} it is trivial. Suppose not, without loss of generality, the outgoing edge north of v_i can only be induced by flow through (v_{i+1}, v_i) and south of v_i by (v_{i-1}, v_i) . If both appear, then these outgoing edges occur just before the incoming bag edges in clockwise or counterclockwise order, hence, we are done. If at most one appears, say v_i has one outgoing edge to v_{i+1} , then the other outgoing edge is just before the incoming edge from v_{i-1} , i.e. it comes from the flow going through the edge (v_i, v_{i-1}) .

It follows that there exists a segment representation in which the contacts of edges of the original graph to v_i lie between the contacts of v_{i-1} and v_{i+1} with v_i . Moreover in the representation we have, we can shift v_{i-1} and v_{i+1} such that all other contacts to v_i lie between them. The coloring extends trivially along the shifting.

Deletion. After the shifting, all endpoints representing edges of G occur between highlighted segments. Therefore, we can delete all non-highlighted parts of the segments without losing edges of the original graph.

Conclusion. It follows from the three steps that each edge (u, v) of G is represented by a non-degenerate contact. If this edge is oriented from u to v then the path of u ends on v . Each vertex v is represented by $\psi(v) + 1$ segments, therefore, it is a grid-path with $\psi(v)$ bends. Hence, the result is a VCPG of G that agrees with α and ψ . \square

With the four steps we have obtained a VCPG from a realizable pair. This completes the proof of Theorem 4.9.

4.2 Bounding the Number of Bends

First we show that every planar $(2, l)$ -tight graph, $l \geq 0$, has a VCPG.

When a planar graph has a 2-orientation it easily follows that it has a VCPG. An example is shown in Figure 4.15.

Lemma 4.16. *Let G be a planar graph, not necessarily simple. If G has a 2-orientation, i.e., an orientation in which every vertex has outdegree at most 2, then G admits a VCPG.*

Proof. Consider an embedding of G and a 2-orientation α of G . Subdivide each loop twice. If a pair of vertices is connected by multiple edges, all but one of the multiple edges are subdivided. The result is a simple plane graph, which has a straight-line drawing by Fàry's theorem. Replace each straight-line edge in such a drawing by an axis-aligned grid-path leaving the start and endpoint intact and such that two grid-paths starting in the same point only coincide in this point. The subdivided edges are merged without changing the grid-paths. A vertex is identified with its outgoing edge(s). The last step is to perturb the last straight part of a grid-path p_v that ends on a grid-path p_w in such a way that this point is not used by any grid-path other than p_v and p_w . This procedure gives a VCPG of G that realizes the chosen embedding. \square




Figure 4.15: A VCPG drawn by approximating a straight-line drawing of a graph and using a 2-orientation.

Lemma 4.5 shows that every planar $(2, l)$ -tight graph has a 2-orientation. It follows that every planar $(2, l)$ -tight graph admits a VCPG.

Using this construction there is no control over the number of bends of a grid-path. In contrast to this, a realizable pair precisely determines the number of bends of each vertex. Moreover, costs can be added to vertices in the flow network to request few bends for certain vertices and allow for more bends at others. The characterization also gives a certificate that the total number of bends is minimized. If all the costs are set to one, and there is a realizable pair that attains a minimum flow, then this is a certificate. If this is not the case, then it is easy to see how far away from the minimum the current representation is. However, it is not yet known whether the minimum can always be attained by a representation. In the following sections we will show some results on the number of bends for a vertex. We will try to locally minimize the number of bends for particular graph classes.

Using the VCPG of a $(2, l)$ -sparse graph G , we can obtain a $(2, l)$ -tight graph that has G as a subgraph. Especially for the cases $l = 0, 1$ the amount of extra vertices added is a lot less with this method than in Lemma 4.2 and 4.3.

Lemma 4.17. *Let $l \in \{0, 1, 2\}$. Every planar $(2, l)$ -sparse graph is the subgraph of a planar $(2, l)$ -tight graph and of a planar $(2, 0)$ -tight graph.*

Proof. Let G be a $(2, l)$ -sparse graph. Let B be the bipartite graph consisting of two copies of each vertex of G in one color class, and the edges of V in the second color class. The edges of B are the incidences between the edges and vertices in G . From the sparsity condition it follows that B has a matching in which all the edge-vertices are matched:

$$\forall A \subseteq E : |A| \leq 2|V \cap A| - l \leq 2|V \cap A| = |N_B(A)|.$$

Hence, the graph has a 2-orientation and by Lemma 4.16 G admits a VCPG. Let R be a VCPG of G . Choose l special free ends, which we do not consider in the following.

Within a face f that contains a free end of the grid-path p_v which represents v , the following step is taken⁴:

- If the f is not a 3-face, p_v is lengthened such that it ends on the grid-path of a non-neighbor of v .
- If the f is a 3-face, a grid-path is added on top of the edge in f that does not include p_v . Then p_v is lengthened so that it ends on the new grid-path.

This takes care of all but l of the free ends, and the resulting graph must be $(2, l)$ -tight. To obtain a $(2, 0)$ -tight supergraph, the l remaining free ends are removed as well. The graph that is represented by the resulting VCPG is a tight graph that has G as a subgraph. \square

4.2.1 B_1 -VCPGs

In a B_1 -VCPG each vertex is represented by a grid-path with at most one bend. Kobourov, Ueckerdt and Verbeek have shown that every $(2, 3)$ -tight graph admits a strict B_1 -VCPG, in which every vertex has precisely one bend. In this section we will give a sketch of their algorithm. We will also show that strict B_1 -VCPG is a proper subset of B_1 -VCPG and that not all planar $(2, 2)$ -tight graphs admit a B_1 -VCPG.

⁴This is the same as in the proof of Lemma 4.1.

4.2.2 Constructing a strict B_1 -VCPG

In this section we will describe the algorithm of Kobourov, Ueckerdt and Verbeek [KUV13]. However, we do not simply duplicate their setup, but convert the algorithm into the notions that have appeared in this chapter, i.e., we will show that the flow and 2-orientation that Kobourov, Ueckerdt and Verbeek build, is a realizable pair.

Let $G = (V, E)$ be a planar (2,3)-tight graph. By Theorem 1.4, G has a construction with Henneberg type 1 and Henneberg type 2 steps, starting from one edge. Moreover, if such a construction exists and the graph is planar, then there is a construction such that in each step the graph is planar. Indeed, start with a planar embedding of G , and reverse the steps until there is only one edge left, to obtain a sequence of steps such that the graph is plane in every step. Instead of starting with an edge, we will start with a triangle, which is an edge together with one Henneberg type 1 step. A (2,3)-tight graph has at least two 3-faces, suppose not, then:

$$4|E| - 4|V| + 8 = 4|F| \leq \sum_i i \cdot f_i = 2|E|$$

which implies $|E| \leq 2|V| - 4$. A (2,3)-tight graph has one edge more, hence, it has at least two 3-faces. One of the 3-faces is chosen as the outer face, and the Henneberg construction is such that the outer face is not changed. Let G_i denote the graph after the i -th Henneberg step, G_0 is the triangle s_1, s_2, s_3 .

The second observation is that for every (2,3)-tight planar graph, there is a flow that satisfies the demands and excesses of the faces, routes through every vertex, except for two of the boundary vertices, and the flow is such that the graph induced by the paths of the flow is a tree (in [KUV13] this is denoted by *angular tree*). This flow can be build along the Henneberg construction, while maintaining the following invariant.

Invariant: The flow ψ is a feasible flow for the G_i , which routes through every vertex, except for s_1, s_2 , precisely once. The underlying paths of ψ form a tree. And for every m , every interior m -face has 1 unit of incoming flow and $m - 3$ units of outgoing flow in ψ (see Figure 4.16).




Figure 4.16: A feasible flow, splitting to obtain a 2-orientation and a proper, strict B_1 -VCPG.

Let ψ_0 be the feasible flow in G_1 , such that, ψ_0 starts in the interior face, goes through s_3 , and enters the outer face. Hence, for ψ_0 the invariant holds. The flow is updated along the construction steps. Let ψ_i , the flow after step i , be such that the invariant holds.

