Geometric Representations of Graphs with low Polygonal Complexity

vorgelegt von
Diplom-Mathematiker
Torsten Ueckerdt
aus Berlin

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. Jan Kratochvil
	Prof. Dr. Stephen G. Kobourov

Tag der wissenschaftlichen Aussprache: 4. November 2011

Berlin 2012

D 83

Preface

This thesis is the fruit of my time in the Discrete Mathematics Group at TU Berlin. This group is a particularly great culture medium for mathematical plants of all kinds. Someone throws in a seed in form of a problem, a thought or an idea, and everybody can watch it either flourishing or wilting. Right from the beginning I was pleased about the possibility to work and learn within the Discrete Mathematics Group, and today I can take some nice flowers along.

I want to thank my advisor Stefan Felsner for always letting me choose my problems, for joining their treatment with well-trained patience and valuable ideas, and for his support of any kind. I am also thankful to Jan Kratochvil and Stephen G. Kobourov for being reviewers of my thesis.

During the years, I had the opportunity to meet and work with a lot of smart and kind researchers, in particular Marie Albenque, Daniel Heldt, Andrea Hoffkamp, and Kolja Knauer, as well as, Thomas Hixon and Irina Mustata here in Berlin, Bartłomiej Bosek, Tomasz Krawczyk, Piotr Micek, and Bartosz Walczak from Jagiellonian University in Kraków, Poland, and Stephen G. Kobourov and Muhammad Jawaherul Alam from University of Arizona.

I would like to emphasize my thanks to Daniel Heldt and Kolja Knauer for our joint research for Chapter 4, but also for listening to my thoughts and discussing even the wild ones. Furthermore, I do thank Stephen G. Kobourov and Muhammad Jawaherul Alam: During their stay in Berlin we did the research for Chapter 3 in a relaxed and productive atmosphere, which I very much enjoyed.

Last and not least, I very much want to thank my wife, *für alle bestehenden und bestandenen Abenteuer*.

Thank you all!

Torsten Ueckerdt
Berlin, August 2011

Contents

Introduction

There are several ways to think of a graph and many of them involve drawing pictures. In the most classical visualization vertices are considered as points in the plane and edges as continuous curves connecting two points, such as in the top-left of Figure 1. Indeed, graph properties of eminent importance, e.g., planarity, are defined with respect to those *drawings*.

Other popular graph visualizations include *intersection representations*. For example, every vertex is depicted as a point set in the plane and an edge between two vertices is described by an intersection of the corresponding point sets, such as in the bottom-left of Figure 1. In a *contact representation* the point set for each vertex is compact and those sets are pairwise interior disjoint. Then intersections involve only boundaries, as in the right of Figure 1, and are thus called contacts.

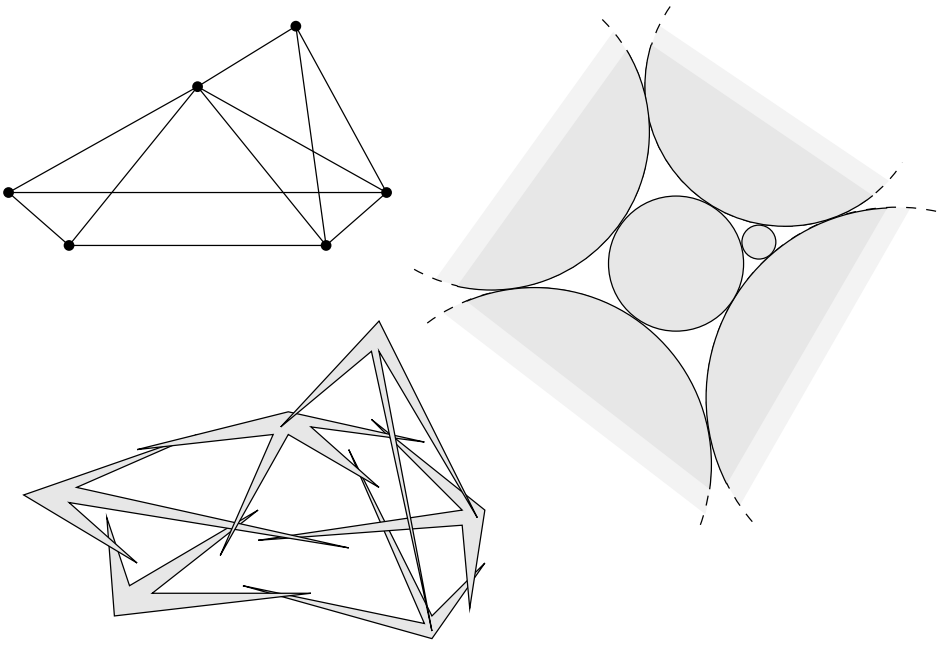

Figure 1: A drawing, an intersection, and a contact representation of a graph.

Many kinds of intersection graphs have been considered, ranging back from Koebe's "Kissing Coins Theorem" [Koe36] in 1936, up to segment representations of planar graphs due to Chapolin and Gonçalves in 2009 [CG09], and further.

Within this thesis we investigate two types of intersection graphs, in both of which vertices are represented by polygonal objects in the plane. We measure the complexity of a polygonal object by the number of its corners. We are then particularly interested in a low polygonal complexity for every vertex, i.e., we want the maximum complexity over all vertices to be as low as possible. Chapter 2 deals with *side contact representations* with one simple polygon for every vertex. Every graph that admits such a representation is necessarily planar. The major part of Chapter 2 concerns hole-free rectilinear representations, i.e., those in which every side of every polygon is either horizontal or vertical, and where the union of all polygons does not leave any holes. The right of Figure 2 shows such a contact representation. We consider here maximally planar graphs only, which is a natural (and almost necessary) assumption in this setting. One of our results is a new proof that polygons of complexity 8 are always sufficient and sometimes necessary for a hole-free rectilinear representation of a maximally planar graph.

In Chapter 3 we investigate what happens if we additionally prescribe the area of each and every polygon in the representation. A representation that respects a set of desired areas is known as a *cartogram*. For example, we prove that one can require any set of areas without increasing the worst-case maximum complexity of a hole-free rectilinear representation, i.e., 8-gons are still sufficient for every cartogram of a maximally planar graph.

The second type of intersection graph is investigated in Chapter 4. In an *EPG representation* vertices are represented as polygonal paths with solely horizontal and vertical segments, and an edge occurs whenever two paths overlap along some part of non-zero length, i.e., neither a touching point nor a crossing causes an edge. An example of such a representation is provided in the left of Figure 2. This time, every graph admits an EPG representation. However, we again want the polygonal complexity, i.e., the number of corners, per path to be low. The least possible maximum complexity over all paths for a given graph is the *bend-number*. We give several new upper and lower bounds on the maximum bend-number for certain graph classes, such as, planar and outer-planar graphs, complete bipartite graphs, graphs of certain tree-width, maximum degree, or degeneracy.

This is how the thesis is organized.

Chapter 1: This chapter introduces the basic concepts and notation we use within this thesis. Section 1.1 is about vertex orderings and in particular building

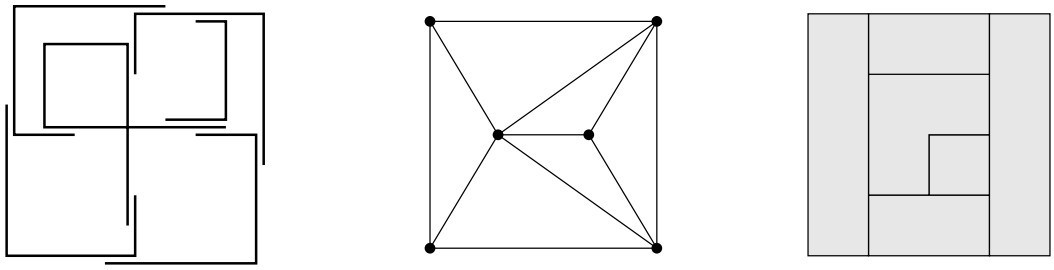

Figure 2: A planar graph with an EPG representation on the left and a side contact representation on the right.

sequences associated with them. Those building sequences underlie many of the constructive proofs presented in subsequent chapters. We define the degeneracy and the tree-width of a graph in terms of vertex orderings. For maximally planar graphs, we review the concepts of canonical orders and Schnyder woods, and outline some aspects of their close relation to each other. Furthermore, we define the separation-tree of an embedded maximally planar graph and deduce the level-i subgraphs from it. In Section 1.2 we consider orientations with prescribed out-degrees, so-called α-orientations, review their most important properties and present an algorithm that computes the minimal α-orientation in near-linear time. We close the preliminaries with Section 1.3, in which we briefly introduce rectangle-representations and transversal structures.

Chapter 2: In this chapter we investigate side contact representations of planar graphs, i.e., vertices are represented by simple polygons which are pairwise interior disjoint, and edges correspond to side contacts. In Section 2.1 we are particularly interested in rectilinear hole-free representations with low polygonal complexity. We present a general method to obtain such a representation for a maximally planar graph from a Schnyder wood. In special cases, we obtain a characterization of those maximally planar graphs that admit a non-rotated rectilinear representation with complexity 4, and 6. Furthermore, we derive a new compact floor plan for maximally planar graphs. In Section 2.2 we improve a result of Sun and Sarrafzadeh [SS93] by presenting a linear-time algorithm that constructs a rectilinear representation with complexity 6 based on a nesting assignment. Our algorithm can be adjusted to construct non-rectilinear representations of complexity 5, which in particular proves their existence under the presence of a nesting assignment. At the end of this chapter in Section 2.3, we provide a general method to compute lower bounds on the complexity of side contact representations. From this we derive matching lower bounds for all classes of planar graphs that we consider here. These bounds were known before, but our examples and argumentation are significantly simpler.

Chapter 3: We are interested here in side contact representations with an additional requirement, that is, we prescribe the area of the representing polygon for each and every vertex in the graph. Such representations are called cartograms. Section 3.1 introduces area-universal layouts, a key-concept in this field. We prove that every maximally planar graph admits an area-universal rectilinear hole-free layout of complexity at most 8, which is an immediate strengthening of the floor plan-result in Chapter 2. In Section 3.2 we present a different such layout with the same complexity for Hamiltonian maximally planar graphs, for which we can compute the actual cartogram in linear time, too. Based on this, we investigate one-sided Hamiltonian cycles, since they reduce the cartogram's complexity to 6. Afterwards, we extend the method for computing lower bounds on the polygonal complexity to cartograms in Section 3.3. We obtain better lower bounds and present a matching upper bound for the case of planar 3-trees. Finally, in Section 3.4 we discuss cartograms for 4-connected maximally planar graphs. This class leaves a lot and the most challenging open questions, for some of which we propose tailored approaches.

Chapter 4: This part of the thesis is not strongly related to the preceding chapters. We are interested in EPG representations, i.e., vertices are represented by polygonal paths in the plane square grid and edges correspond to paths that share a grid edge. Every graph has an EPG representation, but each graph is classified by the required maximum complexity of the paths involved – its bend-number. In Section 4.1 we present lower and upper bounds on the bend-number for complete bipartite graphs. Section 4.2 is concerned with outer-planar and planar graphs. We give a worst-case optimal upper bound in the former and new lower and upper bounds in the latter case. Section 4.3 relates the bend-number of a graph to its degeneracy, tree-width and maximum degree. Again we provide lower and upper bounds in each case, some of which are matching or almost matching. Furthermore, we provide the first NP-completeness result in the field in Section 4.4, i.e., we prove that recognizing single-bend graphs is NP-complete. In the end, we briefly compare the bend-number to other graph parameters, in particular the interval-number, local track-number, and track-number. These connections seem to be worth further investigations.

Open Questions: We close the thesis with a list of selected open questions.

Chapter 1

Preliminaries

This chapter introduces some basic notation and objects that are used throughout the thesis.

Section 1.1: This section is about vertex orderings, which are a very general tool for graphs. Here, we are particularly interested in the use of a vertex ordering as a building sequence of the corresponding graph. Many results in the subsequent chapters are proven constructively along a certain such building sequence. We consider vertex orderings associated with the graph's degeneracy (Subsection 1.1.1) and tree-width (Subsection 1.1.2), as well as canonical orders (Subsection 1.1.3), the closely related Schnyder woods (Subsection 1.1.4), and so-called level-i subgraphs (Subsection 1.1.5).

Section 1.2: We briefly introduce α-orientations and mention some of their most important properties. In Subsection 1.2.1 we present an algorithm that computes an α-orientation with near-linear running time.

Section 1.3: We introduce $\square$-representations, which are also known as rectangular duals. A $\square$-representation of a near-triangulation G is a contact representations of G with axis-aligned rectangles. We as well review the concept of transversal structures, which are closely related to $\square$-representations.

1.1 Vertex Orderings

Let $G = (V, E)$ be some n-vertex graph, and assume the vertices of G are labeled by $v_1, \ldots, v_n$ according to some ordering. Formally, $(v_1, \ldots, v_n)$ is called a *vertex ordering* of G. Of course, for a fixed graph there are $n!$ different vertex orderings, and some

may be more suitable for some purposes than others. Vertex orderings play a central role in graph theory, as many important problems ask for a vertex ordering with certain properties, e.g., the Hamiltonian cycle and path problems, the bandwidth problem, several linear arrangement problems, the elimination degree sequence problem, and others. All these problems are known to be NP-complete [GJ79], but there are equally important vertex orderings that can be computed in polynomial, even linear time, for instance perfect elimination orderings, topological orders, canonical orders, and orderings corresponding to the graphs degeneracy.

For a fixed vertex ordering $(v_1, \ldots, v_n)$ of a n-vertex graph G, we denote the subgraph of G induced by $\{v_1, \ldots, v_i\}$ by G_i, for $i = 1, \ldots, n$. In particular, we have $G_i = G[v_1, \ldots, v_i]$ and $G_n = G$.

- A *building sequence* of an n-vertex graph G is the sequence $G_0 \subset G_1 \subset \cdots \subset G_n$ with respect to an underlying vertex ordering $(v_1, \ldots, v_n)$.

In a building sequence the vertices of the graph are added one at a time together with all their edges to those vertices with smaller index. An example of a building sequence is given in Figure 1.2 in Subsection 1.1.3.

1.1.1 Degeneracy

Definition 1.1.1. The *degeneracy* of a graph $G = (V, E)$, denoted by $d(G)$ is the minimum number k, such that there exist a vertex ordering $(v_1, \ldots, v_n)$, such that for every $i = 1, \ldots, n$ the degree of v_i in G_i is at most k.

The degeneracy was introduced by Erdős and Hajnal [EH66] in 1966. It is not difficult to see that $d(G)$ equals the largest minimum degree of all subgraphs of G. For instance, the degeneracy of a planar graph is at most 5. Figure 1.1 **a)** shows a graph G with a vertex ordering, such that $\deg_{G_i}(v_i) \leq 2$, i.e., every vertex has at most two neighbors with a smaller label, and hence $d(G) \leq 2$. Since the graph has minimum degree 2, we conclude $d(G) = 2$. Note that for a given graph $G = (V, E)$ a vertex ordering with $\deg_{G_i}(v_i) \leq d(G)$ for every $i = 1, \ldots, n$ can be computed in $\mathcal{O}(|V| + |E|)$. To this end, identify a vertex of minimum degree, assign the highest available label to it, remove it from the graph, and iterate. In Theorem 4.3.1 in Section 4.3 we build up the graph along this vertex ordering.

Let us remark, that the concept of degeneracy is also known under the name *coloring number* [KMv+09]. To be precise, $\mathrm{col}(G) = d(G) + 1$, since every graph G can be greedily vertex-colored with $d(G) + 1$ colors using the corresponding vertex ordering.

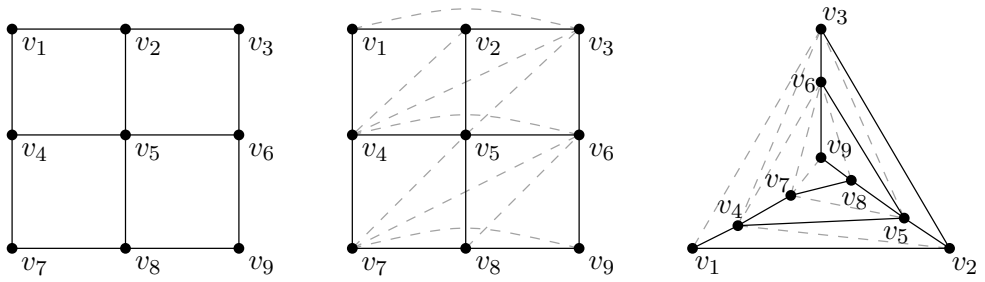

Figure 1.1: **a)** A graph G with a vertex ordering, which shows $d(G) \leq 2$. **b)** The graph G is a subgraph of a 3-tree, which shows $tw(G) \leq 3$. **c)** Another embedding of the graph in **b)**, which shows that G is a subgraph of a planar 3-tree.

1.1.2 Tree-Width

The tree-width was first introduced (under a different name) by Halin in 1976 [Hal76] and independently by Robertson and Seymour in 1986 [RS86]. (However, there are even earlier references [Wag37, HP68].) Tree-width and tree-decompositions play a very central role in graph minor theory, and are intimately related to planar graphs. For further reading we refer to the book of Diestel [Die10]. Within this thesis, we consider tree-width with superficial attention only.

For a number $k \geq 1$, a *k-tree* G with n vertices is either a complete graph with $k+1 = n$ vertices, or it is obtained from a k-tree G' with $n-1$ vertices by adding a vertex v to G' and k edges joining v and all vertices of a k-clique in G'. Note that 1-trees are exactly trees. The recursive definition directly implies that any k-tree has a vertex ordering $(v_1, \ldots, v_n)$, such that the neighborhood of v_i in G_i is a clique of size k. In particular, the degeneracy of a k-tree is at most k. (Since K_{k+1} is a subgraph of every k-tree, its degeneracy is exactly k.)

Definition 1.1.2. The *tree-width* of a graph G, denoted by $tw(G)$, is the minimum k such that G is a subgraph of a k-tree.

Given G with $tw(G) = k$, we usually denote by $\tilde{G}$ a k-tree that is a supergraph of G. For example, Figure 1.1 **b)** shows that the graph G in Figure 1.1 **a)** is a subgraph of a 3-tree $\tilde{G}$ drawn dashed and straight, i.e., $tw(G) \leq 3$. (Indeed, since G contains a K_4-minor, its tree-width is exactly 3 [APC90].) By a result of El-Mallah and Colbourn [EMC90], planar graphs of tree-width at most 3 are exactly the subgraphs of those 3-trees that are planar. Planar 3-trees are specific maximally planar graphs, sometimes called *stacked triangulations* [AC84], *graphs of stack 3-polytopes* [She74] or *Apollonian networks* [AHAdS05], which are defined a follows. The complete graph

on four vertices K_4 is a planar 3-tree, and any graph on at least five vertices is a planar 3-tree if and only if there is a vertex of degree 3 in G, whose removal leaves a planar 3-tree. Figure 1.1 **c)** shows another embedding of the graph in Figure 1.1 **b)** illustrating that it is a subgraph of a planar 3-tree.