Every face has an outgoing edge and, hence, belongs to the underlying paths of ψ_{i+1} . The flow is routed through every vertex, except s_1 and s_2 , precisely once. Hence, the number of

edges vf in the underlying paths, e_ψ , is given by

$$e_\psi = 2|VG_{i+1}| - 2 = |EG_{i+1}| + 1 = |FG_{i+1}| + |VG_{i+1}| - 1.$$

As the number of vertices in the paths is equal to $|FG_{i+1}| + |VG_{i+1}|$, it suffices to show that the graph induced by the paths of the flow remains connected at each construction step. Then, it follows that the underlying paths of ψ_{i+1} form a tree.

Suppose the next step is a Henneberg type 1 step in a face f of size k (see Figure 4.17 (a)). The k -face f is splitted into an l -face f_u and a $(k-l+4)$ -face f_v . The face f has 1 unit of outgoing flow and $k-3$ units of incoming flow in ψ . Without loss of generality, let the unit of incoming flow be routed through a vertex that belongs to f_u after the step. To obtain ψ_{i+1} , the flow in f is divided to f_u and f_v in such a way that, f_u has one unit of outgoing flow and $l-4$ units of incoming flow, it follows that f_v has $k-3-l+4 = k-l+1$ units of incoming flow. This is always possible, as the flow is routed through each vertex precisely once. Then one unit of flow is routed from f_v through the new vertex to f_u . This ensures that the graph induced by the paths of the flow, remains connected. It follows from the construction, that ψ_{i+1} routes through every vertex, except for two of the boundary vertices, precisely once, as ψ_i does before the step and ψ_{i+1} routes through the only new vertex. Moreover, $k-l+1 = |f_v| - 3$, and as there is one unit of flow leaving f_v routed through the new vertex into f_u , the flow ψ_{i+1} is feasible and for every m , every interior m -face has 1 unit of incoming flow and $m-3$ units of outgoing flow in ψ . Hence, the invariant holds.

Suppose the next step is a Henneberg type 2 step, that subdivides the edge uv and connects the new vertex x to w (see Figure 4.17 (b,c)). The k -face f is splitted into an l -face f_u and a $(k-l+3)$ -face f_v , the face f' gets one extra vertex since uv is subdivided. The k -face has 1 unit of outgoing flow and $k-3$ units of incoming flow in ψ . Without loss of generality, let the unit of incoming flow be routed through a vertex that belongs to f_u after the step. To obtain ψ_{i+1} , the flow in f is divided to f_u and f_v in such a way that, f_u has one unit of outgoing flow and $l-3$ units of incoming flow, it follows that f_v has $k-3-l+3 = k-l$ units of incoming flow. If, in ψ_i there is a flow $f' \rightarrow v \rightarrow f$, then this unit of flow is removed and replaced by $f_v \rightarrow v \rightarrow f'$ and $f' \rightarrow x \rightarrow f_u$ (Figure 4.17 (b)). Otherwise, one unit of flow is routed from f_v through the new vertex to f' (Figure 4.17 (c)). It follows that the graph induced by the paths of the flow remains connected. The flow ψ_{i+1} routes through every vertex, except for two of the boundary vertices, precisely once, as ψ_i does before the step and ψ_{i+1} routes through the only new vertex. Moreover, $k-l = |f_v| - 3$, and as there is one unit of flow leaving f_v routed through the new vertex into f' , the flow ψ_{i+1} is feasible and for every m , every interior m -face has 1 unit of incoming flow and $m-3$ units of outgoing flow in ψ . Hence, the invariant holds.




Figure 4.17: Extending the flow along the Henneberg steps.

At each step i , ψ_i is a feasible flow. We will deduce a 2-orientation of G , using the feasible flow, such that, the pair is realizable. Start by expanding every vertex, except for the two vertices on the boundary, through which no flow is routed.

Claim. The expanded graph, with only the edges of G , is a tree.

This follows from the fact that the underlying paths of the flow, form a tree. Therefore the expanded graph, with only the edges of G , is connected. Moreover, this graph has $2|V| - 2$ vertices (two vertices for every expanded vertex) and $2|V| - 3$ edges (all edges of G). Therefore, the graph must be a tree. \triangle

The edges of the expanded graph $G' \setminus s_1s_2$ are oriented towards the two special vertices, s_1 and s_2 . The edge s_1s_2 is oriented arbitrarily. Let this orientation be denoted by α . It is trivial that α is a 2-orientation on G . Every vertex that is splitted has precisely one outgoing edge on either side, all but the special vertices are splitted. The two special vertices have outdegree 0 and 1. Hence, α is a 2-orientation.

Theorem 4.18. *The pair (α, ψ) , as defined above, is a realizable pair.*

Proof. From construction it follows that ψ is a feasible flow and α is a 2-orientation. Since there is at most one unit of flow routed through each vertex, there is a decomposition of the flow into non-crossing paths. Moreover, the 2-orientation is obtained from expanding the vertices, therefore, the flow cuts the path between the outgoing neighbors of a vertex. Hence, the pair is realizable. \square

To obtain a strict B_1 -VCPG, the grid-paths of the two special vertices get a bend in the outer face, both the convex and the concave angle of these bends are in the outer face. As a special feature, the representation that comes from this realizable pair has a nice property. Every interior face has precisely one unit of outgoing flow, therefore, it has precisely one convex angle of a vertex. The authors denote this property by ‘proper’ and they show that plane Laman graphs are precisely the graphs that admit a proper, strict B_1 -VCPG.

4.2.3 The class B_1 -VCPG

There are graphs that are not (2,3)-sparse which admit a B_1 -VCPG, for example K_4 . In this section we will give some insight into the class of graphs that admit a B_1 -VCPG. As mentioned before, it is known that plane (2,3)-tight graphs are precisely the graphs that admit a proper, strict B_1 -VCPG [KUV13]. It follows that planar (2,3)-sparse graphs admit a strict B_1 -VCPG. The complete graph on four vertices, K_4 , shows that proper, strict B_1 -VCPG is a proper subclass of strict B_1 -VCPG. We will show that strict B_1 -VCPG is a proper subclass of B_1 -VCPG, i.e., there exists graphs that can be represented when segments are allowed. We will also show that not all planar (2,2)-tight graphs admit a B_1 -VCPG.

Lemma 4.19. *Strict B_1 -VCPG is a proper subclass of B_1 -VCPG.*

Proof. The graph in Figure 4.18 is a graph that admits a B_1 -VCPG but no strict B_1 -VCPG. A B_1 -VCPG is given. Note that to satisfy the excess of all 3-faces, the six boundary vertices in total have six bends. If x had a bend, then there would be a unit of flow leaving one of the two 4-faces incident to x . This implies that at least one of the boundary vertices would get one more bend, and thus, at least one of the boundary vertices will have two bends. Therefore, x has to be represented as a segment in a B_1 -VCPG of this graph. \square

Lemma 4.20. *Not every planar (2,2)-tight graph admits a B_1 -VCPG.*



Figure 4.18: A graph that has a B_1 -VCPG in which the vertex x must be represented by a segment.

Proof. Suppose a graph has a B_1 -VCPG, then, by Theorem 4.9, there must be a realizable pair (α, ψ) such that $0 \leq \psi(v) \leq 1$ for all vertices v . Consider the plane graph on the right of Figure 4.19. Suppose it has a flow ψ which satisfies $0 \leq \psi(v) \leq 1$ for all vertices v . The two grey-colored K_4 subgraphs both have excess $c(K_4) = 3$. So there must be 6 units of flow going out of the two K_4 subgraphs, however, there are only 5 different vertices bounding the two K_4 subgraphs. We conclude that there is no feasible flow such that there is at most one unit of flow going through each vertex. \square





Figure 4.19: Two embeddings of a planar $(2,2)$ -tight graph that is not B_1 -VCPG.

The vertex m is incident to three K_4 subgraphs, K_4 is a critical set of a $(2,2)$ -tight graph. Recall that a critical set of a $(2,l)$ -sparse graph is a set of vertices H which induces precisely $2|H| - l$ edges. Graphs that do not have a vertex that is incident to many (partly disjoint) critical sets, might admit a B_1 -VCPG.

A Laman-plus-one graph G is a $(2,2)$ -tight graph, such that there exists an edge e in G for which $G - e$ is $(2,3)$ -tight, i.e. a Laman graph [HOR⁺05]. Generic circuits (Section 2.2.3 on page 51) are a subclass of Laman-plus-one graphs. It is a proper subclass, as shown for example by a Henneberg type 1 extension of K_4 , which is Laman-plus-one but not a generic circuit. It is not difficult to extend the result of Kobourov, Ueckerdt and Verbeek from planar $(2,3)$ -tight graphs to planar Laman-plus-one graphs.

Theorem 4.21. *Every planar Laman-plus-one graph has a strict B_1 -VCPG. In this representation precisely one face has two convex corners, the outer face has no convex corner, all other faces have one convex corner.*

Proof. Let G be a planar Laman-plus-one graph and $e = ab$ an edge such that $G - e$ is a Laman graph. Consider an embedding of G such that e is incident to the outer face. We construct an extended graph which has a triangular outer face (see Figure 4.20 (a)). Three new vertices are added, x, y and z .


 Figure 4.20: Extension of $G - e$ and adding e .

1. Add x and the edges xa and xb .
2. Add y and the edges ya and yb in such a way that the outer face now consists of a, x, b and y .
3. Subdivide the edge by , call the new vertex z and add the edge xz such that the outer face now consists of x, b and z .

The addition of x and y are Henneberg type 1 steps and the addition of z is a Henneberg type 2 step. Therefore, the graph $G - e$ with the extension, denoted by $G' - e$ is a Laman graph. We construct a strict B_1 -VCPG of $G' - e$ according to the method of Kobourov, Ueckerdt and Verbeek [KUV13], in such a way that x and z are the special vertices. Hence, in this VCPG every face has precisely one convex corner and x and z have their bend in the outer face.

Consider the three interior faces that are incident to at least one of x, y and z . All three have one convex corner, which is not the corner of x or z . We want to show that a, b and y have their convex corner in one of these faces. Suppose the convex corner of b is not in one of the new faces. Consider the representation of $G - e$ that arises by removing x, y and z . There are precisely three free ends, two at b and one at a . This follows from the fact that x and z together have three free ends and xz must use the fourth of their ends, therefore, the edges of a and b are oriented such that $b \rightarrow z$, $b \rightarrow x$ and $a \rightarrow x$. The boundary of $G - e$ in the representation forms a closed cycle with three free ends. From this it follows that all other vertices on this boundary have their concave angle in the outer face. From this it follows that a, b and y have their convex corner in one of the new faces.

This implies also that the rest of the graph is drawn on one side of b , as shown in Figure 4.20 (b) and (c). The two cases are similar and therefore we only consider the first one, where the rest of the graph lives on the vertical part of the grid-path of b . Suppose a has its convex corner in $axyz$, then there cannot be a path from a to b that bounds $G - e$ (clockwise), see Figure 4.20(d). Therefore, the corners must be precisely as in Figure 4.20 (b).

It follows that y is the unique convex corner for the face $axyz$. The convex corner of a lies in the face with b, x, a and a path from a to b on the boundary. Removing x, y and z from the representation, leaves a strict B_1 -VCPG in which the leg of the free end of a has no vertices ending on it. Moreover, no vertex but one of the free ends of b is to the right of this leg of a . Hence, we move this leg in such a way that it ends on b , the result is a strict B_1 -VCPG of G . It follows from the construction of Kobourov, Ueckerdt and Verbeek that all faces except the two incident to the bend of a have precisely one convex corner, the face incident to the convex corner of a has also a convex corner from b and the outer face has no convex corners. \square

The graph in Figure 4.18 deleted the vertex x , is not a Laman-plus-one graph, however, it has a strict B_1 -VCPG. Hence, it is still open to characterize all graphs that have a strict B_1 -VCPG.

4.2.4 B_2 -VCPGs

For simple (2,2)-tight graphs we show that for every 2-orientation α , there exists a flow ψ such that $\psi(v) \leq 2$ for all vertices v and the pair (α, ψ) is realizable.

Theorem 4.22. *Every planar (2,2)-tight graph admits a B_2 -VCPG.*

Proof. Let G be a planar (2,2)-tight graph and \mathcal{E} a planar embedding of G . Hence we have a dual graph G^* . For simplicity, we denote the objects of the dual graph by their name in G . The excess of the faces in the dual graph is given by $c(f) = 4 - |f|$ and for the outer face it is $c(f_\infty) = -|f_\infty|$.

For every subset of faces H , there are at least $|\sum_{f \in H} c(f)|$ edges leaving H in the dual graph. Let b denote the number of edges leaving H , i.e., the number of boundary edges in the primal graph with respect to H . Let f_H , e_H and v_H be the number of vertices, edges and faces in H , respectively. Using Euler's formula (where we only count interior faces) and the upper bound on the edges in H we obtain the following.

$$\begin{aligned}
\sum_{f \in H} (4 - |f|) &= 4f_H - 2e_H + b & (4.4) \\
&= 4e_H - 4v_H + 4 - 2e_H + b & \text{(Euler's Formula)} \\
&= 2e_H - 4v_H + 4 + b \\
&\leq 4v_H - 2l - 4v_H + 4 + b & \text{(since } e_H \leq 2v_H - l) \\
&= b + 4 - 2l
\end{aligned}$$

The following equations shows that the number of edges in H is at least $2v_H - b$. Let B be the set of boundary vertices of H .

$$e_H = b + |E(G[v_H \setminus B])| \geq b + 2v - l - 2v_{V \setminus V_{H \cup B}} + l = 2v_{H \setminus B} + b = 2v_H - b$$

Therefore, we obtain the following:

$$\begin{aligned}
\sum_{f \in H} (|f| - 4) &= 2e_H - b + 4f_H & (4.5) \\
&= 2v_H + 2f_H - 2 - b - 4f_H & \text{(Euler's Formula)} \\
&= 2v_H - 2f_H - 2 - b \\
&\leq 2v_H - 2v_H + 2b + 2 - 2 - b & \text{(since } f_H \geq v_H - b - 1) \\
&= b
\end{aligned}$$

Putting this together we obtain:

$$\left| \sum_{f \in H} c(f) \right| = \left| \sum_{f \in H} (4 - |f|) \right| \leq b.$$

Hence there is a feasible flow in the dual graph which uses every edge at most once. Consider any orientation α of G such that every vertex has outdegree at most 2. In such a way that

there is either one vertex with outdegree 0 incident to the outer face, or two vertices with outdegree 1 incident to the outer face. We construct the flow ψ in the angle graph as follows: If there is a flow from f_1 to f_2 crossing edge uv , then if $u \rightarrow v$ we add $f_1 \rightarrow u$ and $u \rightarrow f_2$ to ψ . If $v \rightarrow u$ we add $f_1 \rightarrow v$ and $v \rightarrow f_2$ to ψ .

Since the flow in the dual graph is edge-disjoint and each vertex has at most two outgoing edges we have $\psi(v) \leq 2$ for all vertices v . At each vertex the flow cuts off the outgoing edge, hence the expanding condition is satisfied at each vertex.

We conclude that the pair (α, ψ) is realizable. \square

4.2.5 B_k -VCPGs for $k > 2$

For (2,2)-tight graphs, using a flow in the dual gives a tight bound. As we have seen before that there are planar (2,2)-tight graphs which need vertices with more than one bend, and using the dual graph we could show that two bends is sufficient. Unfortunately this is not the case for (2,1)- and (2,0)-tight graphs.

Theorem 4.23. *Every simple planar (2,1)-tight graph admits a B_4 -VCPG.*

Proof. We claim that there is a flow in the dual graph such that each edge has capacity two. Similarly to the proof of Theorem 4.22 we obtain that a simple planar (2,1)-tight graph admits a B_4 -VCPG.

For every subset of face-vertices H there are at least $|\sum_{f \in H} c(f)|$ edges leaving H in the dual graph. Let b be the number of edges leaving H , i.e., the number of boundary edges in the primal with respect to H . From Equations 4.4 and 4.5 we obtain:

$$\left| \sum_{f \in H} c(f) \right| = \left| \sum_{f \in H} (4 - |f|) \right| \leq b + 4 - 2l = b + 2.$$

There are no loops, therefore, $b > 1$. It follows that the demands can be satisfied using b edges with capacity 2. \square

We have only been able to show that the lower bound is two (e.g. the octahedron deleted an interior edge), hence, we conjecture that simple planar (2,1)-tight graphs admit a B_2 -VCPG.

Theorem 4.24. *Every simple planar (2,0)-tight graph admits a B_6 -VCPG.*

Proof. Let G be a simple planar (2,0)-tight graph and \mathcal{E} a planar embedding of G . Hence we have a dual graph G^* . The excess of the face-vertices in the dual is given by $c(f) = 4 - |f|$ and for the outer face it is $c(f_\infty) = -4 - |f_\infty|$.