In general, we have $tw(G) \geq d(G)$, but both numbers can be far apart. For instance, the planar $n \times n$ grid seen as a graph on n^2 vertices has tree-width n and degeneracy 2 [Hal76]. We remark, that many hard problems can be solved in polynomial, often even linear, time in case of small tree-width, e.g., the Hamiltonian cycle problem, the maximum independent set problem, or the 3-coloring problem. Indeed, a graph of constant tree-width can be tested in linear time for every property that can be defined in so-called monadic second-order logic [Cou90]. Unfortunately, testing whether the tree-width of a given graph is at most some number k is NP-complete [ACP87]. However, the recognition of graphs with tree-width at most k can be done in polynomial time *for fixed k (!)* [RS86, ACP87].

Let us remark that an equivalent definition for the tree-width says that $tw(G) + 1$ equals the minimum size of a largest clique among all chordal supergraphs of G, where a *chordal graph* is one without induced cycles of length four or more.

Within this thesis, we sometimes assume that the graph G of interest has (additionally) small tree-width, which enables us to derive certain representations for G. There maybe no such representation for larger tree-width, or we simply fail to find it. Anyways, these proofs rely on the building sequence for G, which is given from the recursive definition of a k-tree $\tilde{G}$ that is a supergraph of G. In such a building sequence, we maintain control over the neighborhood of vertex v_i in G_i, i.e., it is a clique in $\tilde{G}$, although not necessarily in G.

Considering planar graphs, we often analyze the case of tree-width 2 (c.f. Theorem 4.2.1 and Lemma 4.3.5) and tree-width 3 (c.f. Theorem 3.3.4, Theorem 4.2.2, and Lemma 4.2.3). Most importantly, outer-planar graphs have tree-width at most 2, and maximally outer-planar graphs are 2-trees. Hence we get a building sequence here, in which every new vertex is connected to at most two vertices in the already constructed graph.

1.1.3 Canonical Orders

Canonical orders were first introduced by de Fraysseix, Pach, and Pollack [dFPP90] for maximally planar graphs, and later generalized by Kant [Kan92, Kan96] to triconnected plane graphs. A further generalization are *orderly spanning trees* [CLL05]. Canonical orders have been proven to be a very valuable tool for many problems

about planar graphs, such as straight line drawings [dFPP90, Kan96, BFM07], compact graph representations [LLY03, CLL05], graph encoding [CGH+98], graph sampling [PS03], and many others. Within this thesis, we use canonical orders only for maximally planar graphs and hence define them only for this case. Recall that for a given vertex ordering $(v_1, \ldots, v_n)$ of a graph G, the subgraph of G induced by $\{v_1, \ldots, v_i\}$ is denoted by G_i, for $i = 1, \ldots, n$.

Definition 1.1.3. Let $G = (V, E)$ be an embedded maximally planar n-vertex graph with outer vertices u, v, w in counterclockwise order. A vertex ordering $(v_1, \ldots, v_n)$ is called a *canonical order* of G if $v_1 = u$, $v_2 = v$ and $v_n = w$, and the following conditions are met for every $4 \leq i \leq n$.

- The subgraph G_{i-1} of G is bi-connected, and the boundary of its outer face is a cycle C_{i-1} containing the edge (v_1, v_2).
- The vertex v_i lies in the outer face of G_{i-1}, and its neighbors in G_{i-1} form an (at least 2-element) subpath of the v_1-to-v_2 path $C_{i-1} \setminus (v_1, v_2)$.

Every maximally planar graph admits a canonical order and it can be computed in linear time [dFPP90]. A canonical order can be easily used to construct a straight line drawing of G without crossings. Let us consider the building sequence $G_3 \subset G_4 \subset \cdots \subset G_n$ of G. Putting $i = 4$ in Definition 1.1.3, it follows that $G_{i-1} = G_3$ is the triangle $\{v_1, v_2, v_3\}$. We embed G_3 as an acute triangle with the edge (v_1, v_2) drawn horizontally (and straight). For $i \geq 4$, the vertex v_i is connected to the vertices of a subpath P_i (of length at least 2) of the v_1-to-v_2 path $C_{i-1} \setminus (v_1, v_2)$ in G_{i-1}. Let $out_1(v_i)$ and $out_2(v_i)$ denote the start-vertex and end-vertex of P_i, respectively. For convenience, we define $out_1(v_3) = v_1$ and $out_2(v_3) = v_2$. It is easy to see that there is a position for the vertex v_i in the plane, to the right of $out_1(v_i)$, to the left of $out_2(v_i)$ and above all vertices in P_i, such that the resulting straight line embedding remains planar. See Figure 1.2 for an example. However, for better readability the embedding in Figure 1.2 is slightly stretched in some steps and some edges incident to v_{11} and v_{12} are not drawn straight.

Consider the graph G_i for $i = 3, \ldots, n$. We distinguish three kinds of vertices on the outer cycle C_i of G_i. A vertex $v \in C_i$ is called a *hill vertex* if all its neighbors in G_i have smaller y-coordinate. A *valley vertex* is one that has to the left and to the right a neighbor with larger y-coordinate. All other vertices on C_i are neither hill nor valley vertices. We give a more formal definition of hill and valley vertices in Lemma 1.1.7 in Subsection 1.1.4, from which follows that these terms depend only on the canonical order and not on the particular straight line embedding. However, for intuition one may think of the embedded G_i as having mountain shape with the hill and valley vertices being the peaks and valleys, respectively.

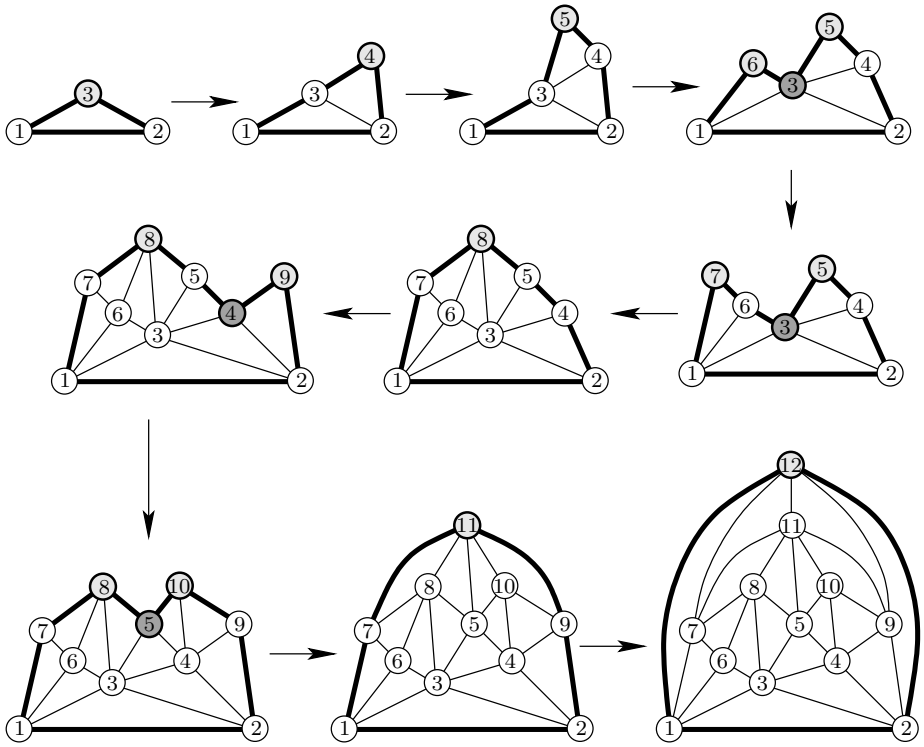

Figure 1.2: A maximally planar graph is built up using a canonical order. Hill and valley vertices in each G_i are highlighted in light and dark grey, respectively. The outer face cycle C_i is drawn bold.

1.1.4 Schnyder Woods

In 1989, Schnyder [Sch89, Sch90] introduced *Schnyder woods* and equivalent angle labellings, so-called *Schnyder labellings*, for maximally planar graphs. He proved a characterization of planar graphs in terms of the (order) dimension of the vertex-edge incidence order. Moreover, he used Schnyder woods to give the first proof that any n-vertex planar graph admits a straight line embedding on the $(n-2) \times (n-2)$ grid without crossings. Today it is known [dFOdM01], that Schnyder woods are in bijection with 3-orientations. (We define 3-orientations and explain this bijection in Section 1.2.) Based on this bijection it was shown [OdM94, Bre00], that the set of all Schnyder woods of a fixed plane graph carries the structure of a distributive lattice. Felsner [Fel01, Fel03] presents a natural way to generalize Schnyder woods to all tri-connected plane graphs.

For a comprehensive introduction to Schnyder woods and related objects we refer to the PhD thesis of É. Fusy [Fus07]. The following definition is taken from there.

Definition 1.1.4. Let $G = (V, E)$ be an embedded maximally planar n-vertex graph

with outer vertices a_1, a_2, a_3 in counterclockwise order. A *Schnyder wood* of G is an orientation and labeling of the inner edges of G with labels in $\{1,2,3\}$, satisfying the following rules.

- Each inner vertex v has exactly one outgoing edge of each label. The outgoing edges e_1, e_2, e_3 in each label $\{1,2,3\}$ occur in counterclockwise order around v. For $i \in \{1,2,3\}$, all edges entering v with label i are in the counterclockwise sector between e_{i+1} and e_{i-1}.
- For $i \in \{1,2,3\}$, all inner edges incident to a_i are ingoing and have label i.

It is convenient to associate the three colors blue, green, and red with the three labels 1, 2, and 3 in a Schnyder wood, respectively. Hence, every inner vertex v has exactly three outgoing edges in every Schnyder wood, one blue edge (labeled 1), one green edge (labeled 2), and one red edge (labeled 3), which appear in this counterclockwise order around v. We denote the end-vertex of the blue, green, and red edge by $out_1(v)$, $out_2(v)$, and $out_3(v)$, respectively. If v has incoming blue edges, then these appear in the sector at v between the green and the red outgoing edge not containing the blue outgoing edge. And analogous statements hold for the incoming green and red edges, respectively. We denote the set of neighbors of v that are connected by a blue edge to v, which is incoming at v, by $in_1(v)$. If there is no incoming blue edge at v, then $in_1(v) = \varnothing$. Similarly, $in_2(v)$ and $in_3(v)$ are defined w.r.t. green and red edges, respectively. Figure 1.3 **a)** illustrates the local rule at an inner vertex v and the notation with out_i and in_i, for $i = 1,2,3$. In Figure 1.3 **b)** the local rule at the outer vertices a_1, a_2, and a_3 is indicated.

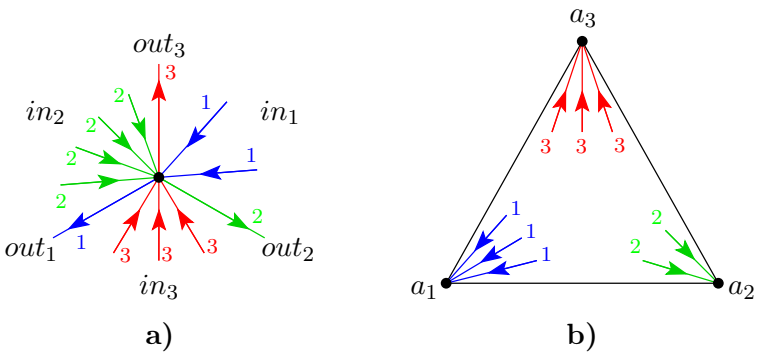

a) **b)**

Figure 1.3: **a)** The local rule of a Schnyder wood at an inner vertex. **b)** The local rule of a Schnyder wood at the outer vertices.

Schnyder [Sch89] proves that every maximally planar graph has a Schnyder wood. An example is shown in Figure 1.4 below. It is called a wood because for each $i \in \{1,2,3\}$ the set of edges with label i forms a directed tree spanning all inner vertices and a_i, with a_i being the root and every edge being directed towards the

root. Hence, the direction of the edges can be recovered from the labels. (This holds even the other way around, i.e., the labels of the edges can be recovered from their directions [dFOdM01].) For $i = 1, 2, 3$, we denote the tree rooted a_i by T_i, formally, $T_i = \{(v, out_i(v)) \mid v \in V \setminus \{a_1, a_2, a_3\}\}$. A Schnyder wood is then denoted by (T_1, T_2, T_3).

It turns out that the leaves of the three trees in a Schnyder wood play an important role. For $i = 1, 2, 3$, let us denote *set of leaves* of T_i by L_i. Sometimes, for example in Theorem 2.1.8 in Section 2.1, it is desirable to have a Schnyder wood in which one tree has few leaves.

Lemma 1.1.5 ([CLL05])**.** *In every Schnyder wood (T_1, T_2, T_3) of a n-vertex graph at least one tree T_i has at most $\lfloor \frac{2n-5}{3} \rfloor$ leaves, i.e., $|L_i| \leq \lfloor \frac{2n-5}{3} \rfloor$.*

Proof. The leaves of the Schnyder wood are in bijection with the acyclic inner faces of the graph, where the orientation of edges is given by the Schnyder wood and a vertex that is a leaf in more than one tree is counted with multiplicity. The bijection is the following. For an inner vertex v that is a leaf in T_i, consider the face Δ incident to v, $out_{i-1}(v)$, and $out_{i+1}(v)$. The vertex v is the unique source in Δ, i.e., the only vertex of which both incident edges in Δ are outgoing. On the other hand, if v is a source in some face Δ, then it has two outgoing edges that are consecutive in the circular order around v. Hence, there is no incoming edge in the corresponding sector and v is a leaf in one tree.

The statement now follows, since a maximally planar graph has precisely $2n - 5$ inner faces and there are three trees in a Schnyder wood. □

Let us remark that Lemma 1.1.5 is tight. Planar 3-trees are exactly those maximally planar graphs that have a unique Schnyder wood, and this implies that every inner face is acyclic. If the planar 3-tree is symmetric w.r.t. a_1, a_2, and a_3, then each tree in the Schnyder wood has the same number of leaves, namely exactly $\frac{2n-5}{3}$, where n is the number of vertices.

With a canonical order $(v_1, v_2, \ldots, v_n)$ we can associate a Schnyder wood in the following natural way. Let $a_1 = v_1$, $a_2 = v_2$, and $a_3 = v_n$. (Indeed, whenever we consider a canonical order and a Schnyder wood of one and the same (embedded) graph, we assume that $a_1 = v_1$, $a_2 = v_2$, and $a_3 = v_n$.) Note that a_1, a_2, a_3 appear in this counterclockwise order around the outer face. Recall that for $i \geq 3$ we have defined $out_1(v_i)$ and $out_2(v_i)$ to be the first and last vertex on $C_{i-1} \setminus (v_1, v_2)$ (when going from v_1 to v_2), that is a neighbor of v_i. Giving every inner edge (v_j, v_k) that is not labeled this way, the label 3 and orienting it from v_j to v_k if $j < k$, we have defined a labeling and orientation of the inner edges of the graph. It is not difficult to check,

that this actually is a Schnyder wood. Let us provide an example, which shows how to get a Schnyder wood from the canonical order in Figure 1.2. See Figure 1.4.

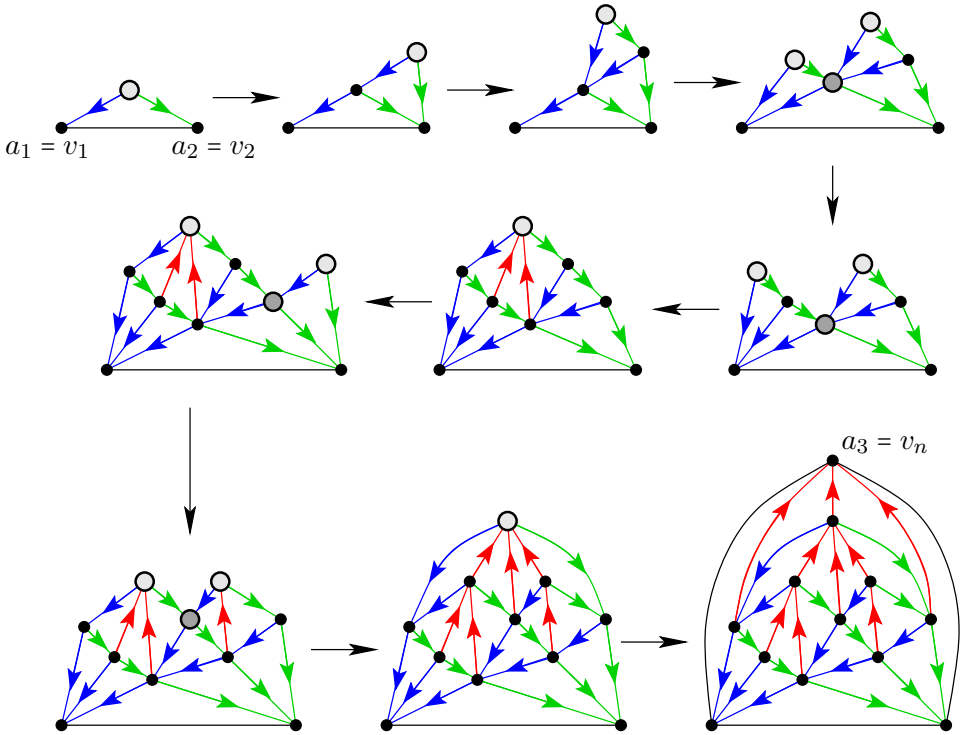

Figure 1.4: How to obtain a Schnyder wood from a canonical order. Hill and valley vertices are highlighted in light and dark grey, respectively.

Indeed, every Schnyder wood is obtained from a canonical order via the above procedure. In general, there are several different canonical orders that give the same Schnyder wood. It is known [Fel04], that reversing all edges in any two trees of a Schnyder wood results in an acyclic subgraph containing all inner edges. A *topological order* of an acyclic graph on n vertices is a vertex ordering $(v_1, \ldots, v_n)$, such that if there is a directed edge from v_i to v_j, then $i < j$.

Lemma 1.1.6. *Let G be an embedded maximally planar graph with outer vertices a_1, a_2, a_3 in counterclockwise order. Then each of the following holds.*

(a) A canonical order of G defines a Schnyder wood of G, where the outgoing edges of a vertex v are to its first and last neighbor with smaller label in the counterclockwise order around v, and to its neighbor with the highest label.