For every subset of face-vertices H there are at least $|\sum_{f \in H} c(f)|$ edges leaving H in the dual graph. Let b be the number of edges leaving H , i.e. the number of boundary edges in the primal with respect to H . From Equations 4.4 and 4.5 we obtain:

$$\left| \sum_{f \in H} c(f) \right| = \left| \sum_{f \in H} (4 - |f|) \right| \leq 4 + b. \quad (4.6)$$

As there are no loops, we can satisfy all excess in the dual graph by using every edge at most thrice.

Similarly as before this induces a realizable pair, in which every vertex has at most 6 units of flow proceeding through. Hence, we find a B_6 -VCPG. \square

4.2.6 Obtaining Better Bounds

In the previous section we have used the dual graph to obtain a feasible flow. This way, we are certain that there exists a 2-orientation such that the pair is realizable. Even stronger, every 2-orientation can be selected. Unfortunately, this method does not always give an optimal solution. In Figure 4.21 two VCPGs of the octahedron are shown, the left of which does not relate to a feasible flow in the dual graph.




Figure 4.21: Two VCPGs of the octahedron that induce the same 2-orientation. The VCPG on the left does not relate to a flow in the dual graph as the rightmost bend on the bottom, of the grid-path of vertex b , cannot be transformed into a flow in the dual graph.

It has been fruitful to use the dual graph, in order to bound the number of bends locally at each vertex, but what if we want to globally minimize the number of bends. It would be ideal, if there always exists a 2-orientation such that the minimum feasible flow is realizable. However, for the graph in Figure 4.22, a minimum flow is given (the value of the flow is equal to the number of 3-faces). For this flow, there does not exist a 2-orientation such that the pair is realizable. Suppose there is such a 2-orientation. The realizability condition requires the orientations $a \rightarrow c$ and $b \rightarrow c$. Then c has only one outgoing edge, but the flow requires the free ends to be in the outer face, contradiction. Note that, rerouting the flow through a or c to go through b , leaves a flow for which there does exist a 2-orientation such that the pair is realizable.

The construction of a flow, as in the algorithm of Kobourov, Ueckerdt and Verbeek [KUV13], allows for constructing a 2-orientation, that is realizable together with this flow. Unfortunately, for $(2,0)$ -tight graphs, no simple construction is known, i.e., the construction may require to use loops and double edges. Representing a vertex with two loops, requires the grid-path to have five bends if the loops are nested (six bends are necessary for two loops that are not nested, however, this cannot be forced). Hence, using the non-simple construction, the bound on the number of bends cannot be guaranteed at any intermediate graph.

4.2.7 Local Flow Decreasing Steps

There exist some (trivial) flow decreasing steps. Given a graph G and a realizable pair (α, ψ) . If one of the steps is possible, we construct a realizable pair (α, ψ') such that $w(\psi') < w(\psi)$.




Figure 4.22: A minimum feasible flow for which there is no 2-orientation such that the pair is realizable.

Let uv be an edge between two faces f_1 and f_2 . If there is flow going from f_1 to f_2 as well as from f_2 to f_1 , then this cancels out (see Figure 4.23 (a)). This is a trivial flow decreasing step which leaves a realizable pair. The same holds for cyclic flow that consists of more than two units of flow.




Figure 4.23: Routes of flow that can be changed to decrease the number of bends.

Detour removal. Given a graph G and a realizable pair (α, ψ) . Given an edge (u, v) , the two faces adjacent to this edges, f_1 and f_2 , and a face f_u adjacent to u (see Figure 4.23 (b)). If there is flow in ψ going from $f_1 \rightarrow v \rightarrow f_2$ and from $f_2 \rightarrow u \rightarrow f_u$, then it can be replaced by $f_1 \rightarrow u \rightarrow f_u$. The orientation does not change, and since $f_1 \rightarrow v \rightarrow f_2 \in \psi$, the edge uv is oriented from v to u in α . Hence, the expansion condition is not violated by this change. Moreover, if the following flow is in ψ : $f_2 \rightarrow v \rightarrow f_1$ and $f_u \rightarrow u \rightarrow f_2$, similarly this flow can be replaced by $f_u \rightarrow u \rightarrow f_1$.

In general, flow decreasing steps might not give a realizable pair. There are examples such that given a minimum flow ψ in the angle graph that satisfies the facial demands, there is no 2-orientation which together with ψ is a realizable pair (see previous section). Given a realizable pair, there are examples for which the cycle and detour removals do not give the minimum flow.




Figure 4.24: A VCPG (a) of a graph G (b), induced by a flow that cannot be reduced by the given steps and a VCPG of G with the minimal number of bends (c).

The graph in Figure 4.24 (b) and the given realizable pair, can not be improved using the two local decreasing steps described above. The VCPG in Figure 4.24 (a) belongs to this realizable pair. The VCPG in Figure 4.24 (c), however, has less bends than the VCPG in (a). This shows that the minimum cannot always be found using the two local decreasing steps.

4.2.8 Constructive Argument for (2,2)-tight graphs

A constructive argument can be build upon the construction of simple (2,2)-tight graphs [Nix11, NOP12].




Figure 4.25: The four steps: Henneberg type 1, Henneberg type 2, edge-to- K_3 and vertex-to- K_4 . In the latter two the sets of neighbors are given by the grey colored triangles.

Construction of a B_2 -VCPG for (2,2)-tight planar graphs.

We will build an edge-disjoint flow in the dual graph. Then given any orientation, the flow of the dual graph can be mapped to a flow in the angle graph such that this flow together with the orientation is realizable.

Nixon has shown that every simple (2,2)-tight graph G is derivable from K_4 by the Henneberg type 1, Henneberg type 2, edge-to- K_3 and vertex-to- K_4 moves. Moreover a planar (2,2)-tight graph has such a construction in which all intermediate graphs are planar. We will build a flow in the dual graph along the plane construction of G .

We start with K_4 and the flow in the dual graph that leaves the three inner faces through the edges that are incident to the outer face (as given in Figure 4.26). Throughout the construction we maintain an edge-disjoint flow in the dual graph.

Extension along Henneberg type 1. Vertex x is added in face f and connected to two vertices of f say u and v . The face f is now splitted into two faces f_1, f_2 and the




Figure 4.26: K_4 and a flow in the dual graph, and the extension of the flow in the dual graph along the construction steps: Henneberg type 1 (b), Henneberg type 2 (c), edge-to- K_3 (d) and vertex-to- K_4 (e).

discrepancy is at most two, i.e., in the worst case f_1 now has excess 2 and f_2 has demand 2. To prove this we consider the dual graph, and in particular the vertex f and its neighbors. The demand of f is 0 before the Henneberg type 1 step. The degrees of f , f_1 and f_2 in the dual graph satisfy $\deg(f) = \deg(f_1) + \deg(f_2) - 4$. The demand of f is equal to the sum of the demands of f_1 and f_2 , so at most one of the faces has an excess, say f_1 . We also know that the current flow is edge-disjoint. Suppose the excess of f_1 is larger than 2, hence it has more than $\deg(f_1) - 2$ incoming edges. The two edges to f_2 have not been used in the current flow. Therefore f_1 has at most $\deg(f_1) - 2$ incoming edges and thus at most excess 2. Therefore at most two edges need to be directed from f_1 to f_2 . As there are two edges not yet used in the flow, that connect f_1 with f_2 , we can extend the flow in an edge-disjoint way (see Figure 4.26 (b)).

Extension along Henneberg type 2. Edge uv is subdivided and the new vertex x is connected to w . The face f is splitted in f_1 and f_2 , the face f_x on the other side of the edge uv has gotten one more vertex, x . Consider the faces in the dual graph and the flow, deleted the flow that uses the edge (f, f_x) . The excess of f_x is at least 0 and at most 2 and the sum of the demands $c(f_x) + c(f_u) + c(f_v)$ is zero.

Suppose $c(f_x) = 0$, then, in the worst case $c(f_1) = -c(f_2) = 2$ and we add flow $f_1 \rightarrow f_x, f_2$ and $f_x \rightarrow f_2$ (see Figure 4.26 (c)).

Suppose $c(f_x) = 1$, then, in the worst case $c(f_1) = 1, c(f_2) = -2$ and we add flow $f_1 \rightarrow f_2$ and $f_x \rightarrow f_2$.