(b) For a Schnyder wood (T_1, T_2, T_3) every topological order of the acyclic graph $T_1^{-1} \cup T_2^{-1} \cup T_3$ defines a canonical order, where T_k^{-1} is the tree T_k with the direction of all its edges reversed.

To be precise, we insert the edge (a_1, a_2) directed from a_1 to a_2 into $T_1^{-1} \cup T_2^{-1} \cup T_3$, so that every topological order starts with a_1. Moreover, if $(v_1, \ldots, v_n)$ is the canonical order defined by the Schnyder wood (T_1, T_2, T_3), then the Schnyder wood that is defined by $(v_1, \ldots, v_n)$ is again (T_1, T_2, T_3). This way, we may associate to every Schnyder wood (T_1, T_2, T_3) the set of those canonical orders that define (T_1, T_2, T_3).

The following lemma can be taken as the formal definition of hill and valley vertices that was promised in the preceding subsection.

Lemma 1.1.7. *Let (T_1, T_2, T_3) be a Schnyder wood of a maximally planar graph G and $(v_1, v_2, \ldots, v_n)$ a topological order of $T_1^{-1} \cup T_2^{-1} \cup T_3$. Then for every $i = 3, \ldots, n-1$ each of the following holds.*

- *The edge-set $T_1 \cap E(G_i)$ and $T_2 \cap E(G_i)$ is a sub-tree of T_1 with root a_1 and of T_2 with root a_2, respectively.*
- *The v_1-to-v_2 path $C_i \setminus (v_1, v_2)$ along the outer face of G_i consists of an alternating sequence of paths in T_1 oriented towards v_1 and paths in T_2 oriented towards v_2, beginning with a T_1-path and ending with a T_2-path.*
- *The* hill vertices *in G_i are those vertices on C_i that are the start-vertex of a T_1-path and a T_2-path on C_i. A hill vertex is a leaf in $T_1 \cap E(G_i)$ and $T_2 \cap E(G_i)$.*
- *The* valley vertices *in G_i are those vertices on C_i that are the end-vertex of a T_1-path and a T_2-path on C_i. A valley vertex is neither a leaf in $T_1 \cap E(G_i)$ nor in $T_2 \cap E(G_i)$.*

1.1.5 Separation-Trees and Level-i Subgraphs

We introduce the separation-tree of a maximally planar graph with a fixed plane embedding. Based on this we define the level-i subgraph $G[i]$ of G, which enables us to build up G by iteratively inserting 4-connected pieces to the already constructed part of G.

Given a fixed embedding of a maximally planar graph G, a triangle Δ in G, i.e., a set $\{u, v, w\}$ of three pairwise adjacent vertices, is called *non-empty* if it is not an inner face in G. In particular, there is at least one further vertex inside the bounded region enclosed by Δ. The non-empty triangles are precisely the separating triangles and the outer triangle (if $|V(G)| \geq 4$), where a *separating triangle* is a set of three pairwise adjacent vertices that do not form a face in any embedding of G. We say that a triangle Δ_1 is *contained* in a triangle Δ_2, if the bounded region enclosed by Δ_1 is strictly contained in the one enclosed by Δ_2. For example, the outer triangle contains every triangle in the graph (except itself), and no triangle in G is contained in an inner facial triangle.

Definition 1.1.8. Let G be a maximally planar graph with a fixed plane embedding. The *separation-tree of* G is the rooted tree $\mathcal{T}_G$ whose vertices are the non-empty triangles in G, with Δ being a descendant of Δ' if and only if Δ is contained in Δ'.

The separation-tree has been considered for example in [SS93]. It is easy to see, that $\mathcal{T}_G$ is indeed a tree with the outer triangle as a root (provided $|V(G)| \geq 4$). For example, Figure 1.5 **c)** shows the separation-tree of the maximally planar graph in Figure 1.5 **a)**.

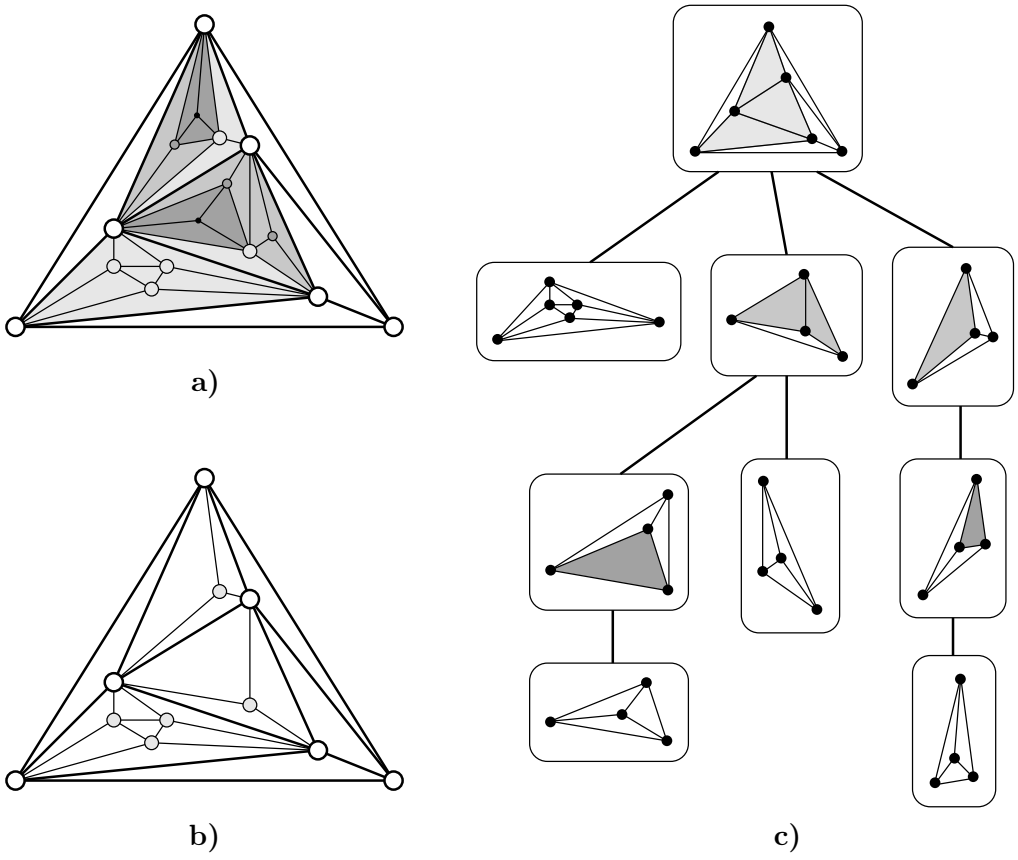

Figure 1.5: **a)** A maximally planar graph G with a fixed plane embedding. **b)** The level-1 subgraph $G[1]$ of G. **c)** The separation-tree $\mathcal{T}_G$. For each non-empty triangle Δ the graph G_Δ is depicted "inside" the vertex corresponding to Δ.

For a non-empty triangle $\Delta = \{u, v, w\}$, equivalently a vertex in $\mathcal{T}_G$, we define the graph G_Δ as follows.

- The vertex set of G_Δ consists of u, v and w, and every vertex x that lies inside Δ but not inside any triangle Δ', which is contained in Δ.
- The edge set of G_Δ consists of all the edges of G induced by the vertices of G_Δ.

It is not difficult to see that G_Δ is a maximally planar graph with outer triangle Δ. Moreover, G_Δ does not contain any separating triangles and hence G_Δ is either the complete graph on four vertices or it is the (unique) 4-connected maximally planar subgraph of G containing u, v, w, and at least one further vertex x that lies inside Δ. In Figure 1.5 **c)** the graph G_Δ for each non-empty triangle Δ is shown inside the corresponding vertex in the separation-tree.

The *depth* of a vertex in a rooted tree is its distance (measured by the number of edges) from the root. The depth of the rooted tree is the maximum depth of its vertices. We now define the *level* of a vertex $v \in V(G)$ based on the *depth* of a certain vertex $\Delta \in \mathcal{T}_G$.

Definition 1.1.9. Let $\mathcal{T}_G$ be the separation-tree of G, Δ_0 be its root, and d be its depth. The *level of a vertex* $v \in V(G)$ is the minimum depth of a non-empty triangle Δ in $\mathcal{T}_G$, such that $v \in V(G_\Delta)$.

For $i = 0, \dots, d$, the *level-i subgraph* of G is the subgraph $G[i]$ of G induced by all vertices of level at most i.

In other words, the level-i subgraph $G[i]$ of G is the union of all graphs G_Δ with depth of Δ in $\mathcal{T}_G$ at most i. For example, Figure 1.5 **b)** shows the level-1 subgraph $G[1]$ of the maximally planar graph G in Figure 1.5 **a)**. If Δ_0 denotes the outer triangle of G, then $G[0] = G_{\Delta_0}$. For example, the level-0 subgraph of G is depicted in the root of $\mathcal{T}_G$ in Figure 1.5 **c)**.

Sometimes, we prove a statement for a maximally planar graph G by an iterative procedure based on the level-i subgraphs of G. Suppose, we can show our statement for 4-connected maximally planar graphs. We may start with $G[0]$, and insert all graphs G_Δ with Δ being at depth 1 in the separation-tree one after another, while always maintaining the required statement. We have then proven the statement for $G[1]$, and iterating this procedure, we may end up with a proof for $G[d] = G$, where d is the depth of the separation-tree. Theorem 2.2.4 in Section 2.2 and Theorem 4.2.6 in Section 4.2 are proven in this fashion.

1.2 Orientations with Prescribed Out-Degrees

This section briefly introduces orientations with prescribed out-degrees, so-called α-orientations. These have been introduced and investigated by Felsner [Fel04] and independently by Ossona de Mendez [OdM94]. An *orientation* of an undirected graph G is a directed graph, whose underlying undirected graph is G. In other words, an orientation of G fixes a direction/orientation for each edge in G.

Definition 1.2.1. Given a graph $G = (V, E)$ and a mapping $\alpha : V \to \mathbb{N}$, an orientation of the graph's edges is called an *α-orientation* if out-deg$(v) = \alpha(v)$ for every vertex $v \in V$, i.e., the mapping α prescribes the out-degree at every vertex.

We restrict our attention to planar graphs here. There are several reasons for this, one being the following result, which is of significant importance. It was independently proved by Felsner [Fel04] and Ossona de Mendez [OdM94], and many more times for particularly special cases, as discussed further below.

Theorem 1.2.2 ([Fel04, OdM94])**.** *Let G be a plane graph and $\alpha : V(G) \to \mathbb{N}$ be a mapping. The set of α-orientations of G carries an order-relation which is a distributive lattice.*

Figure 1.6 shows an example of a plane graph and its set of five different α-orientations for a fixed mapping α. The orientations are depicted with the distributive lattice structure given by Theorem 1.2.2. We remark that different plane embeddings of G give rise to different distributive lattices for the α-orientations. However, we fix here any plane embedding of G.

Consider some α-orientation of G and a directed cycle in this, now directed, graph. Then reorienting every edge on the cycle results in a new orientation of G, which again is a (different) α-orientation. It is not difficult to see that every α-orientation of G can be transferred into any other by a sequence of cycle reversals. (This holds even in the non-planar case.) The *essential cycles* of G w.r.t. α are an inclusion-minimal set of cycles that is needed to get from every α-orientation to every other. The essential cycles can be chosen in such a way that the interiors of any two such cycles are either disjoint or contained in each other. In the latter case, the cycles are edge-disjoint. An edge is called *rigid w.r.t.* α if it has the same direction in every α-orientation of G. For example, both edges incident to the vertex in the bottom-right in Figure 1.6 are rigid. In case there are no rigid edges, we may choose the set of all inner faces as the essential cycles of G w.r.t. α. The essential cycles for the graph in Figure 1.6 w.r.t. the chosen α are the three inner facial cycles that are highlighted in grey in the top-most orientation.

The *minimal α-orientation* is defined to be the unique α-orientation of G in which no essential cycle is oriented counterclockwise. Indeed, it then follows that no cycle in G is counterclockwise. Moreover, an α-orientation *covers* another α-orientation in the distributive lattice if and only if the first arises from the second by reversing an essential cycle from clockwise to counterclockwise. Here covering means that the first α-orientation lies above the second in the distributive lattice and that there is no third, which lies above the second and below the first.

- A reversal of an essential cycle from clockwise to counterclockwise is called a *flip*.

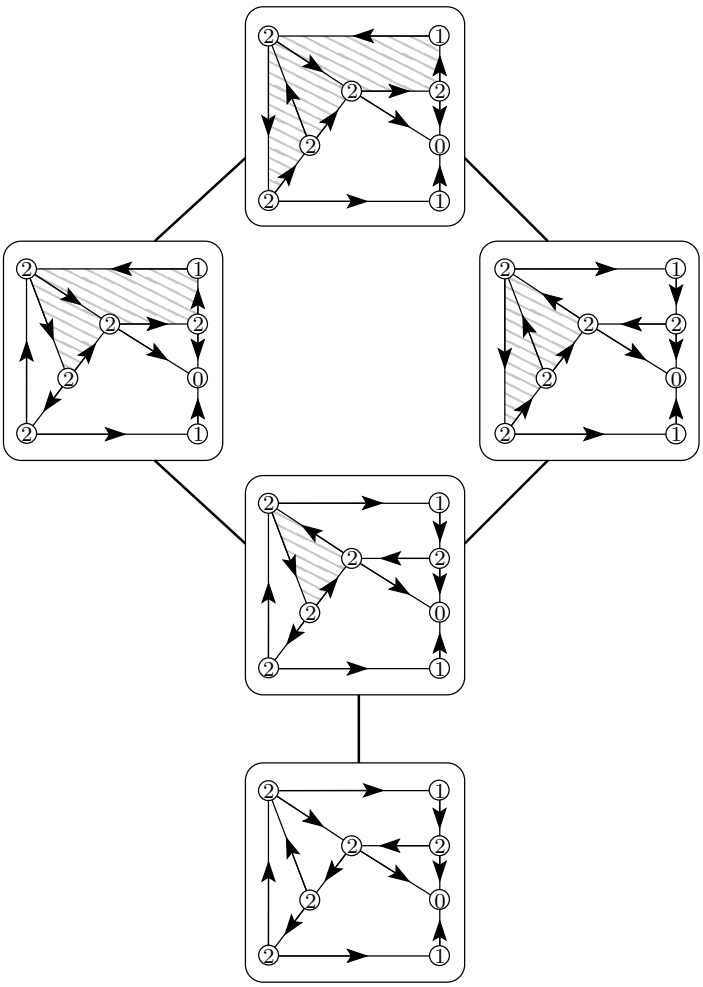

Figure 1.6: The distributive lattice on the set of all α-orientations of a plane graph G. The mapping $\alpha : V(G) \to \mathbb{N}$ is indicated at every vertex and the essential cycles that are reversed w.r.t. the minimal α-orientation are highlighted in grey.

- The inverse operation, i.e., from reversing from counterclockwise to clockwise, is called a *flop*.

In Figure 1.6, in each graph the set of essential cycles that are flipped w.r.t. the minimal orientation is highlighted in gray. Note that in general, this set is a multi-set.

Many combinatorial graph structures can be encoded as α-orientations. We have already mentioned in Subsection 1.1.4, that Schnyder woods are encoded by 3-orientations [dFOdM01]. To be precise, if G is an embedded maximally planar graph with outer vertices a_1, a_2, and a_3, we remove the three outer edges from G and define $\alpha(v) = 3$ for $v \notin \{a_1, a_2, a_3\}$, as well as $\alpha(a_1) = \alpha(a_2) = \alpha(a_3) = 0$. Such an α-orientation is called a 3-*orientation of* G. Clearly, every Schnyder wood induces a 3-orientation by disregarding all edge-labels. Conversely, one can show that every 3-orientation induces a Schnyder wood and that this is a bijection between Schnyder woods and 3-orientations.

There are many more examples of bijections between α-orientations for a particular mapping $\alpha : V(G) \to \mathbb{N}$ and certain combinatorial objects associated with G. In particular, every such bijection gives a distributive lattice structure on the set of these combinatorial objects by Theorem 1.2.2. The following lists structures that are in bijection with α-orientations of an associated graph. For many of them the distributive lattice given by Theorem 1.2.2 was already known before – the corresponding results are listed as well.

- Domino and lozenge tilings of a plane region (Rémila [Rém04] and others based on Thurston [Thu90])
- Spanning trees in planar graphs (Gilmer and Litherland [GL86], Propp [Pro93])
- Perfect matchings in planar bipartite graphs (Lam and Zhang [LZ03])
- d-factors in planar bipartite graphs (Felsner [Fel04], Propp [Pro93])
- Schnyder woods of a maximally planar graph (Brehm [Bre00], Ossona de Mendez [OdM94])
- Eulerian orientations of a planar graph (Felsner [Fel04])

Let us remark, that there are more general graph objects that carry a distributive lattice structure, which include α-orientations as a special case. For instance, Bernardi and Fusy introduce k-fractional orientations with prescribed out-degrees of a planar graph [BF10a], as well as Schnyder decompositions of a plane d-angulations of girth d [BF10b]. And already in 1993, Propp [Pro93] introduced c-orientations, which we address in the next subsection. For a comprehensive investigation of lattice structures on planar and non-planar graphs we refer to the work of Felsner and Knauer [FK09, FK11]. The PhD thesis of K. Knauer [Kna10] gives a nice and comprehensive survey.

1.2.1 Near-Linear Time Computation

We present an algorithm that given an n-vertex planar graph G and a mapping $\alpha : V(G) \to \mathbb{N}$ computes an α-orientation of G, or decides that one does not exist. Actually, our algorithm computes the minimal α-orientation of G. We also show how to compute the minimal α-orientation in linear time, provided we are given any α-orientation of G.

Although our algorithm is an immediate application of known results, it has not been stated in relation with α-orientations before (to the best of our knowledge). Previous algorithms for computing an α-orientation [Fel04, Fus07] rely on flow computations with running time $\mathcal{O}(n^{3/2})$. Our algorithm solves a single-source shortest path problem in a directed planar graph with possibly negative integer edge-lengths, which can be done in $\mathcal{O}(n \log^2(n)/\log\log n)$ due to a recent result of Mozer and Wulff-Nilsen [MWN10], which improves the (equally recent) $\mathcal{O}(n \log^2(n))$-algorithm of Klein, Mozer and Weimann from 2010 [KMW10]. We remark that, there is a linear-time algorithm for directed planar graphs with *non-negative* edge-lengths [KRRS94].