Suppose $c(f_x) = 2$, then, $c(f_1) = c(f_2) = -1$ and we add flow $f_x \rightarrow f_1, f_2$.

Extension along edge-to- K_3 . The neighbors of a vertex x are splitted into three (possibly empty) sets N_1, N_2, N_3 , one of the three sets contains precisely one neighbor, say $N_3 = \{u\}$. The rotation system at x is maintained. Remove x and u , add a K_3 . Two vertices of the K_3 are connected with N_1 and N_2 respectively. The third vertex of the K_3 gets the neighbors of u . The extension of the flow is trivial, the negative demand of the face between the splitted neighbor-sets N_1 and N_2 is satisfied with flow coming from the new face of the K_3 (see Figure 4.26 (d)). Note that, if there was a unit of flow routed over the edge before the step, then now this flow is routed through the K_3 .

Extension along vertex-to- K_4 . The neighbors of a vertex x are splitted into three (possibly empty) sets, maintaining the rotation system at x . Remove x , add a K_4 such that the graph stays plane, and the three outermost vertices of the K_4 are connected with the respective neighbor-sets of x . The extension of the flow is trivial, the negative demand of the faces between the splitted neighbor-sets of x is satisfied with flow coming from the three interior faces of the introduced K_4 (see Figure 4.26 (e)).

Now we have obtained an edge-disjoint flow in the dual graph, such that the demand of all faces is satisfied. Similarly as before, given any 2-orientation α , this flow can be transformed to a flow ψ in the angle graph such that there is at most two units of flow going through a vertex. The pair (α, ψ) is realizable, this gives a B_2 -VCPG.

Extension to simple planar (2,1)-tight graphs

Nixon and Owen have also shown that simple planar (2,1)-tight graphs have a similar construction [NO], starting from one of the graphs in Figure 4.27. Both graphs admit a B_2 -VCPG. A (2,1)-tight graph is constructed from one of the starting graphs using Henneberg type 1, Henneberg type 2, edge-to- K_3 , vertex-to- K_4 and edge-joining moves. The *edge-joining move* connects two disjoint (2,1)-tight graphs with an edge. Unfortunately, even for the starting graphs, it is not possible to design a feasible flow in the dual graph that uses every edge at most once. When instead a feasible flow in the angle graph is constructed, there is no guarantee that within a face the flow is “nicely” distributed, i.e., a Henneberg type 1 step might induce two faces, one of which contains all vertices through which incoming flow is routed, the other contains none.




Figure 4.27: Two K_4 s glued at an edge and a B_2 -VCPG and K_5 deleted one edge and a B_2 -VCPG

Is there a way to use the construction steps to build a feasible flow? Can it be done such that every vertex has at most 2 units of flow routed through?

4.3 Conclusion

We have shown that there exist planar (2,2)-tight graphs that do not admit a B_1 -VCPG. However, the only type of (2,2)-tight, planar graph that we found not to have a B_1 -VCPG, has at least one vertex which is the intersection of ‘many’ critical subsets. Do all planar (2,2)-tight graphs that have no such vertex admit a B_1 -VCPG?

We have also obtained bounds for simple (2,1)-tight and (2,0)-tight planar graphs, however, we believe that these bounds are not tight. A lower bound of three bends for simple, planar (2,0)-tight graphs is given by the octahedron.

Conjecture 4.25. *Every simple planar $(2,0)$ -tight graph admits a B_3 -VCPG.*

Conjecture 4.26. *Every simple planar $(2,1)$ -tight graph admits a B_2 -VCPG.*

The bounds that we have shown do not depend on a chosen 2-orientation (i.e. the bounds hold for every 2-orientation). It would be interesting to find a sufficient condition on a flow such that, when satisfied, there exists a 2-orientation such that the pair is realizable. For $(2,3)$ -tight graphs, the algorithm of Kobourov, Ueckerdt and Verbeek, takes a particular flow and a particular network. Is there a way to construct a realizable pair simultaneously for all $(2,0)$ -sparse graphs?

“Weet je wat dat is, vergetelheid? Het is de grote prijs die aan het eind van je leven op je ligt te wachten. Nooit meer herinnerd worden aan je fouten, aan je gemiste kansen, de schaamte die in je verleden ligt te etteren en nooit meer zal genezen, al het ongewroken verraad: alles lost uiteindelijk op in de vergetelheid.”

Adriaan Jaeggi, De laatste duik van de dag - Nagelaten proza

5

List of Open Problems

In this chapter some open problems and the progress (to the best of my knowledge) are reported. By no means, do I try to write an comprehensive list. One could say that this is a list of problems that I have encountered and hope(d) to solve at some point.

Straight-Line Triangle Representations

Question 1. Is the recognition of graphs that have an SLTR (GFAA) in P ?

Question 2. Does a feasible solution of the network in Section 2.3 (page 54) imply the existence of an integral feasible solution? For example, there is a theorem of Seymour about binary hypergraphs [Sey77]. From this it follows that a 2-commodity flow problem in the undirected case, is totally dual integral unless it has a certain subgraph. Perhaps the set of paths that comes from a feasible solution gives rise to a binary hypergraph, and we can use the result of Seymour.

Question 3. This is posed as a conjecture by Gonçalves, Lévêque and Pinlou [GLP12]. Does every 3-connected planar graph have a primal-dual contact representation by right-angled triangles? Can the harmonic system of equations be used for this, or is there a way to write a system of equations that always has a nondegenerate solution yielding a contact representation by right-angled triangles?

Question 4. We can obtain an SLTR with integer coordinates. Is there a way to bound the size of the grid necessary? The solution of the system of equations will give the coordinates of the vertices in a SLTR. Is there a way to select the parameters and locations of the suspensions such that all coordinates are integers and the size of the grid is small?

Touching Triangle Graphs

Question 5. Characterize the biconnected outerplanar graphs that admit a k TTR.

Question 6. Is it possible to characterize biconnected internally cubic planar graphs that admit a cTTR? Chang and Yen have a polynomial decision algorithm that exploits the structure of the auxiliary graph of biconnected internally cubic graphs [CY]. They also identify a strict obstruction, which is based on the number of assignments needed “inside” some area. To be more specific: if a certain subgraph needs all but two of its boundary vertices to be assigned inside this subgraph, there is no good assignment. Does this give a counting argument that characterizes the biconnected internally cubic planar graphs that admit a cTTR?

Question 7. Characterize the 2-outerplanar graphs that admit a cTTR. Halin graphs are 2-outerplanar graphs. All Halin graphs admit a 3TTR, but 2-outerplanar graphs do not necessarily have a nice structure like Halin graphs. Can the result for biconnected outerplanar graphs be used to extract a result on 2-outerplanar graphs?

Question 8. Are k TTRs area-universal? A representation is area-universal if any prescription of the areas can be realized.

Vertex Contact representations of Paths on a Grid

Question 9. Characterize the class B_1 -VCPG. We have shown that not all planar (2,2)-tight graphs admit a B_1 -VCPG. The examples that we have found, all have at least one

vertex that belongs to many critical sets that are almost disjoint. Suppose the critical sets in a graph can be ordered by inclusion in such a way that either a critical set is contained in the other or the two sets are disjoint. Does such a graph always admit a B_1 -VCPG? Can such a graph be splitted into Laman-plus-one graphs in such a way that the representations of the subgraphs can be glued together without introducing more bends?

Question 10. Does every simple planar $(2,0)$ -tight graph admit a B_3 -VCPG?

In this thesis we have shown an upper bound of 6. We believe that using the orthogonal drawing of Fößmeier, Kant and Kaufmann [FKK96], it is possible to obtain an upper bound of 4. First the drawing is obtained, then the vertices are replaced by boxes. A vertex is identified by its outgoing edges, the ends of its neighbors have to be adjusted to ensure proper contacts. Then the number of bends should be reduced, as there will be many bends with Z-shapes.




Figure 5.1: The orthogonal drawing that is obtained by the algorithm of Fößmeier, Kant and Kaufmann.

However, it seems unlikely that this idea will reduce the number of bends to three bends per grid-path. Three bends per grid-path would be best possible, as shown by the octahedron.

Question 11. Does every simple planar $(2,1)$ -tight graph admit a B_2 -VCPG? So far we have only been able to show an upper bound of 4. Can the construction for simple $(2,1)$ -tight graphs as given by Nixon and Owen [NO] be of use?