The main idea is the following: α-orientations of a plane graph G are in bijection with the so-called c-orientations of the dual graph G^* [Fel04]. The c-orientations, introduced by Propp [Pro93], in turn are in bijection with particular flow circulations with upper and lower capacities [Fel04]. Miller and Naor [MN95] reduced the problem of finding such a flow circulation to a single source shortest path problem. Hence, by this chain of reasoning finding an α-orientation can be reduced to finding a shortest-path tree in an appropriate directed planar graph. Here we summarize the main steps of this reduction, starting with the definition of a c-orientation.

Definition 1.2.3. Let $G = (V, E)$ be a graph, with a number $c(C) \geq 0$ and a fixed order of traversal for every cycle C in G. An orientation of the graphs edges is called a *c-orientation* if for every cycle C the number of forward-edges of C w.r.t. its traversal order equals $c(C)$.

A c-orientation prescribes the number of forward-edges of every cycle. It is enough to prescribe these numbers for the cycles of a cycle base of the graph. In a plane graph, we may choose the set of inner facial cycles as a cycle base. Now let G be a plane graph and G^* be its dual. Let v and f_v be a vertex and the corresponding face in G and G^*, respectively. Then the α-orientations of G are in bijection with the c-orientations of G^* by fixing the clockwise traversal order for every cycle and putting $c(f_v) = \alpha(v)$ for every $v \in V(G)$. The bijection is then given by the right-hand-rule:

- Traversing a primal edge along its direction in an α-orientation of G, the dual edge in the corresponding c-orientation of G^* is crossing from left to right.

We give an example in Figure 1.7 **a)**. It shows the plane dual of the graph in Figure 1.6 equipped with the minimal α-orientation and the corresponding c-orientation.

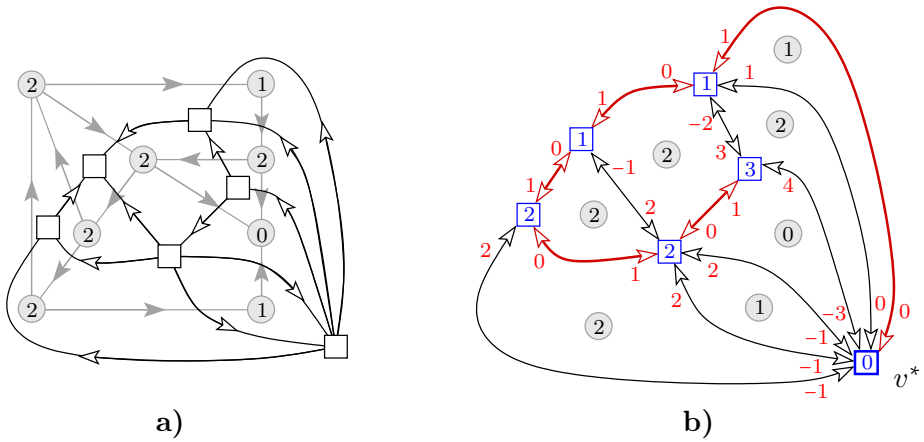

Figure 1.7: **a)** A plane graph with an α-orientation and its plane dual with the corresponding c-orientation. **b)** The corresponding bi-directed graph with the RFS-tree and the edge-lengths in red and π-values in blue.

It is known [FK09], that the set of c-orientations of a plane graph are in turn in bijection with certain vertex potentials, called Δ-*bonds*. We make a series of observations here from which the vertex potentials can be extracted. For the precise definitions and proofs we refer to the work of Felsner and Knauer [FK09].

Let v^* be an outer vertex in the dual graph G^*, for instance the vertex corresponding to the outer face in the primal graph G. Let T be the *right-first-search tree*, RFS-tree for short, of G^*. For example, in Figure 1.7 **b)** this tree is highlighted in red. Suppose for the moment, that we know some c-orientation of G^*.

- For every vertex $v \in V(G^*)$ we define π_v to be the number of edges on the unique v-to-v^* path in T that are directed towards v^*.

The numbers π_v are the blue numbers in Figure 1.7 **b)**. For a non-tree edge (v, w) we denote the unique cycle in $T \cup (v, w)$ by $C_T(v, w)$. Note that since T is a depth-first-search tree, every non-tree edge connects two vertices that have ancestor-descendant relation. Moreover, since T is a *right*-first-search tree, the clockwise traversal order of $C_T(v, w)$ goes along (v, w) from the descendant to the ancestor. We make the following three crucial observations.

- For every edge $(v, w) \in T$ with v being the parent of w in T we have:

$$\pi_v \leq \pi_w \leq \pi_v + 1 \tag{1.1}$$

- For every edge $(v,w) \notin T$ with v being an ancestor of w in T we have:

$$\pi_v + c(C_T(v,w)) - 1 \leq \pi_w \leq \pi_v + c(C_T(v,w)) \tag{1.2}$$

- For the vertex v^* we have:

$$\pi_{v^*} = 0 \tag{1.3}$$

Observation (1.1) is immediate, since (v,w) is directed either towards w, in which case $\pi_w = \pi_v$, or towards v, in which case $\pi_w = \pi_v + 1$. Observation (1.2) is similar, i.e., (v,w) is directed either towards w, in which case $\pi_w = \pi_v + c(C_T(v,w))$, or towards v, in which case $\pi_w = \pi_v + c(C_T(v,w)) - 1$.

Note that conditions (1.1), (1.2), (1.3) do not depend on the c-orientation whose knowledge we assumed. We have argued that our definition of $\{\pi_v \mid v \in V(G^*)\}$ results in a set of numbers satisfying the above conditions. In general, there are many such sets, and even many with solely non-negative integer values. The c-orientation of the dual graph corresponding to the minimal α-orientation of the primal graph is called the *minimal c-orientation*. Then, the following holds.

Lemma 1.2.4 ([FK09])**.** *The unique solution of* (1.1), (1.2), (1.3), *which minimizes* $\sum_{v \in V(G^*)} \pi_v$ *is obtained from the minimal c-orientation of G^* by*

$$\pi_v = \#\{\textit{edges in the } v\textit{-to-}v^* \textit{ path in } T \textit{ that are directed towards } v^*\}$$

By Lemma 1.2.4 the π-values corresponding to the minimal c-orientation are the solution to the following problem. Once the π-values are known, the minimal c-orientation can be easily recovered in linear time.

$$\begin{array}{lllll}
\text{minimize} & & \sum_{v \in V(G^*)} \pi_v & & \\
\text{such that:} & \pi_v & \leq \pi_w & & \forall (v,w) \in T \\
& \pi_w & \leq \pi_v + 1 & & \forall (v,w) \in T \\
& \pi_v & \leq \pi_w + 1 - c(C_T(v,w)) & & \forall (v,w) \notin T \\
& \pi_w & \leq \pi_v + c(C_T(v,w)) & & \forall (v,w) \notin T \\
& \pi_{v^*} & = 0 & &
\end{array}$$

The above has an interpretation as a single source shortest path problem in a directed planar graph $\tilde{G^*}$, which is defined as follows. Consider an edge (v,w) in G^* and let v be an ancestor of w in T. (If $(v,w) \in T$, then v is the parent of w in T.) There are two anti-parallel edges between v and w in $\tilde{G^*}$. The length of the edge directed from v to w is 1 if (v,w) is a tree-edge, and $C_T(v,w)$ if it is not, and the length of the edge directed from w to v is 0 if (v,w) is a tree-edge, and $1 - C_T(v,w)$ if it is not. Now it is easy to see that the solution of the above linear program equals the shortest distances from v^* in $\tilde{G^*}$ w.r.t. these edge-lengths. Figure 1.7 **b)** shows

the bi-directed graph $\tilde{G^*}$, the RFS-tree T in red, the edge-lengths in red, and for each vertex its distance from v^* in blue.

We conclude the following theorem.

Theorem 1.2.5. *For every planar n-vertex graph $G = (V, E)$ and every mapping $\alpha : V(G) \to \mathbb{N}$ the minimal α-orientation problem for (G, α) can be solved by a single source shortest path problem in a planar directed graph with (possibly negative) integer edge-lengths. Moreover, the currently fastest known algorithm takes $\mathcal{O}(n \log^2(n)/\log\log n)$ time [MWN10].*

Remark 1.2.6. Above we have chosen the tree T to be the RFS-tree in G^* rooted at some outer vertex v^*. Actually, we may choose *any* tree T rooted at *any* vertex v^*, and this would give us a (slightly more complicated) set of inequalities similar to (1.1), (1.2), (1.3), and in consequence different edge-lengths for $\tilde{G^*}$. (Clearly, the graph $\tilde{G^*}$ depends only on G^*.) Still, the distance of every vertex v from v^* w.r.t. these edge-lengths equals the number of edges on the v-to-v^* path in T that are oriented towards v^* in the minimal c-orientation of G^*. Interestingly, every choice of a tree T and a root vertex v^* results in a set of edge-lengths of $\tilde{G^*}$, such that of every two anti-parallel edges exactly one is on a shortest path from v^*.

One may ask in which cases the edge-lengths happen to be non-negative. We could then apply a linear-time algorithm for the shortest path problem [HKRS97]. However, it is not difficult to see, that for a pair (T, v^*) of a tree and a root vertex, the corresponding edge-lengths are non-negative if and only if every edge in T is oriented away from v^* in the minimal c-orientation. (Indeed, all edge-lengths are either 0 or 1 in this case.) However, knowing such a tree, one can directly read off the orientation of every non-tree edge in the minimal c-orientation.

Furthermore, if some c-orientation is given, we can compute in linear-time the corresponding π-values w.r.t. any rooted tree by counting edges that are oriented towards the root. These π-values satisfy (1.1), (1.2), and (1.3), i.e., are a feasible vertex potential of the shortest path problem. Based on this potential one can compute the so-called *reduced edge-lengths*, which are non-negative and keep all the shortest paths the same. (Indeed, every edge has length either 0 or 1 after the modification.) Hence, we can compute the minimal c-orientation in linear time, provided some c-orientation is given. Applying the bijection to α-orientations, we can compute the minimal α-orientation in linear time, provided some α-orientation is given.

Let us close this section by remarking that the α-orientation problem is equivalent to the arc-disjoint Menger problem. In the *arc-disjoint Menger problem* we are given a *directed* graph $G = (V, E)$ and an integer $b(v)$ for every vertex $v \in V$. The task is to compute a set $\mathcal{P}$ of arc-disjoint directed paths in G, such that the number of path starting at v minus the number of paths ending at v equals $b(v)$.

Suppose we are given a graph $G = (V, E)$ and a mapping $\alpha : V \to \mathbb{N}$. Then we can compute some directed graph G' that is an orientation of G. Let $\mathcal{P}$ be a solution of the arc-disjoint Menger problem for (G', b'), where $b'(v) = \text{out-deg}_{G'}(v) - \alpha(v)$ for every $v \in V$. Reversing every directed path in $\mathcal{P}$ then results in an α-orientation of G.

On the other hand, suppose we are given a directed graph $G = (V, E)$ and a mapping $b : V \to \mathbb{Z}$. Let G' the underlying undirected graph and G'' be a solution of the α-orientation problem for (G', α'), where $\alpha'(v) = \text{out-deg}_G(v) + b(v)$ for every $v \in V$. Then the symmetric difference $G \Delta G''$ is a set $\mathcal{P}$ of directed paths and cycles, where the set of paths is a solution of the arc-disjoint Menger problem for (G, b).

Note that we require planarity neither for the α-orientation problem, nor for the arc-disjoint Menger problem. However, having a planar graph in one problem translates into a planar graph in the other problem. We summarize without formal proof.

Theorem 1.2.7. *The α-orientation problem for directed and directed planar graphs is equivalent to the arc-disjoint Menger problem for directed and directed planar graphs, respectively.*

The arc-disjoint Menger problem has been considered in the literature, and linear-time algorithms are known only for very special cases. For instance, Brandes and Wagner [BW00] present a linear-time algorithm in case $b(v) = 0$ for all but two vertices in the graph.

1.3 Rectangle-Representations and Transversal Structures

In preparation for Chapter 2, this section is concerned with a special class of side contact representations. Let us first define a *near-triangulation* to be a planar graph G with at least five vertices, which admits a plane embedding with quadrangular outer face and only triangular inner faces. In other words, a near-triangulation is a maximally planar graph minus one outer edge. Throughout this thesis we are interested only in near-triangulations that are 4-*connected*, i.e., that remain connected under the removal of any set of three vertices. Sometimes and in especially within this section, we simply say near-triangulation although we always mean 4-connected near-triangulation.

For a near-triangulation G on n vertices, we always consider the embedding in which the outer face has degree 4, and denote the outer vertices by v_1, v_2, v_n, and v_{n-1}, in this counterclockwise order. For convenience, we assume that there is no

inner vertex that is a common neighbor of v_1 and v_n, unless $n = 5$. Then the graph $G' = G \cup (v_1, v_n)$ is a 4-connected maximally planar graph, provided $n > 5$, i.e., G' contains no separating triangle.

Definition 1.3.1. Let $G = (V, E)$ be an embedded near-triangulation with outer vertices v_1, v_2, v_n, v_{n-1} in counterclockwise order. A $\square$*-representation*, or *rectangle-representation* of G is a set $\Gamma = \{\mathcal{R}(v) \mid v \in V\}$ of axis-aligned rectangles in $\mathbb{R}^2$, one for each vertex, such that

- Any two rectangles $\mathcal{R}(v)$, $\mathcal{R}(w)$, for $v \neq w$, are interior disjoint.
- Two vertices v and w are connected by an edge in G if and only if $\mathcal{R}(v)$ and $\mathcal{R}(w)$ have a side contact.
- The union of all rectangles in Γ is a rectangle itself whose left side and right side is constituted by the left side of $\mathcal{R}(v_1)$ and the right side on $\mathcal{R}(v_n)$, respectively.

We define side contacts more formally in Definition 2.0.4 in Chapter 2. With the definitions and notation from Chapter 2, a $\square$-representation is the same as a hole-free rectilinear representation of G with complexity 4. For now, we just give an example of a near-triangulation and a $\square$-representation of it in Figure 1.8 and leave it at the intuitive meaning. We remark that a $\square$-representation of a near-triangulation G is also known as a *rectangular dual* of G. For a nice survey on rectangle-representations of planar graphs, we refer to the work of Felsner [Fel11].

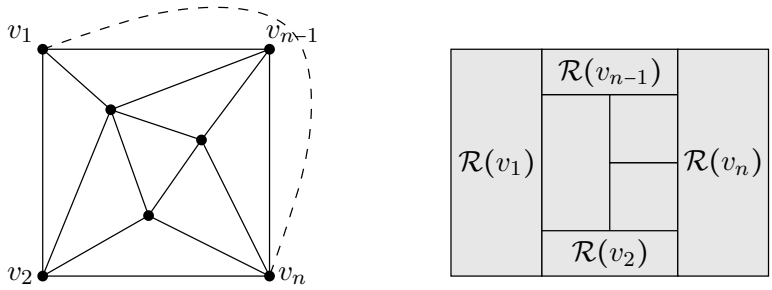

Figure 1.8: A near-triangulation and a rectangle-representation of it.

The following theorem appears in several independent sources [Ung53, LL84, KK85, Tho86, RT86], or can at least be derived from those.

Theorem 1.3.2. *Every near-triangulation has a* $\square$*-representation and it can be computed in linear time.*

We make frequent use of Theorem 1.3.2 within this thesis, e.g., in Lemma 2.1.4, Lemma 2.2.3, Theorem 4.2.4, and Theorem 4.2.6. Note that for instance, bipartite and planar graphs, as well as outer-planar graphs, are subgraphs of 4-connected near-triangulations, and thus some of the results in the thesis hold for these graphs as well.

It is known that $\square$-representations of near-triangulations are encoded by structures, which are similar to Schnyder woods. We define here transversal structures as introduced by Fusy [Fus07], which were independently considered by He [He93] under the name *regular edge labellings*. For a nice overview about regular edge labellings and their relations to geometric structures we refer to the introductory article by D. Eppstein [Epp10].

Definition 1.3.3. Let $G = (V, E)$ be a near-triangulation on $n \geq 5$ vertices with outer vertices v_1, v_2, v_n, v_{n-1} is counterclockwise order. A *transversal structure* of G is an orientation and labeling of the inner edges of G with labels $\{1, 2\}$, satisfying the following rules.

- Each inner vertex v has at least one outgoing and incoming edge of each label. In counterclockwise order around v occurs a set of outgoing edges of label 1, a set of outgoing edges of label 2, a set of incoming edges of label 1, and a set of incoming edges of label 2.
- For $i \in \{1, 2\}$, all inner edges incident to v_i are outgoing and have label i and all inner edges incident to v_{n+1-i} are incoming and have label i.

As for Schnyder woods, we associate colors with the labels, red for label 1 and blue for label 2. Figure 1.9 **a)** illustrates the local rule at an inner vertex and Figure 1.9 **b)** the local rule at outer vertices. For example, a (indeed the unique) transversal structure of the near-triangulation from Figure 1.8 is shown in Figure 1.9 **c)**. A transversal structure of a near-triangulation G can be interpreted as a combinatorial description of a $\square$-representation Γ of G. Red and blue edges correspond to vertical and horizontal side contacts, while the edge is oriented from the rectangle of the left or the bottom to the rectangle on the right or on the top, respectively. Figure 1.9 **d)** illustrates this correspondence. It is easy to see that every $\square$-representation defines a transversal structure this way. In fact, this holds the other around, i.e., every transversal structure comes from a $\square$-representation [KH97].

Fusy [Fus07] presents a bijection between transversal structures of G and the α_4-orientations of the angular map Q_G of G. Given an embedded near-triangulation G the *angular map* Q_G is the graph on $V(G) \cup F_{in}(G)$, where $F_{in}(G)$ denotes the set of inner faces of G, whose edges are of the form (v, f) with $f \in F_{in}(G)$ and $v \in V(G)$, v being incident to f. We define a mapping $\alpha_4 : V(Q_G) \to \mathbb{N}$ as follows.