Question 12. Is there a way to construct a realizable pair simultaneously for all $(2,0)$ -sparse graphs? The construction of Kobourov, Ueckerdt and Verbeek [KUV13] constructs the flow in such a way that a fitting 2-orientation exists. Even more, the fitting 2-orientation can be obtained from the flow. Is there a similar way to construct a flow, in the case where the flow does not necessarily live on a tree?

Question 13. Is there is construction for simple planar $(2,0)$ -tight graphs? To the best of our knowledge there is no construction known. In the Ph. D. thesis of Nixon a start has been made [Nix11].

Question 14. Is every (k,l) -sparse graph a subgraph of a (k,l) -tight graph with the same number of vertices? This is not true if we require the graph to be simple and planar, as $K_5 - e$ is a $(2,0)$ -sparse graph, but there does not exist a simple, planar $(2,0)$ -tight graph on five vertices.

The answer is yes if we ask for a planar $(2,0)$ -tight graph, not necessarily simple, such that a simple planar $(2,0)$ -sparse graph is a subgraph. The free ends in a VCPG can be extended to end on some other grid-path, possibly itself, but loops and multiple edges are allowed. The answer is also yes if we ask for a simple planar $(3,6)$ -tight graph of a $(3,6)$ -sparse graph, as we can triangulate each face.

Question 15. Can every planar graph be represented as intersection representation of grid-paths with one bend? Or, even stronger, with grid-paths with one bend that are oriented the same way, e.g. as L's. This has been conjectured by Chaplick and Ueckerdt [CU13]. They also conjectured that every triangle-free planar graph admits a B_1 -VCPG, where each grid-path that has a bend, is oriented the same way. This is shown to be true by Francis [Fra].

Bibliography

- [ACG⁺12] A. ASINOWSKI, E. COHEN, M. C. GOLUMBIC, V. LIMOUZY, M. LIPSHTEYN, AND M. STERN, *Vertex Intersection Graphs of Paths on a Grid*, J. Graph Alg. and Appl., 16 (2012), pp. 129–150.
- [AEK⁺15] M. J. ALAM, D. EPPSTEIN, M. KAUFMANN, S. G. KOBOUROV, S. PUPYREV, A. SCHULZ, AND T. UECKERDT, *Contact representations of sparse planar graphs*, CoRR, abs/1501.00318 (2015).
- [Aer14] N. AERTS, *Biconnected outerplanar graphs that have a $kTTR$* . Appeared as a poster at GD2014, 2014.
- [AF] N. AERTS AND S. FELSNER, *Straight line triangle representations*. Submitted, extended version of GD2013.
- [AF13a] N. AERTS AND S. FELSNER, *Henneberg steps for triangle representations*, in Proc. EuroComb, J. Nešetřil and M. Pellegrini, eds., vol. 16 of CRM Series, Scuola Normale Superiore, 2013, pp. 503–509.
- [AF13b] N. AERTS AND S. FELSNER, *Straight line triangle representations*, in Proc. Graph Drawing, S. K. Wismath and A. Wolff, eds., vol. 8242 of Lecture Notes in Computer Science, Springer, 2013, pp. 119–130.
- [AF14] N. AERTS AND S. FELSNER, *Vertex Contact Graphs of Paths on a Grid*, in Proc. WG, D. Kratsch and I. Todinca, eds., vol. 8747 of Lecture Notes in Computer Science, Springer, 2014, pp. 56–68.
- [BBC11] M. BADENT, U. BRANDES, AND S. CORNELSEN, *More canonical ordering*, J. Graph Alg. and Appl., 15 (2011), pp. 97–126.
- [BFM07] N. BONICHON, S. FELSNER, AND M. MOSBAH, *Convex drawings of 3-connected plane graphs*, Algorithmica, 47 (2007), pp. 399–420.
- [BJ03] A. R. BERG AND T. JORDÁN, *A proof of Connelly’s conjecture on 3-connected circuits of the rigidity matroid*, J. Comb. Th. Ser. B, 88 (2003), pp. 77–97.
- [Bre00] E. BREHM, *3-orientations and Schnyder 3-tree-decompositions*, master’s thesis, Freie Universität Berlin, 2000.
- [CG09] J. CHALOPIN AND D. GONÇALVES, *Every planar graph is the intersection graph of segments in the plane: extended abstract*, in Proc. ACM Symp. Theor. of Comp., M. Mitzenmacher, ed., ACM, 2009, pp. 631–638.
- [CGO10] J. CHALOPIN, D. GONÇALVES, AND P. OCHEM, *Planar graphs have 1-string representations*, Discr. and Comput. Geom., 43 (2010), pp. 626–647.
- [CK97] M. CHROBAK AND G. KANT, *Convex grid drawings of 3-connected planar graphs*, Int. J. Comput. Geometry Appl., 7 (1997), pp. 211–223.
- [CKU98] 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 (1998).
- [CU13] S. CHAPLICK AND T. UECKERDT, *Planar graphs as VPG-graphs*, J. Graph Alg. and Appl., 17 (2013), pp. 475–494.

BIBLIOGRAPHY

- [CY] Y.-J. CHANG AND H.-C. YEN, *Touching triangle representations of biconnected internally cubic plane graphs*. Private Communication.
- [dCCD⁺02] N. DE CASTRO, F. J. COBOS, J. C. DANA, A. MÁRQUEZ, AND M. NOY, *Triangle-free planar graphs as segment intersection graphs*, J. Graph Alg. and Appl., 6 (2002), pp. 7–26.
- [dFdM01] H. DE FRAYSSEIX AND P. O. DE MENDEZ, *On topological aspects of orientations*, Discr. Math., 229 (2001), pp. 57–72.
- [dFdM03] H. DE FRAYSSEIX AND P. O. DE MENDEZ, *Stretching of Jordan arc contact systems*, in Proc. Graph Drawing, 2003, pp. 71–85.
- [dFdM04] H. DE FRAYSSEIX AND P. O. DE MENDEZ, *Contact and intersection representations*, in Proc. Graph Drawing, J. Pach, ed., vol. 3383 of Lecture Notes in Computer Science, Springer, 2004, pp. 217–227.
- [dFdM07a] H. DE FRAYSSEIX AND P. O. DE MENDEZ, *Barycentric systems and stretchability*, Discr. Appl. Math., 155 (2007), pp. 1079–1095.
- [dFdM07b] H. DE FRAYSSEIX AND P. O. DE MENDEZ, *Representations by contact and intersection of segments*, Algorithmica, 47 (2007), pp. 453–463.
- [dFdMP94] H. DE FRAYSSEIX, P. O. DE MENDEZ, AND J. PACH, *Representation of planar graphs by segments*, in Intuitive Geometry, K. Böröczky and G. F. Tóth, eds., North-Holland, Amsterdam, 1994, Coll. Math. Soc. J. Bolyai 63, pp. 109–117.
- [dFdMP95] H. DE FRAYSSEIX, P. O. DE MENDEZ, AND J. PACH, *A left-first search algorithm for planar graphs*, Discr. and Comput. Geom., 13 (1995), pp. 459–468.
- [dFdMR94] H. DE FRAYSSEIX, P. O. DE MENDEZ, AND P. ROSENSTIEHL, *On triangle contact graphs*, Comb., Probab. and Comput., 3 (1994), pp. 233–246.
- [dFPP90] H. DE FRAYSSEIX, J. PACH, AND R. POLLACK, *How to draw a planar graph on a grid*, Combinatorica, 10 (1990), pp. 41–51.
- [EIS76] S. EVEN, A. ITAI, AND A. SHAMIR, *On the complexity of timetable and multicommodity flow problems*, SIAM J. Comput., 5 (1976), pp. 691–703.
- [Fel01] S. FELSNER, *Convex drawings of planar graphs and the order dimension of 3-polytopes*, Order, 18 (2001), pp. 19–37.
- [Fel03] S. FELSNER, *Geodesic embeddings and planar graphs*, Order, 20 (2003), pp. 135–150.
- [Fel04] S. FELSNER, *Geometric Graphs and Arrangements*, Advanced Lectures in Mathematics, Vieweg, 2004. Some Chapters from Combinatorial Geometry.
- [Fel13] S. FELSNER, *Rectangle and square representations of planar graphs*, in Thirty Essays on Geometric Graph Theory, J. Pach, ed., Springer, 2013, pp. 213–248.
- [FFNO11] S. FELSNER, É. FUSY, M. NOY, AND D. ORDEN, *Bijections for baxter families and related objects*, J. Comb. Th. Ser. A, 118 (2011), pp. 993–1020.
- [FHKO10] S. FELSNER, C. HUEMER, S. KAPPES, AND D. ORDEN, *Binary labelings for plane quadrangulations and their relatives*, Discr. Math. and Theor. Comp. Sci., 12 (2010), pp. 115–138.
- [FKK96] U. FÖSSMEIER, G. KANT, AND M. KAUFMANN, *2-visibility drawings of planar graphs*, in Proc. Graph Drawing, S. C. North, ed., vol. 1190 of Lecture Notes in Computer Science, Springer, 1996, pp. 155–168.
- [Fow13] J. J. FOWLER, *Strongly-connected outerplanar graphs with proper touching triangle representations*, in Proc. Graph Drawing, S. K. Wismath and A. Wolff, eds., vol. 8242 of Lecture Notes in Computer Science, Springer, 2013, pp. 156–161.
- [Fra] M. C. FRANCIS. Private Communication.