- $\alpha_4(v) = 4$ for every $v \in V(G) \setminus \{v_1, v_2, v_{n-1}, v_n\}$
- $\alpha_4(v_1) = \alpha_4(v_n) = 0$
- $\alpha_4(v_2) = \alpha_4(v_{n-1}) = 2$
- $\alpha_4(f) = 1$ for every $f \in F_{in}(G)$

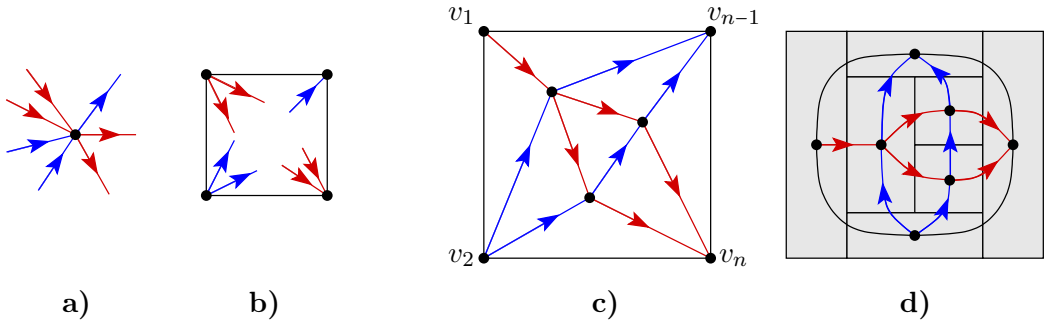

Figure 1.9: **a)** The local rule at an inner vertex. **b)** The local rule at the outer vertices. **c)** A transversal structure of a near-triangulation. **d)** The corresponding rectangle-representation.

In Figure 1.10 we show the angular map for the near-triangulation in Figure 1.8 and an α_4-orientation of it. The α_4-orientation has the following interpretation. Suppose Γ is a $\square$-representation of G, then the vertices of G correspond to the rectangles in Γ. The inner (triangular) faces $f = \{u, v, w\}$ of G, which are drawn as white vertices in Figure 1.10, correspond to those points p_{uvw} in the plane that are the common intersection of the three rectangles $\mathcal{R}(u)$, $\mathcal{R}(v)$, $\mathcal{R}(w)$ in Γ. Exactly two of these rectangles have a corner at p_{uvw}. Now the four outgoing edges at an inner vertex v point to those faces f where $\mathcal{R}(v)$ has a corner, and the single outgoing edge at an inner face points to the vertex, whose rectangle that does not have a corner at the face. It is easy to check that this way every $\square$-representation of G defines an α_4-orientation of Q_G. Again, the converse is true as well, i.e., every α_4-orientation of Q_G comes from a $\square$-representation of G. For the details and proofs we refer to the work of Fusy [Fus07] and just provide an example in Figure 1.10.

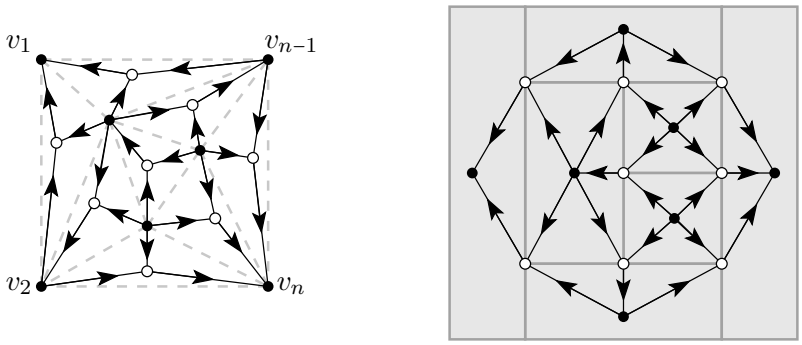

Figure 1.10: The angular map of a near-triangulation together with an α_4-orientation.

Recall from Section 1.2 that the set of all α-orientations of a planar graph is connected under flips an flops of essential cycles, i.e., under the reversal of directed such

cycles from clockwise to counterclockwise or vice versa. Moreover, the minimal α-orientation is the unique such orientation, that contains no counterclockwise directed cycle. The $\square$-representation of G that corresponds to the minimal α_4-orientation of Q_G is called the *minimal $\square$-representation* of G. Fusy [Fus07] has shown that every essential cycle of the α_4-orientations of an angular map Q_G is either a facial cycle (of length 4), or a non-facial cycle of length 8. We illustrate the flip at such an essential cycle and the corresponding local change in the $\square$-representation in the top row and the bottom row of Figure 1.11, respectively. A flop, i.e., a cycle reversal from counterclockwise to clockwise, is just the reverse operation. To close this section, note that the minimal α_4-orientation contains no counterclockwise directed cycle. Correspondingly, the right-hand versions of the local details in Figure 1.11 do *not* appear in the minimal $\square$-representation of G.

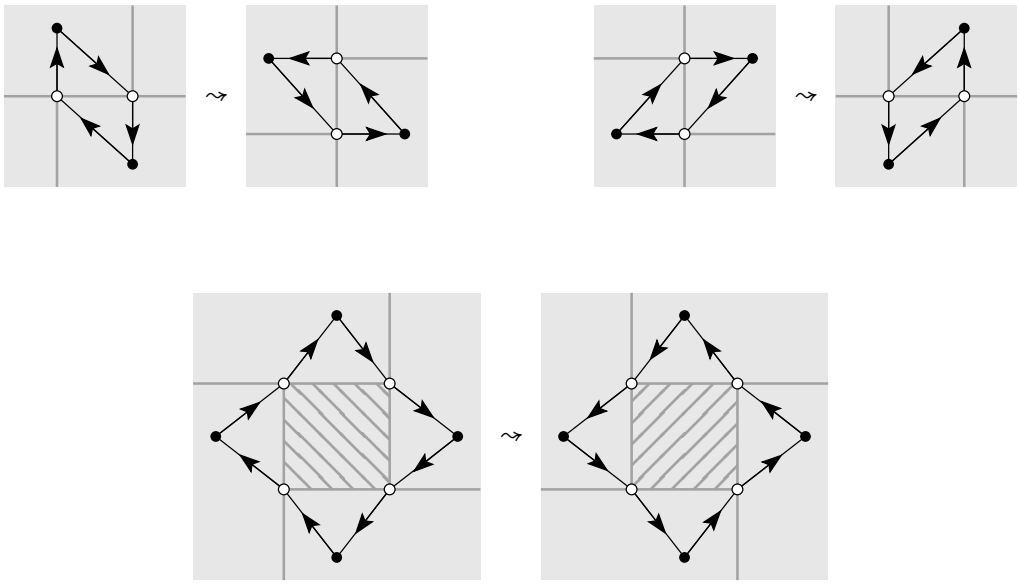

Figure 1.11: **Top row:** The two possibilities for facial cycle flips. **Bottom row:** A flip of an essential cycle of length 8.

Chapter 2

Side Contact Representations of Planar Graphs

In the present and the next chapter we are dealing with side contact representations of planar graphs with simple polygons. Two simple polygons $\mathcal{P}$ and $\mathcal{P}'$ in the plane may intersect in several ways. The intersection $\mathcal{P} \cap \mathcal{P}'$ may contain an interior point of $\mathcal{P}$ (and hence of $\mathcal{P}$ as well), or not. Let us assume that the polygons are interior disjoint, i.e., $\mathcal{P} \cap \mathcal{P}'$ is empty or consists of points that are on the boundary of both polygons. Seen as a point set in the plane, $\mathcal{P} \cap \mathcal{P}'$ is the union of some straight segments and isolated points. Every such segment or isolated point is called a *contact between* $\mathcal{P}$ *and* $\mathcal{P}'$, where segments are called *side contacts* and points are called *point contacts*. A point contact is either a corner of both polygons or a corner of one is contained in a side of the other polygon. Every side contact contains at least two corners of the two polygons, and at most two corners from either of the two. For example the intersection of the two interior disjoint polygons in Figure 2.1 consists of three point contacts and three side contacts.

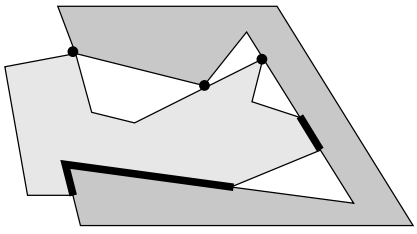

Figure 2.1: Two interior disjoint polygons with three point contacts and three side contacts.

We consider side contact representations of planar graphs, i.e., vertices are represented by simple and pairwise interior disjoint polygons, and edges correspond to side

contacts between the polygons of two adjacent vertices. We proceed with the formal definition of a representation. To be precise, a representation here is a *side contact representation*, while a *point contact representation* is defined analogous using point contacts instead of side contacts.

Definition 2.0.4. A *representation* of a planar graph $G = (V, E)$ is a set $\Gamma = \{\mathcal{P}(v) \mid v \in V\}$ of simple polygons, one for each vertex, with the following properties.

- Any two polygons $\mathcal{P}(v)$, $\mathcal{P}(w)$, for $v \neq w$, are interior disjoint.
- Two vertices v and w are connected by an edge in G if and only if $\mathcal{P}(v) \cap \mathcal{P}(w)$ contains a side contact.

For example, Figure 2.2 **b)** shows a representation of the planar graph G in Figure 2.2 **a)**. If (v, w) is an edge in G, then there may be several side contacts between the two corresponding polygons. For example, in Figure 2.2 **b)** $\mathcal{P}(5) \cap \mathcal{P}(6)$ consists of three consecutive side contacts, and $\mathcal{P}(1) \cap \mathcal{P}(4)$ consists of two non-consecutive side contacts. Note that for instance $\mathcal{P}(3) \cap \mathcal{P}(6)$ and $\mathcal{P}(1) \cap \mathcal{P}(5)$ consists of a point contact only, and indeed each of $(3, 6)$, $(1, 5)$ is *not* an edge in G.

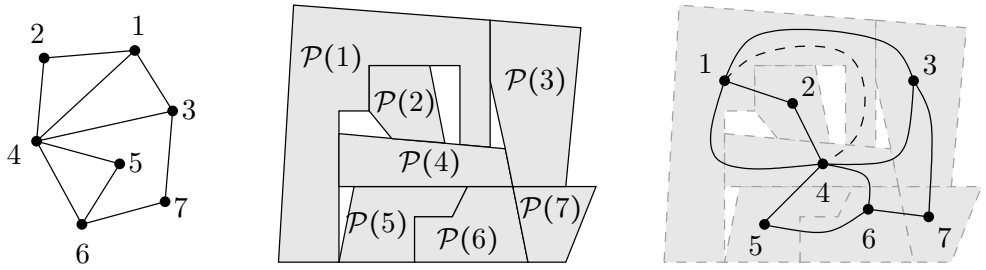

a) **b)** **c)**

Figure 2.2: **a)** A planar graph G. **b)** A representation Γ of G. **c)** A plane embedding of G inherited from Γ.

With every representation Γ we can associate a plane embedding of the represented graph G. We say that an embedding of G is *inherited* from the representation Γ, if every (v, w) in G can be assigned to one side contact in $\mathcal{P}(v) \cap \mathcal{P}(w)$, such that the clockwise order of assigned side contacts around $\mathcal{P}(v)$ is the same as the clockwise order of incident edges around v in the embedding, for every $v \in V(G)$. The presence of multiple side contacts between polygons in Γ may cause different embeddings to be inherited from the same representation. For example, Figure 2.2 **c)** shows an embedding of the graph G in Figure 2.2 **a)**, which is inherited from the representation Γ of G in Figure 2.2 **b)**. In the (necessarily plane) embedding every edge (v, w) is drawn as a curve from v to w that passes (exactly once) through $\mathcal{P}(v) \cap \mathcal{P}(w)$ at the

side contact the edge is assigned to. Note that assigning the edge $(5, 6)$ to any of the three side contacts between $\mathcal{P}(5)$ and $\mathcal{P}(6)$ gives the same embedding. However, assigning $(1, 4)$ to the other side contact between $\mathcal{P}(1)$ and $\mathcal{P}(4)$ corresponds to replacing the embedded edge $(1, 4)$ in Figure 2.2 **c)** with the dashed edge, which gives another, different, embedding of G. Moreover, note that the embedding of G in Figure 2.2 **a)** is *not* inherited from Γ.

The *complexity* of a polygon is defined as the number of sides it has (or equivalently, the number of its corners). For most purposes, polygons of low complexity are desired, or even required. Moreover, every bounded and unbounded region of $\mathbb{R}^2 \setminus \Gamma$ is a polygon itself, whose complexity also should be low. We call the polygon corresponding to the unbounded region the *(outer) boundary* of Γ. The *(polygonal) complexity of a representation* is the maximum complexity among its polygons.

Most of the results in this and the next chapter are of the following form. Given a planar graph G, we show the existence of a representation Γ of G with certain (desirable) properties and certain (low) complexity. We remark that in almost every case we assume G to be given with a fixed embedding, and construct the representation Γ in such a way that the embedding of G is inherited from it. The only exception is Theorem 3.2.2 and its corollaries Lemma 3.2.4 and Corollary 3.2.8.

We now define the set of desirable properties mentioned above. For a representation Γ, a *hole* is a bounded component of $\mathbb{R}^2 \setminus \Gamma$, i.e., a bounded subset of the plane surrounded by polygons in Γ. The representation in Figure 2.2 **b)** has four holes, two are surrounded by $\mathcal{P}(1)$, $\mathcal{P}(2)$, and $\mathcal{P}(4)$, and two further are surrounded by $\mathcal{P}(1)$, $\mathcal{P}(3)$, $\mathcal{P}(4)$, and $\mathcal{P}(1)$, $\mathcal{P}(4)$, $\mathcal{P}(5)$, respectively. Some holes correspond to faces of some embedding inherit from Γ, e.g., the first three mentioned above, and some do not, e.g., the fourth of these holes. On the other hand, the polygons of even more complex inner faces, such as $\{3, 4, 6, 7\}$ in the embedding in Figure 2.2 **c)**, need not surround a hole.

Of particular interest are representations without holes, which we call *hole-free representations*. If the polygons corresponding to the vertices of an inner face do not surround a hole, then there is a unique point in the plane that is the common intersection of all those polygons. This point may be considered as the dual vertex representing the inner face and the boundaries of the polygons as dual edges drawn as polygonal lines. This way, a hole-free representation Γ of G can be interpreted as a plane embedding of the dual graph of G minus the vertex for the exterior face. For example Figure 2.3 **c)** shows this embedding of the dual of the graph G in Figure 2.3 **a)** obtained from the hole-free embedding in Figure 2.3 **b)**. However, it is convenient to treat the unbounded subset $\mathbb{R}^2 \setminus \Gamma$ as (the complement of) a polygon itself. Then Γ can be seen as an embedding of a graph G' by putting a vertex on every point in the plane that is shared by three or more polygons and considering the boundaries of

polygons as polyline edges. The plane dual of G' without the vertex corresponding to the outer face is again the graph G that is represented by Γ. Hence, G' can be seen as a "polygonal dual" of G. Figure 2.3 **d)** shows the embedded graph G' associated with the representation Γ from Figure 2.3 **b)**.

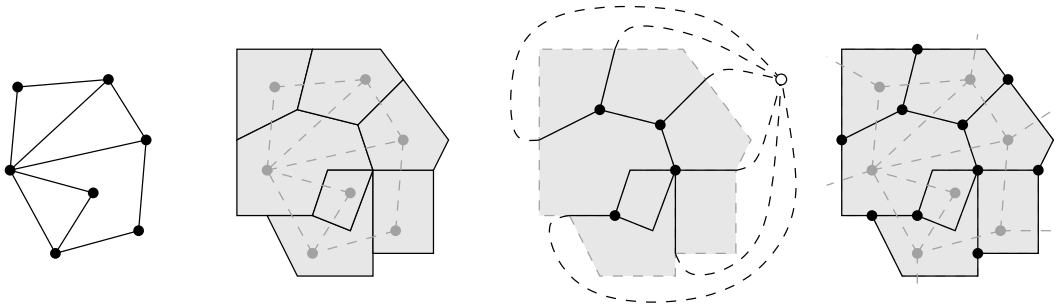

Figure 2.3: **a)** A planar graph G. **b)** A hole-free representation Γ of G. **c)** A plane embedding of the dual of G inherited from Γ. **d)** A "polygonal dual" of G.

In consequence, the study of hole-free representations can be translated into the study of plane embeddings with polyline edges. Indeed, several publications in this area are written in the language of those embeddings rather than in the language of side contact representations. We give some examples below.

As representations are often used for visualization purposes, their quality is often measured by the "readability". This term is vague and most likely means different things in different settings. However, one natural attempt to get "readable" representations is to restrict oneself to polygons with right angles, so-called *rectilinear polygons*. A polygon is rectilinear if every angle is a multiple of $\frac{\pi}{2}$ and a *rectilinear representation* is one in which every polygon is rectilinear. In rectilinear representations, every side contact is either horizontal or vertical. Moreover, the minimum complexity of a polygon is four, rather than three as in the general case. We usually denote a 4-sided polygon in a rectilinear representation by $\mathcal{R}(v)$ rather than $\mathcal{P}(v)$ to emphasize that it is a rectangle.

Another measure for readability is the so-called *minimum feature size*, which we define here as the minimum length of a side of a polygon. A large minimum feature size is supposed to increase the readability, and for instance rules out the possibility of polygons with "long, skinny arms". Clearly, nothing prevents us so far from scaling a representation so that it has a huge feature size. However, if we require the representation to fit into a small integer grid or have a prescribed area, it becomes interesting to give lower bounds on the minimum feature size.

Most of the work concerning side contact representations actually deals with rec-

tilinear hole-free representations. In the language of polyline embeddings these are called *rectilinear duals* or *floor plans*. Rectilinear duals were studied in graph theoretic context [Tam87], and in the context of Very-Large-Scale Integrated (VLSI) layouts and floor planning [Ott88, KP88]. Let us also refer to the survey of Eiglsperger *et al.* [EFK01]. The case when every polygon is a rectangle, i.e., so-called *rectangular duals*, received particular attention. The class of planar graphs that admit rectangular duals has been independently characterized several times [Ung53, LL84, KK85]. Buchsbaum *et al.* [BGPV08] provide some historical background and a summary of the rectangle contact graphs literature.

Many researchers [Rin87, SS93, SY93, GHK10, GHKK10] are interested in the following question.

Question 2.0.5. Given a graph class $\mathcal{G}$, what is the minimum number k, such that every graph $G \in \mathcal{G}$ has a representation with complexity at most k? What if we require the representation to be hole-free and/or rectilinear?

A complete answer to Question 2.0.5 for a particular graph class $\mathcal{G}$, would consist of two parts. First, for *every* graph $G \in \mathcal{G}$ we have to find a representation of G with polygonal complexity at most k. We sometimes call such a proof an "upper bound" for the graph class. And second, for *some* graph $G^* \in \mathcal{G}$ we have to show that *every* representation of G^* has polygonal complexity at least k, i.e., at least one polygon has complexity k or more. Consequently, we call this a "lower bound".