- [FZ06] S. FELSNER AND F. ZICKFELD, *Schnyder woods and orthogonal surfaces*, in Proc. Graph Drawing, 2006, pp. 417–429.
- [FZ08] S. FELSNER AND F. ZICKFELD, *On the number of planar orientations with prescribed degrees*, Electr. J. Combin., 15 (2008).
- [GHK10] E. R. GANSNER, Y. HU, AND S. G. KOBOUROV, *On touching triangle graphs*, in Proc. Graph Drawing, 2010, pp. 250–261.
- [GLP12] D. GONÇALVES, B. LÉVÊQUE, AND A. PINLOU, *Triangle contact representations and duality*, Discr. and Comput. Geom., 48 (2012), pp. 239–254.
- [GLS09] M. C. GOLUMBIC, M. LIPSHTEYN, AND M. STERN, *Edge intersection graphs of single bend paths on a grid*, Networks, 54 (2009), pp. 130–138.
- [Hen11] L. HENNEBERG, *Die graphische Statik der starren Systeme*, Johnson Reprint 1968, Leipzig 1911.
- [HNZ91] I. B.-A. HARTMAN, I. NEWMAN, AND R. ZIV, *On grid intersection graphs*, Discr. Math., 87 (1991), pp. 41–52.
- [HOR⁺05] R. HAAS, D. ORDEN, G. ROTE, F. SANTOS, B. SERVATIUS, H. SERVATIUS, D. L. SOUVAINE, I. STREINU, AND W. WHITELEY, *Planar minimally rigid graphs and pseudo-triangulations*, Comp. Geom.: Theory and Appl., 31 (2005), pp. 31–61.
- [Kan96] G. KANT, *Drawing planar graphs using the canonical ordering*, Algorithmica, 16 (1996), pp. 4–32.
- [KMBW02] E. KRUIJA, J. MARKS, A. BLAIR, AND R. C. WATERS, *A short note on the history of graph drawing*, in Proc. Graph Drawing, P. Mutzel, M. Jünger, and S. Leipert, eds., vol. 2265 of Lecture Notes in Computer Science, 2002, pp. 272–286.
- [Koe36] P. KOEBE, *Kontaktprobleme der konformen Abbildung*, Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl., 88 (1936), pp. 141–164.
- [KUV13] S. G. KOBOUROV, T. UECKERDT, AND K. VERBEEK, *Combinatorial and geometric properties of planar Laman graphs*, in Proc. ACM-SIAM Symp. Discr. Algo., 2013, pp. 1668–1678.
- [Lam70] G. LAMAN, *On graphs and rigidity of plane skeletal structures*, Journal of Engineering Mathematics, 4 (1970), pp. 331–340.
- [Lov09] L. LOVÁSZ, *Geometric representations of graphs*, Draft version December 11, 2009. <http://www.cs.elte.hu/~lovasz/geomrep.pdf>.
- [Mil02] E. MILLER, *Planar graphs as minimal resolutions of trivariate monomial ideals*, Docum. Math., 7 (2002), pp. 43–90.
- [Nix11] A. NIXON, *Rigidity on Surfaces*, PhD thesis, Lancaster University, Lancaster, 2011.
- [NO] A. NIXON AND J. C. OWEN, *An inductive construction of (2,1)-tight graphs*. <http://arxiv.org/abs/1103.2967v2>.
- [NOP12] A. NIXON, J. C. OWEN, AND S. C. POWER, *Rigidity of frameworks supported on surfaces*, SIAM J. Discr. Math., 26 (2012), pp. 1733–1757.
- [NR04] T. NISHIZEKI AND M. S. RAHMAN, *Planar graph drawing*, vol. 12 of Lecture notes series on computing, World Scientific Publishing Co. Pte. Ltd., 2004.
- [NW64] C. S. J. A. NASH-WILLIAMS, *Decomposition of finite graphs into forests*, J. London Math. Soc., 39 (1964).
- [RT86] P. ROSENSTIEHL AND R. E. TARJAN, *Rectilinear planar layouts and bipolar orientations of planar graphs*, DCG, 1 (1986), pp. 343–353.
- [Sch89] W. SCHNYDER, *Planar graphs and poset dimension*, Order, 5 (1989), pp. 323–343.

BIBLIOGRAPHY


- [Sch90] W. SCHNYDER, *Embedding planar graphs on the grid*, in Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms, Proc. ACM-SIAM Symp. Discr. Algo., Philadelphia, PA, USA, 1990, Society for Industrial and Applied Mathematics, pp. 138–148.
- [Sch95] M. W. SCHÄFFTER, *Drawing graphs on rectangular grids*, Discr. Appl. Math., 63 (1995), pp. 75–89.
- [Sey77] P. D. SEYMOUR, *The matroids with the max-flow min-cut property*, J. Comb. Th. Ser. B, 23 (1977), pp. 189–222.
- [Tam87] R. TAMASSIA, *On embedding a graph in the grid with the minimum number of bends*, SIAM J. Comput., 16 (1987), pp. 421–444.
- [Tar85] E. TARDOS, *A strongly polynomial minimum cost circulation algorithm*, Combinatorica, 5 (1985).
- [Tut63] W. T. TUTTE, *How to draw a graph*, Proc. of the London Math. Soc., 13 (1963), pp. 743–767.
- [Uec13] T. UECKERDT, *Geometric representations of graphs with low polygonal complexity*, PhD thesis, Technische Universität Berlin, 2013.
- [Zic07] F. ZICKFELD, *Geometric and combinatorial structures on graphs*, PhD thesis, Technische Universität Berlin, 2007.

Index

- (k, l) -Sparse and (k, l) -Tight Graphs, 4
- B_k -VCPG, 97
- α -orientation, 3
- Almost 4-regular Graphs, 50
- Angle Graph, 103
- Auxiliary Graph, 76
- Bag-Vertices, 109
- Belonging to a component, 77
- Bipolar Orientation, 7
- Boundary Path Coloring, 90
- Canonical Order, 2
- cFAA, 72
- Chord-to-Endpoint Assignment, 69
- Combinatorial Convex Corner, 21
- Cone, 13
- Contact Family of Pseudo-Segments, 22
- Contact Representation, 8
- Corner Compatibility, 30
- Corner Compatible Pair, 30
- Critical Set, 4
- Dividing Segment, 32
- Dual Vein, 79
- Elbow Geodesic, 13
- Extremal Point, 29
- FAA, 20
- Face Counting, a drawing obtained by, 13
- Filter, 13
- Flats, of a geodesic embedding, 13
- Free End of a Grid-Path, 99
- Free Point, 23
- Generalized Schnyder Wood, 10
- Generic Circuit, 51
- Geodesic Embedding, 14
- Geometric Convex Corners, 21
- GFAA, 22
- Grid Intersection Graphs, 97
- Half-Edge, 10
- Halin Graph, 73
- Henneberg Steps, 5
- Internally k -connected, 2
- Internally 3-connected, 19
- Intersection Representation, 8
- Laman Graph, 5
- Medial Graph, 48
- Minimally Rigid Generic Framework, 4
- One-Sided-Perfect Matching, 20, 46
- Orthogonal Arc, 13
- Orthogonal Graph Drawing, 98
- Orthogonal Surface, 13
- Outline Cycle, 21
- Parallel Components, 89
- Path-Edges, 109
- Path-Vertices, 109
- Primal-Dual Contact Representation, 47
- Quadrangulation, 6
- Realizability Condition, 106
- Realizable Pair, 105
- Reversed Peeling Order, 68
- Ridge, 13
- Rigid Orthogonal Surface, 15
- Schnyder Labeling, 10
- Schnyder Wood, 9
- Separating (Subdivided) Triangle, 32
- Separating Decomposition, 6
- Simple Outline Cycle, 28
- SLTR, 18
- Spanning Tree Decompositions, 5
- Stretchable Contact Family of Pseudosegments, 29
- Strongly Connected Outerplanar Graphs, 69
- Suspended Graph, 10
- Suspension Vertex, 18
- Suspension Vertices, 10
- Touching Triangle Representation, 66
- Valid Orientation of $\text{VENATION}(G)$, 78
- VCPG, 97
- Venation Graph, 76
- Vertex Reduction, 19