The table below summarizes all upper and lower bounds to Question 2.0.5 that we know of for a set of some graph classes of interest. The columns in Table 2.1 labeled LB and UB contain the lower and upper bounds, respectively. We consider here rectilinear and not necessarily rectilinear, as well as hole-free and not necessarily hole-free representations. Note that we have matching upper and lower bounds in every single case. Nonetheless, many interesting questions remain open – We address some of these questions in Section 2.1. For example, we do not know a characterization of those graphs that admit a non-rectilinear representation with polygonal complexity at most k for $k = 4, 5$. The case $k = 3$ becomes interesting when considering planar graphs that are not necessarily maximally planar. For rectilinear representations, the minimum size of the underlying grid is still unknown. We discuss this issue in more detail in Section 2.1.

This chapter is organized as follows.

Section 2.1: We introduce non-rotated rectilinear representations and reveal a correspondence of those to Schnyder woods. This enables us to characterize the existence of a non-rotated rectilinear representation of polygonal complexity at

Graph Class	Non-Rectilinear				Rectilinear			
	Holes		Hole-Free		Holes		Hole-Free	
	LB	UB	LB	UB	LB	UB	LB	UB
Maximally Outer-Planar	3	3	3	3 [GHK10]	4	4	6 [Rin87]	6 [ABF+11a]
4-Connected Near-Triangulation	4 [GHK10], Lem. 2.3.2	4	4	4	4	4	4	4 [KK85]
Hamiltonian Maximally Planar	5 [GHKK10], Lem. 2.3.2	5	5	5 Cor. 2.2.5	6	6	6	6 [SS93], Cor. 2.2.5
Planar 3-Tree	6 [GHKK10], Lem. 2.3.2	6	6	6 [SY93]	8	8	8	8
Maximally Planar	6	6	6	6 [GHKK10]	8	8	8	8 [SY93, LLY03], Thm. 2.1.8

Table 2.1: Summary of lower bounds and upper bounds for the polygonal complexity required for representations of all graphs within some graph class. (LB = lower bound, UB = upper bound)

most 4, 6 and 8, in terms of the existence of a Schnyder wood with certain properties. We prove sufficiency of our characterization by presenting a linear-time algorithm, that given a Schnyder wood with a corresponding additional property constructs a rectilinear representation with polygonal complexity at most 4, 6, or 8. For general maximally planar graphs, this representation requires the same worst-case area as previously known representations. In the end of the section, we prove that our representation (as well as the best previously known) is moreover worst-case optimal w.r.t. the number of segments and the so-called overall complexity, i.e., the total number of corners of all polygons in the representation.

Section 2.2: We review nesting assignments as introduced by Sun and Sarrafzedeh [SS93] and give a new proof for the sufficiency in their main result, i.e., that a maximally planar graph admits a rectilinear representation of complexity at most 6 if and only if it admits a nesting assignment. Our proof is a much simpler construction, which moreover runs in linear time. To our knowledge, only quadratic-time algorithms where known before. Additionally, a slight modification of our construction provides a representation of every such graph with convex pentagons. Even with the modification, the algorithm remains simple and linear-time. It was not known, whether every Hamiltonian maximally planar graph has a side contact representation with pentagons. Our result gives such a representation, since every Hamiltonian maximally planar graph has a nesting assignment.

Section 2.3: We present a general counting argument, which allows us to give lower bounds on the complexity of certain representations. This way, we reproduce all (tight) lower bounds for general (not necessarily rectilinear or hole-free) representations considered here, i.e., the first column of Table 2.1. Our proofs are much simpler than previous ones e.g. in [GHKK10], and our tight examples have significantly fewer vertices. To be precise, we provide a Hamiltonian maximally planar graph on 9 vertices and a planar 3-tree on 25 vertices of which every representation has complexity at least 5 and 6, respectively. The graphs constructed in [GHKK10] consist of 55 and 157 vertices, respectively, and the former graph is not Hamiltonian. Moreover, our new method allows us to derive new lower bounds in Chapter 3 (c.f. Lemma 3.3.2).

2.1 Non-Rotated Representations and Schnyder Woods

Throughout this section we consider hole-free rectilinear representations of low complexity, i.e., every vertex is represented by a rectilinear polygon with at most eight

sides, and the union of all polygons is again a simple polygon (indeed a rectangle). Not every planar graph can be represented with rectilinear polygons, such that the representation is hole-free. In particular, in a hole-free rectilinear representation every inner face has length 3 or 4. If every inner face is a triangle, the corresponding three polygons may have a point in common and this way the representation may be hole-free. Thus, in this section we consider maximally planar graphs only. Moreover, we will often drop the words hole-free and rectilinear, as all representations considered here have these properties.

Figure 2.4 shows a complete list of all rectilinear polygons with complexity at most 8. To our knowledge, every rectilinear representation from the literature with complexity at most 8 involves just a proper subset of the polygons in Figure 2.4. The most important shapes are rectangles, L-shapes, and T-shapes. Those are depicted in the first two rows in Figure 2.4. It is known [LLY03], that every maximally planar graph admits a rectilinear representation based on only one of the four T-shape polygons. For instance in the currently most compact floor-plan [LLY03] every vertex is represented by an upside-down T-shape or a degenerate version of it, where some "arms" have zero-length. In particular, no rotation of this fixed T-shape is required. We are going to derive a combinatorially different representation, which has all these good features, too. To underline the difference, we use an upright T-shape as the basic polygon. In Figure 2.4 this polygon and its four degenerate versions are highlighted in dark-grey.

Moreover, for special classes of maximally planar graphs, we derive a representation with complexity at most 6, which again relies on only one polygon with six sides. This L-shape polygon and its degenerate version, the rectangle, is marked with stripes in Figure 2.4. Let us formalize these concepts.

Definition 2.1.1. Let ⊤ and ⌜ denote the dark-grey T-shape polygon and the striped L-shape polygon in Figure 2.4, respectively.

- A *non-rotated* ⊤*-representation* of a maximally planar graph G is a hole-free representation Γ of G in which every polygon has a shape equivalent to a possibly degenerate ⊤.
- A *non-rotated* ⌜*-representation* of a maximally planar graph G is a hole-free representation Γ of G in which every polygon has a shape equivalent to a possibly degenerate ⌜.

First, consider a non-rotated ⌜-representation, i.e., every polygon is either a rectangle or has L-shape rotated as the striped L-shaped polygon in Figure 2.4. It is known [SY93], that not every maximally planar graph admits such a representation. Indeed, even allowing all rotations of an L-shape is not sufficient. However, the class of graphs, that *do have* a non-rotated ⌜-representation is non-trivial.

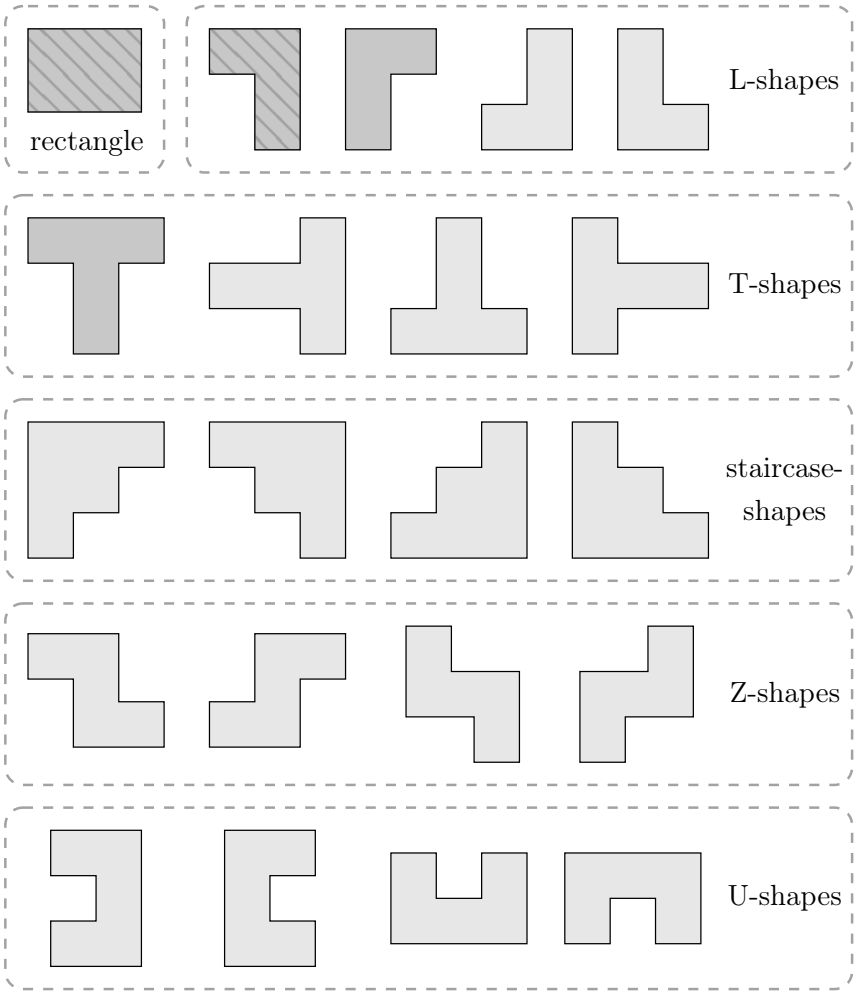

Figure 2.4: All rectilinear polygons with complexity at most 8.

Two non-rotated (possibly degenerate) L-shape polygons can have a side contact in only a few ways. A crucial observation, which is rigorously proved below, states that every such side contact involves the top-left, top-right, or bottom-right corner of one of the two polygons. The combinatorics of side contacts in a non-rotated $\sqsupset$-representation turns out to have a rich structure.

Definition 2.1.2. For a maximally plane graph G with outer triangle $\{v_1, v_2, v_n\}$ let Γ be a non-rotated $\sqsupset$-representation of G with $\mathcal{P}(v_n)$ at the top-right corner. Then the *Schnyder wood* (T_1, T_2, T_3) *inherited from* Γ is defined as follows.

- For an inner vertex v define $out_1(v)$ to be the vertex, whose polygon lies immediately left of the top-left corner of $\mathcal{P}(v)$.
- For an inner vertex v define $out_2(v)$ to be the vertex, whose polygon lies immediately underneath the bottom-right corner of $\mathcal{P}(v)$.

- For an inner vertex v define $out_3(v)$ to be the vertex, whose polygon lies immediately top-right of the top-right corner of $\mathcal{P}(v)$.
- For $i = 1, 2, n$ define T_i to be the set $\{(v, out_i(v)) \mid v \text{ inner vertex }\}$.

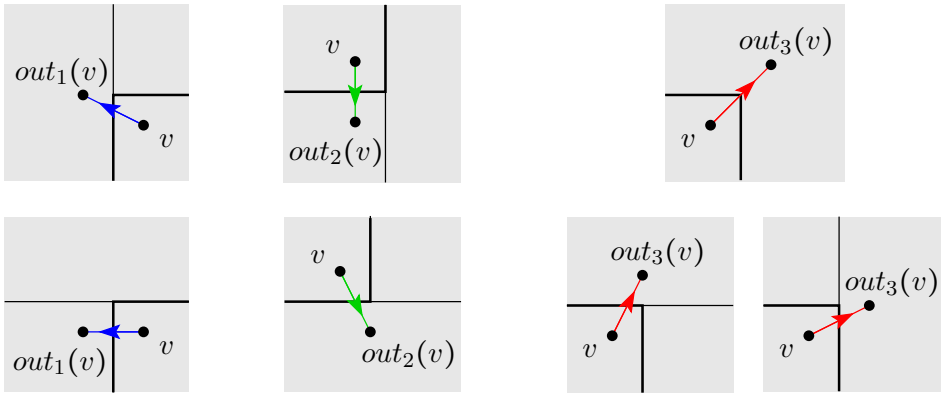

Figure 2.5: The definition of $out_1(v)$, $out_2(v)$, and $out_3(v)$ w.r.t. a non-rotated L-representation.

Lemma 2.1.3. *For any non-rotated ◱-representation Γ of a maximally planar graph G the Schnyder wood (T_1, T_2, T_3) inherited from Γ forms a Schnyder wood of G. Moreover, it holds $|L_1 \cap L_2 \cap in_3(v)| \leq 1$ for every vertex $v \in V(G)$.*

Proof. First we prove that every inner edge in G is contained in exactly one T_i. Therefore let v and w be two inner vertices of G. In case there are two distinct side contacts between $\mathcal{P}(v)$ and $\mathcal{P}(w)$, the top-right corner of one, say $\mathcal{P}(v)$, matches the concave corner of the other, $\mathcal{P}(w)$. By definition we have $out_3(v) = w$. Now if $\mathcal{P}(v)$ and $\mathcal{P}(w)$ have only a horizontal side contact, say $\mathcal{P}(v)$ lies above and $\mathcal{P}(w)$ below, then consider the rightmost contact point p. If it is a corner of $\mathcal{P}(v)$, it is its bottom-right corner, since we dealt with the concave case before. By definition we have $out_2(v) = w$. If p is no corner of $\mathcal{P}(v)$, it is one of $\mathcal{P}(w)$, namely its top-right corner. By definition we have $out_3(w) = v$. The case that $\mathcal{P}(v)$ and $\mathcal{P}(w)$ have only a vertical side contact is similar to the previous one, by considering the topmost contact point.

Note that only topmost and rightmost contact points result in defining some $out_i(v)$. From this it follows that every edge is contained in at most one T_i, and only in one direction.

Next we consider the polygon $\mathcal{P}(v)$ of an inner vertex v. Figure 2.5 depicts all possible types of side contacts and hence directly verifies the upcoming argumentation. An example of v and all its incident edges is given in the left of Figure 2.6.

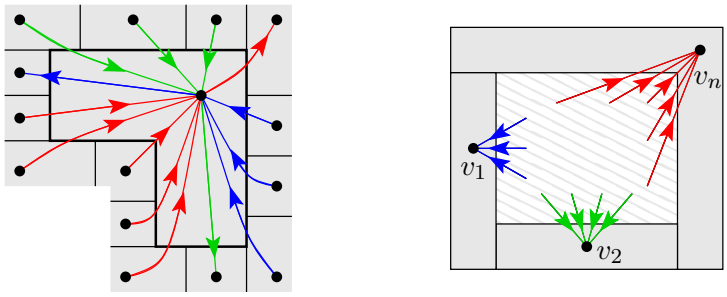

Figure 2.6: The Schnyder coloring at an inner vertex and the three outer vertices v_1, v_2, and v_n.

We trace the boundary of $\mathcal{P}(v)$ in clockwise direction, starting at its top-right corner, and investigate the T_i that a present side contact belongs to. At the top-right corner of $\mathcal{P}(v)$ is one outgoing edge in T_3. Along the right border is a set of incoming edges in T_1. Tracing along the bottom-left borders starts with one outgoing edge in T_2, proceeds with a set of incoming edges in T_3, and ends with one outgoing edge in T_1 at the top-left corner of $\mathcal{P}(v)$. Finally, along the top border there is a set of incoming edges in T_2. Note that the circular order of edges in T_1, T_2, and T_3 around v is consistent with the the definition of a Schnyder wood, c.f. Definition 1.1.4.

As illustrated in the right of Figure 2.6, for $i = 1, 2, 3$ every inner edge incident to the outer vertex v_i is incoming and contained in T_i. Thus (T_1, T_2, T_3) indeed is a Schnyder wood of G.

It remains to show that $|L_1 \cap L_2 \cap in_3(v)| \leq 1$ for every vertex v, i.e., there is at most one vertex in $in_3(v)$ that is a leaf in T_1 and T_2. To this end, consider an inner vertex w. By definition, the top-right corner of $\mathcal{P}(w)$ is shared by the polygon of $v = out_3(w)$. The right of Figure 2.5 shows that there is no third polygon sharing this point only if $\mathcal{P}(v)$ bends around $\mathcal{P}(w)$, i.e., has its concave corner there. On the other hand, if there is a third polygon $\mathcal{P}(z)$, then either $w = out_1(z)$ or $w = out_2(z)$, which means that $w \notin L_1 \cap L_2$. Hence we have $w \in L_1 \cap L_2$ only if the top-right corner of $\mathcal{P}(w)$ matches the concave corner of $\mathcal{P}(v)$. In other words, for every v there is at most one vertex in $in_3(v) \cap L_1 \cap L_2$, which proves the lemma. □

A □-representation (c.f. Definition 1.3.1) can be seen as a non-rotated ◱-representation in which every L-shape polygon is degenerated to a rectangle. To be precise, we have defined a non-rotated ◱-representation and a □-representation only for a maximally planar graph and a near-triangulation, i.e., a maximally planar graph *minus one edge*, respectively. However, Lemma 2.1.3, and in particular the additional property of the Schnyder wood inherited from a non-rotated ◱-representation, allows us to associate a Schnyder wood of G to a □-representation of $G \setminus (v_1, v_n)$. In con-

sequence, we obtain such a Schnyder wood whenever G is a 4-connected maximally planar graph.

Lemma 2.1.4. *Let G be a 4-connected maximally planar graph with outer triangle $\{v_1, v_2, v_n\}$ and v_{n-1} be the fourth outer vertex of $G \setminus (v_1, v_n)$, then G admits a Schnyder wood (T_1, T_2, T_3) with $L_1 \cap L_2 = \{v_{n-1}\}$.*

Proof. Consider a $\square$-representation Γ of $G \setminus (v_1, v_n)$. According to Definition 1.3.1 the rectangles $\mathcal{R}(v_1)$, $\mathcal{R}(v_2)$, $\mathcal{R}(v_{n-1})$, and $\mathcal{R}(v_n)$ are represented in Γ as in the middle of Figure 2.7, i.e., $\mathcal{R}(v_1)$, respectively $\mathcal{R}(v_2)$, is the leftmost, respectively bottommost, rectangle in Γ. Moreover, the side contact between $\mathcal{R}(v_{n-1})$ and $\mathcal{R}(v_1)$ as well as $\mathcal{R}(v_n)$ is horizontal. We establish a side contact between $\mathcal{R}(v_1)$ and $\mathcal{R}(v_n)$ by bending $\mathcal{R}(v_n)$ around $\mathcal{R}(v_{n-1})$ as illustrated in the right of Figure 2.7. The resulting is a non-rotated $\sqsupset$-representation Γ' of the original graph G.

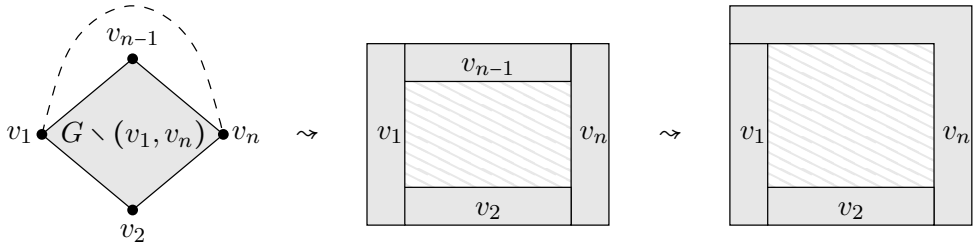

Figure 2.7: The graph $G \setminus (v_1, v_n)$, a rectangle-representation of it, and a non-rotated L-representation of G.

Consider the Schnyder wood (T_1, T_2, T_3) inherited from Γ'. In the proof of Lemma 2.1.3 we show $v \in L_1 \cap L_2$ for an inner vertex v only if the top-right corner of $\mathcal{P}(v)$ matches the concave corner of $\mathcal{P}(out_3(v))$. Since $\mathcal{P}(v_n)$ is the only polygon with a concave corner and $v_n = out_3(v_{n-1})$, we have $L_1 \cap L_2 = \{v_{n-1}\}$. $\square$

By Lemma 2.1.4, Definition 2.1.2 associates with every $\square$-representation of $G \setminus (v_1, v_n)$ a Schnyder wood of G. A $\square$-representation in turn is in bijection with a transversal structure. Hence, Definition 2.1.2 can be seen as a mapping ϕ from the set of all transversal structures of $G \setminus (v_1, v_n)$ to the Schnyder woods of G. From Lemma 2.1.4 we know that ϕ maps only onto Schnyder woods with $L_1 \cap L_2 = \{v_{n-1}\}$. Indeed, Lemma 2.1.6 below shows that every such Schnyder wood is obtained from a transversal structure via the map ϕ. However, ϕ is not one-to-one. For instance, consider the two different $\square$-representations Γ_1 and Γ_2 of the same maximally planar graph G in Figure 2.8 **a)**. As shown in the figure both Γ_1 and Γ_2 inherit *the same* Schnyder wood. On the other hand, if we mirror the situation as done in Figure 2.8 **b)**, then the two different $\square$-representations inherit *two different* Schnyder

woods. Moreover, in this case the flip in the α_4-orientation between the transversal structures corresponding Γ_1 and Γ_2 corresponds to a flip in the Schnyder wood.

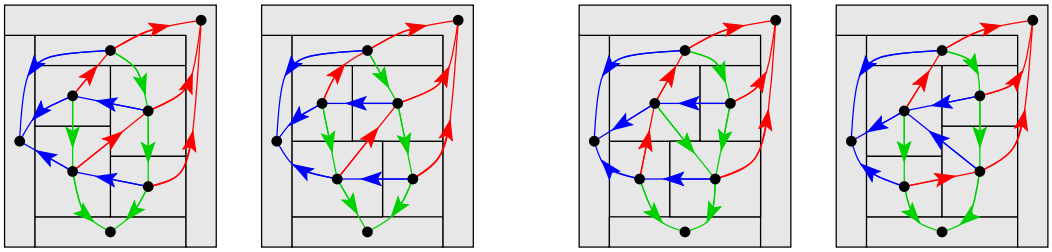

Figure 2.8: **a)** Two different representations inherit *the same* Schnyder wood. **b)** Two different representations inherit *two different* Schnyder woods.

Remark 2.1.5. Within the scope of this thesis we do not study the connection between the traversal structures of $G \setminus (v_1, v_n)$ and the Schnyder woods of G with $L_1 \cap L_2 = \{v_{n-1}\}$ in more detail. However, we think that this is a field worth investigating; see Question 2 at the end of the thesis. For instance, how does the rich structure of the set of all transversal structures and α_4-orientations carry over to those Schnyder woods? What characterizes the Schnyder wood inherited from the minimal or maximal $\square$-representation? What local operations transfer a Schnyder wood with $L_1 \cap L_2 = \{v_{n-1}\}$ into another one with the same property?

In Lemma 2.1.3 and Lemma 2.1.4 we derive from a non-rotated $\sqsupset$-representation and a $\square$-representation a Schnyder wood with one additional property. Next, we show that the reverse direction can be done as well. That is, given a Schnyder wood (T_1, T_2, T_3) of G satisfying $|L_1 \cap L_2 \cap in_3(v)| \leq 1$ for every vertex $v \in V(G)$ or $L_1 \cap L_2 = \{v_{n-1}\}$, we construct a non-rotated $\sqsupset$-representation or a $\square$-representation of G, respectively. Moreover, the given Schnyder wood will be inherited from the constructed representation. Actually, we show something stronger, namely that every Schnyder wood is inherited from some non-rotated $\sqsupset$-representation. We remark that we defined how a Schnyder wood is inherited only from a non-rotated $\sqsupset$-representation. Indeed, we require the non-rotated $\sqsupset$-representation to satisfy an additional property, in order to obtain a Schnyder wood from it.

If polygon $\mathcal{P}$ has the T-shape as highlighted in Figure 2.4, we call the rectangle spanned by the top-right corner and right reflex corner of $\mathcal{P}$, the *right arm* of $\mathcal{P}$, denoted by $\mathcal{P}_R$. Similarly, the *left arm* of $\mathcal{P}$, denoted by $\mathcal{P}_L$, is the rectangle spanned by the top-left corner and the left reflex corner of $\mathcal{P}$. See Figure 2.9 **a)** for an illustration.

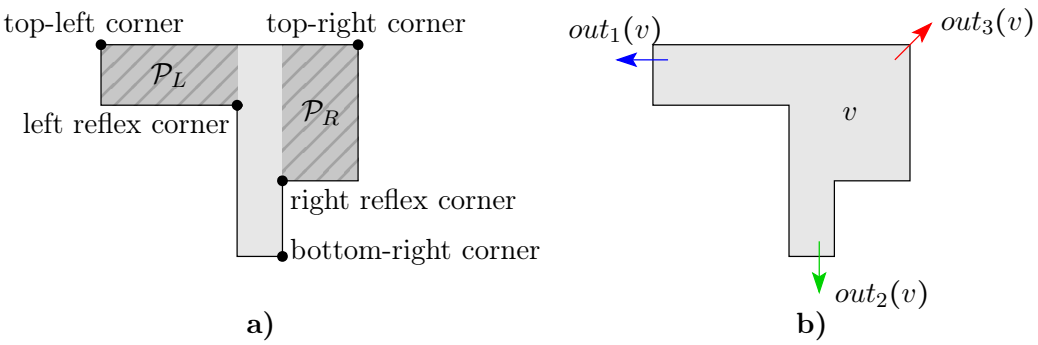

Figure 2.9: **a)** The left arm $\mathcal{P}_L$ and the right arm $\mathcal{P}_R$ of an upright T-shape polygon. **b)** Illustration of how a Schnyder wood is inherited from a non-rotated representation with (possibly degenerate) upright T-shape polygons.

A non-rotated ⊤-representation Γ that is to be constructed in the next lemma, has the following property. If a polygon $\mathcal{P}$ has a T-shape, then the bottom side of $\mathcal{P}_R$ is a (horizontal) side contact with only one other polygon. The same holds if the left arm of $\mathcal{P}$ is degenerate and hence $\mathcal{P}$ has complexity 6. This property is sufficient (and actually also necessary) to obtain a Schnyder wood from Γ in the "same" way as in Definition 2.1.2.

- A non-rotated ⊤-representation inherits a Schnyder wood as indicated in Figure 2.9 **b)** if and only if the bottom side of every right arm $\mathcal{P}_R$ is a side contact with only one other polygon.

We omit the details here and sketch just the basic idea. First, define the top-right, top-left, and bottom-right corner of a T-shape polygon $\mathcal{P}$ as indicated in Figure 2.9 **a)**, and define $out_1(v)$, $out_2(v)$, and $out_3(v)$ as in Definition 2.1.2. The location of $\mathcal{P}(out_1(v))$, $\mathcal{P}(out_2(v))$, and $\mathcal{P}(out_3(v))$ is then indicated by the blue, green, and red arrows in Figure 2.9 **b)**, respectively.

We are now ready for the following lemma.

Lemma 2.1.6. *Let G be a maximally planar graph with outer triangle $\{v_1, v_2, v_n\}$ and (T_1, T_2, T_3) be a Schnyder wood of G. Then each of the following holds.*

(a) G has a non-rotated ⊤-representation, from which (T_1, T_2, T_3) is inherited.
(b) If $|L_1 \cap L_2 \cap in_3(v)| \leq 1$ for every vertex v, then G has a non-rotated ◱-representation, from which (T_1, T_2, T_3) is inherited.
(c) If $L_1 \cap L_2$ consists solely of the fourth outer vertex in $G \setminus (v_1, v_n)$, then $G \setminus (v_1, v_n)$ has a □-representation, from which (T_1, T_2, T_3) is inherited.

Proof. We construct each required representation based on a Schnyder wood with the given properties. Therefore let (T_1, T_2, T_3) be *any* Schnyder wood of G. We construct a non-rotated $\sqcap$-representation of G along the building sequence w.r.t. any canonical order associated with (T_1, T_2, T_3), i.e., a topological order $(v_1, v_2, \ldots, v_n)$ of $T_1^{-1} \cup T_2^{-1} \cup T_3$. (See Lemma 1.1.6.) We start with the representation of $G_3 := G[v_1, v_2, v_3]$ depicted in Figure 2.10.

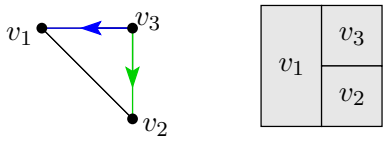

Figure 2.10: The graph G_3 and a (degenerate) non-rotated T-representation.

We maintain the following invariants for $i \geq 3$:

- Γ_i is a non-rotated $\sqcap$-representation of $G_i := G[v_1, v_2, \ldots, v_i]$.
- The outer boundary of Γ_i partitions into the left and bottom border of $\mathcal{R}(v_1)$, the bottom border of $\mathcal{R}(v_2)$, and a monotonously decreasing staircase-path.
- The horizontal segments of the staircase are given by the top border of $\mathcal{R}(v_1)$ and the top border of every $\mathcal{P}(v)$, such that $(v, out_1(v))$ is an outer edge in G_i.
- The vertical segments of the staircase are given by the right border of $\mathcal{R}(v_2)$ and the right border of every $\mathcal{P}(v)$, such that $(v, out_2(v))$ is an outer edge in G_i.

The invariant is illustrated in Figure 2.11.

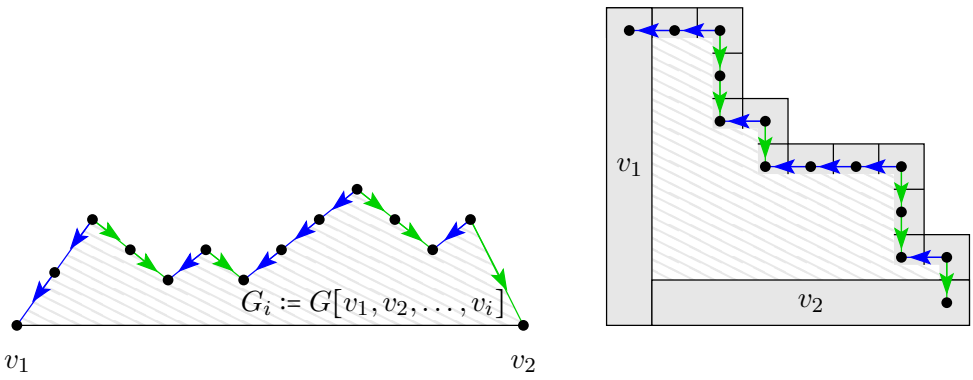

Figure 2.11: The graph G_i and a non-rotated T-representation with a staircase-shape.

Note that the (top-right) convex and concave corners of the staircase correspond to the hill and valley vertices in G_i, respectively. We now show how to insert the polygon $\mathcal{P}(v_{i+1})$ into Γ_i, such that the invariant again holds for the resulting representation Γ_{i+1}. We distinguish the following cases:

Case (1*a*)**:** $in_3(v_{i+1})$ contains no hill but one valley vertex.
Case (1*b*)**:** $in_3(v_{i+1})$ contains no hill and no valley vertex.
Case (2)**:** $in_3(v_{i+1})$ contains one hill vertex.
Case (3)**:** $in_3(v_{i+1})$ contains at least two hill vertices.

Case (1*a*)**:** Let v denote the valley vertex in $in_3(v_{i+1})$. Consider a ray starting right underneath the top-right corner of $\mathcal{P}(v)$ and emerging from there to the left through the entire representation. Then stretch the representation above the ray by some positive amount. The result is still a combinatorially equivalent ⊤-representation with the right border of $\mathcal{P}(v)$ now being partially on the (still staircase) boundary. Now put a rectangle $\mathcal{R}(v_{i+1})$ with its bottom-left corner onto the former position of the top-right corner of $\mathcal{P}(v)$ and its top-right corner horizontally aligned with $\mathcal{P}(out_1(v_{i+1}))$ and vertically aligned with $\mathcal{P}(out_2(v_{i+1}))$. This case is illustrated in the left two images of Figure 2.12.

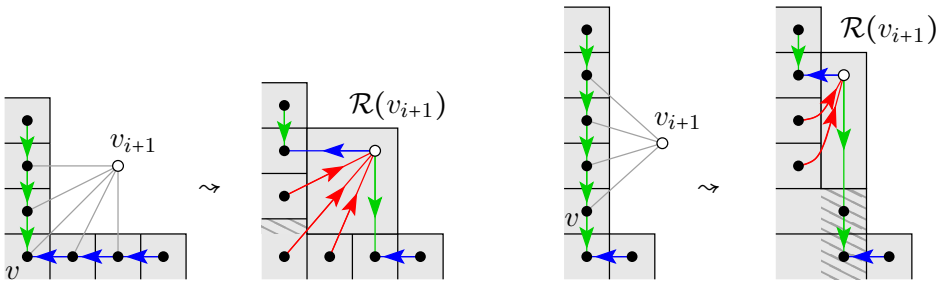

Case (1*a*) Case (1*b*)

Figure 2.12: Adding vertex v_{i+1} to the graph and rectangle $\mathcal{R}(v_{i+1})$ to the non-rotated T-representation according to Case (1).

Case (1*b*)**:** In this case v_{i+1} is connected to a path either in T_1 or T_2. We consider the latter case, which is depicted in the right of Figure 2.12. (The former case is similar.) Denote $v \coloneqq out_2(v_{i+1})$ and consider a ray starting immediately to the left of the top-right corner of $\mathcal{P}(v)$ and emerging from there downwards through the entire representation. Then stretch the representation to the right of the ray by some positive amount. The result is still a combinatorially equivalent ⊤-representation with the top border of $\mathcal{P}(v)$ now being partially on the (still staircase) boundary. Now put a rectangle $\mathcal{R}(v_{i+1})$ in exactly the same way as in the previous case.

Case (2)**:** Let v denote the hill vertex in $in_3(v_{i+1})$. We split the vertex v_{i+1} into two, denoted by v' and v''. Connect v' to neighbors of v_{i+1} starting with $out_1(v_{i+1})$ and ending with the hill vertex v and introduce a rectangle $\mathcal{R}(v')$ according to Case (1). Afterwards connect v'' to v' and the remaining neighbors of v_{i+1} and again introduce a rectangle $\mathcal{R}(v'')$ according to Case (1). Define the union

of $\mathcal{R}(v')$ and $\mathcal{R}(v'')$ to be the L-shape polygon $\mathcal{P}(v_{i+1})$. This case is illustrated in Figure 2.13.

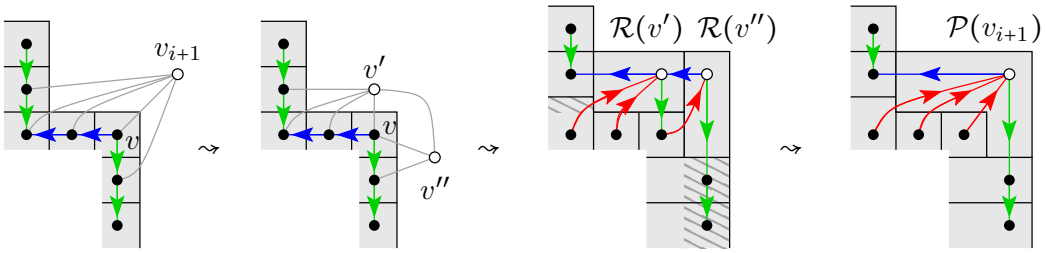

Figure 2.13: Adding vertex v_{i+1} to the graph and an L-shape polygon $\mathcal{P}(v_{i+1})$ to the non-rotated T-representation according to Case (2).

Case (3): Let v and w denote the rightmost and leftmost hill vertex in $in_3(v_{i+1})$, respectively. Extending the idea from Case (2), we define a sequence of artificial vertices, each of which we add to Γ_i according to Case (1) and afterwards merge with an already existing polygon. Consider v, the vertex in $in_3(v_{i+1})$ at the rightmost convex staircase corner, and $u := out_1(v)$, which lies on the boundary of G_i. Introduce a vertex x incident only to u and v and add a rectangle $\mathcal{R}(x)$ according to Case (1). Then merge $\mathcal{P}(u)$ with $\mathcal{R}(x)$, which increases the complexity of $\mathcal{P}(u)$. In particular if $\mathcal{P}(u)$ was an L-shape before, it is a T-shape now. Since $u \in in_3(v_{i+1})$ the polygon $\mathcal{P}(u)$ will disappear from the staircase boundary once $\mathcal{P}(v_{i+1})$ is introduced and hence its complexity will not increase a second time; see the first step in Figure 2.14. Note that the right border of u appears on the staircase boundary now, while the top border of v disappeared. If this intermediate representation would satisfy the invariant (which it does not), we would have $(u, v) \in T_2$ instead of $(u, v) \in T_1$. That is why in Figure 2.14 after the modification the edge (u, v) is drawn dashed.

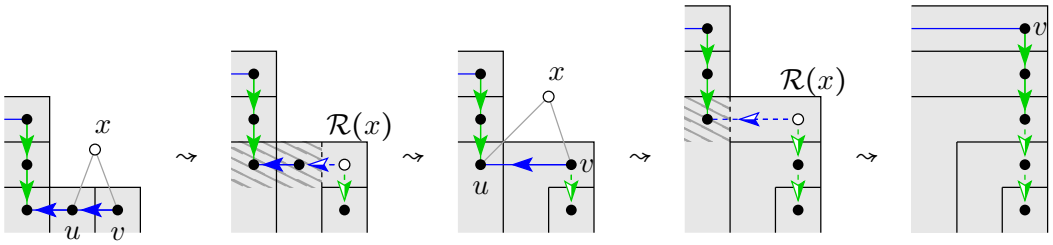

Figure 2.14: Removing one step of the staircase in $in_3(v_{i+1})$ according to Case (3).