Samenvatting

In een klassieke afbeelding van een graaf, worden de knopen weergegeven als punten in het vlak, en de relatie tussen twee punten door middel van een verbindingslijn. Een graaf is planair, wanneer hij zo getekend kan worden, dat geen twee lijnen elkaar kruisen. Een voorbeeld van een graaf is een metronetwerk, de stations zijn de knopen, twee stations zijn verbonden wanneer er een metro van het ene station naar het andere rijdt. Een andere manier om grafen weer te geven, is door middel van intersectie-representatie. In dit geval worden de knopen weergegeven als geometrische objecten, zoals lijnen of cirkels, en twee knopen zijn verbonden precies dan wanneer de twee objecten elkaar snijden. Als de objecten elkaar slechts raken, spreken we van een contact-representatie. In dit proefschrift komen drie representaties aan de orde.




Alleereerst wordt onderzocht welke grafen een zogenaamde Straight-Line Triangle Representation (SLTR) hebben. Dit is een representatie van een planaire graaf, waarbij ieder omsloten vlak een driehoek vormt. We geven twee karakterisering van grafen die een SLTR hebben. Echter, weten we niet hoe we efficient kunnen controleren of een graaf aan de voorwaarden voldoet. Met behulp van de eerste karakterisering, hebben we een stelling van de Fraysseix en Ossona de Mendez op een nieuwe manier bewezen. We geven ook een nieuw bewijs van een stelling van Gonçalves, Lévêque and Pinlou. De tweede karakterisatie laat zich coderen als een maximaal-stroom probleem in een netwerk, echter, benodigen we twee soorten stroom en het is bekend dat dit probleem NP-compleet is. Een positief resultaat laat zien dat SLTRs met Henneberg type II stappen uitgebreid kunnen worden. Hieruit volgt dat grafen die geconstrueerd kunnen worden met deze stap, een SLTR hebben.

Nauw verbonden met SLTRs zijn zogenaamde Touching Triangle Representations (TTRs), dit is een contact-representatie waarbij de knopen als driehoeken worden weergegeven. Twee knopen zijn verbonden, precies dan, wanneer de driehoeken een (deel van een) zijde delen.




Hierbij, mogen er geen ruimtes omsloten worden die geen knoop representeren. Wanneer de vereniging van alle driehoeken een convexe veelhoek is, spreken we van een TTR in een veelhoek. We geven een karakterisatie van de outerplanaire grafen die zo een representatie hebben. Tevens bewijzen we dat iedere Halin graaf een TTR in een driehoek heeft.



Daarna beschouwen we een representatie waarbij de knopen door roosterpaden worden weergegeven. Twee roosterpaden zijn verbonden als het ene roosterpad loodrecht op het andere roosterpad eindigt. Twee roosterpaden mogen elkaar alleen op deze manier raken. Het is eenvoudig te bepalen welke grafen op deze manier getekend kunnen worden, dat zijn alle grafen voor welke iedere subgraaf hoogstens twee keer zoveel lijnen als knopen heeft. Deze grafen worden (2,0)-verspreid genoemd. We beschrijven een methode waarmee we het aantal buigpunten van de roosterpaden minimaliseren. Met gebruik van deze methoden hebben we bovengrenzen bewezen voor het aantal buigpunten in sommige graafklassen, deze graafklassen zijn subklassen van (2,0)-verspreid grafen.

Zusammenfassung



Klassisch werden Graphen durch Zeichnungen in der Ebene repräsentiert, das heißt die Knoten sind Punkte in der Ebene und die Kanten sind Linien, die adjazente Knoten verbinden. Eine Alternative ist, Knoten durch geometrische Objekte der Ebene wie Linien oder Kreise zu repräsentieren, welche sich genau dann schneiden wenn die korrespondierenden Knoten adjazent sind. In diesem Fall sprechen wir von Schnitt-Repräsentationen. Wenn sich die geometrischen Objekte nur berühren, sprechen wir von Kontakt-Repräsentationen.



Eine Straight-Line Triangle Representation (SLTR) ist eine klassische Zeichnung, bei der die Linien Segmente sind und jede Facette ein Dreieck ist. In dieser Arbeit untersuchen wir welche planaren Graphen ein SLTR besitzen. Wir geben dazu zwei Charakterisierungen an,

wobei offen ist, ob sich die ergebnen Bedingungen effizient testen lassen. Mit Hilfe der ersten Charakterisierung können wir eine Satz von de Fraysseix und Ossona de Mendez auf eine neue Weise und einfacher beweisen. Auch geben wir einen einfacheren Beweis eines Theorems von Gonçalves, Lévêque and Pinlou. Die zweite Charakterisierung lässt sich mit einem 2-Güterflussproblem formulieren. Leider ist bekannt, dass das Lösen von 2-Güterflussproblemen NP-schwer ist. Im Positiven haben wir bewiesen, dass SLTRs mit Henneberg-Typ-II-Schritten erweitert werden können. Ein Graph mit n Knoten ist eine generische Schaltung, falls er $2n - 2$ Kanten besitzt und jeder Subgraph mit m Knoten höchstens $2m - 3$ Kanten hat. Es ist bekannt, dass jede generische Schaltung mit Henneberg-Typ-II-Schritten konstruiert werden kann. Daraus folgt, dass jede planare generische Schaltung ein SLTR besitzt.

Eng verwandt mit SLTRs sind Dreieck-Kontakt Darstellungen, das heißt jeder Knoten wird durch ein Dreieck repräsentiert und zwei Dreiecke haben Seitenkontakt genau dann, wenn die Knoten verbunden sind. Dabei darf es keine Lücken in der Repräsentation geben. Wir geben eine Charakterisierung von zweifach zusammenhängende außerplanaren Graphen, die eine solche Repräsentation in einem konvexen Vieleck haben. Die Bedingungen der Charakterisierung können einfach getestet werden. Zweitens haben wir bewiesen, dass jeder Halin-Graph eine Kontakt-Repräsentation mit Dreiecken in einem Dreieck besitzt.



Danach betrachten wir eine Repräsentationsart, bei der die Knoten durch achsenparallele Gitter-Pfade repräsentiert werden. Zwei Knoten sind genau dann adjazent, wenn der Gitter-Pfad des einen Knoten orthogonal auf dem Gitter-Pfad des anderen Knoten endet. Es

ist einfach zu charakterisieren, welche Graphen so repräsentiert werden können: Es handelt sich um genau jene Graphen, bei denen jeder Subgraph höchstens zweimal soviel Kanten wie Knoten hat. Diese Graphen werden $(2,0)$ -karg genannt. Wir geben eine Methode an, mit der die Anzahl der Knicke der Gitter-Pfade minimiert werden kann. Das Minimum wird durch Reduktion auf ein Flussproblem berechnet. Damit konnten wir für einige Graphenklassen, die Subklassen von $(2,0)$ -kargen Graphen sind, obere Schranken für die Knickanzahl beweisen.

Curriculum Vitae

Nieke Aerts was born on January 24 in Riel, The Netherlands. She studied technical mathematics at Eindhoven University of Technology. She obtained her bachelor's degree in 2010, on the topic of secure communication schemes. The project was supervised by Henk van Tilborg. Following the bachelor's degree she studied discrete mathematics and applications and graduated in August 2011 on the topic of hash functions. The master's project was carried out at Macquarie university, Sydney, under supervision of Josef Pieprzyk and Benne de Weger.

In October 2011 Nieke started as a Ph.D. student at the Technical University Berlin. The research that resulted in this thesis was performed at the Discrete Mathematics group at the Technical University Berlin, under supervision of Stefan Felsner.