Now the former u is the vertex in $in_3(v_{i+1})$ at the rightmost convex staircase corner and thus becomes the new v. We iterate this procedure with $v := u$ until $(v, out_1(v))$ is an inner edge, i.e., v is a valley vertex. Then extend every polygon on the outer T_2-path ending at v to the right so that it aligns with $\mathcal{P}(v)$. Let again v denote the vertex in $in_3(v_{i+1})$ at the rightmost convex staircase corner; see the right

of Figure 2.14 for an illustration. Note that the number of steps of the staircase is reduced by one. We iterate this procedure until $v = w$, i.e., only one vertex in $in_3(v_{i+1})$ is located at a convex staircase corner. We now introduce an L-shape polygon for v_{i+1} according to Case (2). In doing so we act as if the invariant would hold for the current representation, i.e., every edge (x, y) between the initial v and w is in T_2.

It is not difficult to see that in each case we get a non-rotated ⊤-representation of G_{i+1} satisfying the invariant above. This proves part (a) of the lemma. In order to prove parts (b) and (c), note that for an inner vertex v we have $v \in L_1 \cap L_2$ if and only if v is a hill vertex at the time $out_3(v)$ is inserted. Hence if the Schnyder wood (T_1, T_2, T_3) satisfies $|L_1 \cap L_2 \cap in_3(v)| \leq 1$ for every vertex v, then no vertex is inserted according to Case (3), which is the only case that alters the complexity of any existing polygon. Since each vertex is introduced either as a rectangle in Case (1) or as a non-rotated L-shape in Case (2) we conclude that we end up with a non-rotated ⌜-representation in case we have $|L_1 \cap L_2 \cap in_3(v)| \leq 1$ for every vertex v, which proves (b).

Now assume $L_1 \cap L_2 = \{v_{n-1}\}$ and $(v_1, v_{n-1}) \in E(G)$. We have $out_3(v_{n-1}) = v_n$ and thus v_n is the only vertex for which Case (1) does not apply. Consequently the representation resulting from the above construction looks like that in the left of Figure 2.15. A local modification, as illustrated in the figure, then gives a □-representation of $G \setminus (v_1, v_n)$, which finally proves (c). □

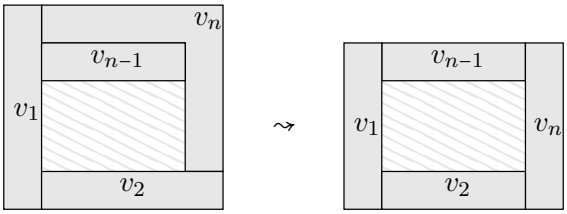

Figure 2.15: The resulting non-rotated L-representation of G in case $L_1 \cap L_2 = \{v_{n-1}\}$ and $(v_1, v_{n-1}) \in E(G)$, and a rectangle-representation of $G \setminus (v_1, v_n)$.

Let us provide an example of the construction procedure from Lemma 2.1.6. Figure 2.16 **a)** depicts a maximally planar graph G on 12 vertices and a Schnyder wood (T_1, T_2, T_3) of G. The vertices are numbered w.r.t. a canonical order corresponding to that Schnyder wood. In Figure 2.16 **b)**–**k)** we show the partial representations of G after every step. The polygons for vertices 11 and 12 are introduced according to Case (3) and Case (2), respectively. Figures 2.16 **j')** and **j")** illustrate the two step procedure for Case (3). The polygons for vertices 8 and 10 are introduced according to Case (1a), and for all other polygons (except the starting triangle $\{1, 2, 3\}$) Case (1b) applies. Note that the representation constructed by this method uses only

integer coordinates. Moreover, only the introduction of $\mathcal{P}(11)$ and $\mathcal{P}(12)$ requires three and two new grid lines, respectively. With every polygon other than $\mathcal{P}(11)$ and $\mathcal{P}(12)$ only one new grid line is introduced. We are going to address this issue in the upcoming Theorem 2.1.8.

Taking Lemma 2.1.4 and Lemma 2.1.6 together, we have proven a characterization of those maximally planar graph that admit a non-rotated $\sqsupset$-representation and those near-triangulations that admit a $\square$-representation. We summarize this in the following theorem.

Theorem 2.1.7. *Each of the following holds.*

(a) A maximally planar graph has a non-rotated $\sqsupset$-representation if and only if it admits a Schnyder wood (T_1, T_2, T_3) with $|L_1 \cap L_2 \cap in_3(v)| \leq 1$ for every vertex v.

(b) A near-triangulation $G \setminus (v_1, v_n)$ has a $\square$-representation if and only if G admits a Schnyder wood (T_1, T_2, T_3) with $L_1 \cap L_2 = \{v_{n-1}\}$.

Moreover, if an adequate Schnyder wood is given, the corresponding representation can be constructed in linear time.

Considering all maximally planar graphs, we have proven the existence of a non-rotated $\sqcap$-representation, which additionally inherits any given Schnyder wood. We can even show that this representation is very compact, i.e., requires only few supporting grid lines.

Theorem 2.1.8. *Every maximally planar graph on n vertices admits a non-rotated $\sqcap$-representation with at most $\lfloor \frac{5n-2}{3} \rfloor$ supporting grid lines. Moreover, this representation can be constructed in $\mathcal{O}(n)$ time.*

Proof. The representation is the one constructed in Lemma 2.1.6. Therefore let (T_1, T_2, T_3) be any Schnyder wood of G, which can be found in linear time. W.l.o.g. we assume that $|L_2| \leq \lfloor \frac{2n-5}{3} \rfloor$ by Lemma 1.1.5. In the first step, the starting triangle $G_3 = \{v_1, v_2, v_3\}$ is represented in a 2×2 grid; see Figure 2.7. (As in previous results [LLY03], we consider the 2×2 grid to consist of four grid lines only.) The resulting representation is denoted by Γ_3. Now, consider step $i+1$ of the algorithm, for $i \geq 3$, i.e., vertex v_{i+1} is added to G_i and a corresponding polygon $\mathcal{P}(v_{i+1})$ to the representation Γ_i. We bound the number of new grid lines required in this step by distinguishing the same cases as in Lemma 2.1.6.

Case (1)**:** In both Case (1a) and Case (1b), one new grid line is required; see Figure 2.6.

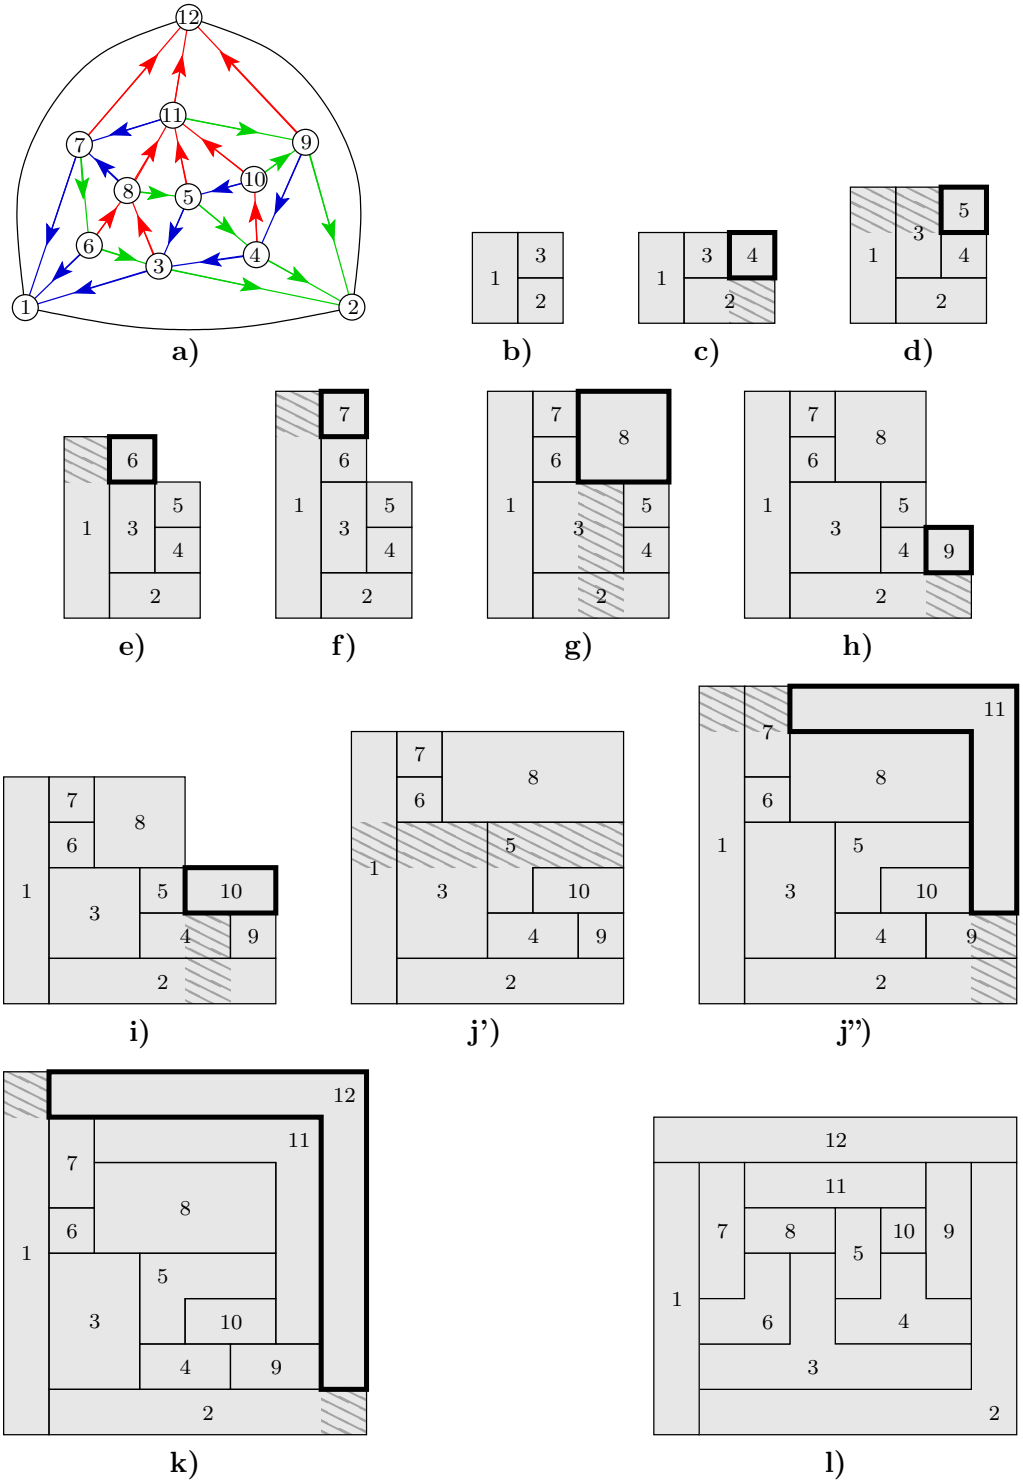

Figure 2.16: **a)** A maximally planar graph G with a Schnyder wood (T_1, T_2, T_3). **b)**–**k)** The construction of a representation with non-rotated (possibly degenerate) T-shaped polygons for the graph G with respect to the given Schnyder wood. **l)** The representation of G obtained with a previously known method [LLY03].

Case (2)**:** As argued in the proof of Lemma 2.1.6, this case can be seen as two applications of Case (1). Each application requires one new grid line; see Figure 2.10. On the other hand, the hill vertex v in $in_3(v_{i+1})$ is a leaf in T_1 and T_2, in particular we have $|in_3(v_{i+1}) \cap L_2| = 1$.

Case (3)**:** First, the steps of the staircase in $in_3(v_{i+1})$ are removed one by one. As evident from Figure 2.14, only new horizontal lines are required for the removal of one step, and the number of such lines equals the number of leaves in T_2 at this step. As soon as there is only one step remaining in $in_3(v)$, Case (2) applies. Hence, two additional grid lines are introduced here. But the hill vertex corresponding to the last staircase step is a leaf in T_2 and not associated with a new horizontal line so far. Thus if k new lines are required, then $|in_3(v_{i+1}) \cap L_2| \geq k - 1$.

From the above case distinction, we get that within step $i+1$ at most $|in_3(v_{i+1}) \cap L_2|+1$ grid lines are newly introduced. Thus after $n-3$ steps, the number of supporting grid lines in the representation is bounded from above by $(2 + (n-3)) + (2 + |L_2|)$, which in turn is at most $(n-1) + \lfloor \frac{2n+1}{3} \rfloor = \lfloor \frac{5n-2}{3} \rfloor$.

Finally, note that it is easy to implement the construction algorithm in linear time. □

2.1.1 Overall Complexity and Number of Segments

Liao, Lu and Yen [LLY03] describe a different way to construct a non-rotated ⊤-representation of a maximally planar graph G based on a given Schnyder wood. Applying their method to the graph G and Schnyder wood (T_1, T_2, T_3) in Figure 2.16 **a)** gives the representation in Figure 2.16 **l)**. Note that the representations in Figure 2.16 **k)** and **l)** look very different, although the underlying Schnyder wood can be read of from either of them. Comparing the two, we note that the representation in Figure 2.16 **l)** has one less grid line, but more segments, and thus higher average polygonal complexity.

It is an interesting and non-trivial task to capture the quality of a rectilinear representation. Let us briefly discuss three such measures, which are in fact closely related. The first is the size of the grid required. In Theorem 2.1.8 above we have shown that the new approach has the same worst-case number of supporting grid lines as the one in [LLY03]. The second measure is given by the number of segments in the representation, as a representation may be considered more complex, or even less readable if it contains a lot of segments. In the proof of Theorem 2.1.8 we actually count the number of segments and use this as an upper bound on the number of supporting grid lines. As the third measure we may take the total number of corners in the representation, which we call the *overall complexity*. The representations considered

so far seem to be designed for a small *maximum* complexity among the polygons involved, i.e., minimizing $\max_{v \in V(G)} |\mathcal{P}(v)|$, where $|\mathcal{P}(v)|$ denotes the complexity of $\mathcal{P}(v)$. In a sense, it is more natural to go for small overall complexity, i.e., minimizing $\sum_{v \in V(G)} |\mathcal{P}(v)|$. For instance, a representation consisting of only quadrangles and a single polygon with, say, 12 corners may be considered less complex than a representation with solely 8-gons. Again, we could bound the overall complexity of the representation constructed in Lemma 2.1.6 by adapting the proof of Theorem 2.1.8. However, we prefer to determine the total number of corners as follows: Every rectilinear polygon has exactly four convex corners more than it has concave corners. In consequence, if a representation with n polygons has k concave corners in total, then its overall complexity is given by $2k + 4n$. Thus it suffices to count the number of concave corners, which turns out to be intimately related to the number of *inner segments*, that is, segments not on the outer boundary of the representation.

Lemma 2.1.9. *In every (not necessarily rectilinear of hole-free) representation of a maximally planar graph on n vertices, the number of inner segments is at least the number of concave corners plus $(n-1)$.*

Proof. We count the number of endpoints of inner segments in a representation Γ. First, the outer boundary of Γ is made of three polygons. Traversing the outer boundary, we encounter three changeovers from one of the three polygons to the next, at each of which lies an endpoint of at least one inner segment. The same holds for every (triangular) face which is represented as a hole in Γ. If a face is not represented as a hole, then the point in the plane corresponding to the face is the endpoint of at least one segment. Lastly, a concave corner of a polygon $\mathcal{P}$ clearly is the endpoint of both segments bounding $\mathcal{P}$ there.

Because there are $2n-5$ inner faces in a maximally plane graph on n vertices, there are at least $2n-2$ endpoints of segments from the faces and additionally twice as many endpoints as there are concave corners. Since every segment has two ends, the statement follows. □

Note that for hole-free rectilinear representations the inequality in Lemma 2.1.9 holds with equality. Hence in the representation constructed with Lemma 2.1.6 the number of supporting grid lines, inner segments, and concave corners is bounded by $\lfloor \frac{5n-2}{3} \rfloor$, $\lfloor \frac{5n-8}{3} \rfloor$, and $\lfloor \frac{2n-5}{3} \rfloor$, respectively. Interestingly, the bound for the overall complexity and the number of segments is worst-case optimal, while the best known [LLY03] lower bound for the number of grid lines is $\lceil \frac{4n}{3} \rceil$.

Theorem 2.1.10. *The non-rotated ⊓-representation in Theorem 2.1.8 has at most $\lfloor \frac{5n-8}{3} \rfloor$ inner segments, and at most $\lfloor \frac{2n-5}{3} \rfloor$ concave corners. Moreover, both bounds are worst-case optimal.*

Proof. We only prove the tight lower bound. The upper bound follows from the reasoning above.

Consider any maximally planar graph on k vertices and insert in each of its $2k-5$ inner faces a new vertex, and connect it by an edge to all three vertices of the face. The resulting graph G has $n = 3k-5$ vertices, $2n-5 = 6k-15$ inner faces, and $2k-5$ non-empty triangles, no two of which have the relation of containment. Consider any rectilinear (not necessarily hole-free) representation Γ of G. The three polygons of a non-empty triangle inscribe a polygonal region of positive area, which clearly has at least four convex corners. Since only three polygons constitute the boundary of this region, at least one of the region's convex corners is a concave corner of a polygon. Since all these $2k-5$ regions are mutually disjoint, there are at least $2k-5 = \frac{6k-15}{3} = \frac{2n-5}{3}$ concave corners in Γ.

The statement about the number of inner segments now follows with Lemma 2.1.9. □

2.2 Representations from Nesting Assignments

This section deals with so-called ⌝-representations and ⬠-representations, which are defined as follows.

Definition 2.2.1.
- A ⌝*-representation* of a maximally planar graph G is a hole-free representation of G with rectilinear polygons of complexity at most 6, whose outer boundary is a rectangle.
- A ⬠*-representation* of a maximally planar graph G is a hole-free representation of G with convex polygons of complexity at most 5, whose outer boundary is a rectangle.

Recall that a *non-empty triangle* (also known as a *complex triangle*) in a plane graph G is a set of three pairwise adjacent vertices, that do not form an inner face in G. Note that if G has an outer triangle (and at least one inner vertex), then this triangle is non-empty although it is not a separating triangle. The following definition is due to Sun and Sarrafzadeh [SS93].

Definition 2.2.2. A *nesting assignment* of a plane graph G is an assignment of every non-empty triangle in G to one of its three vertices, such that any two triangles that are assigned to the same vertex have the relation of containment in the embedding of G.

Sun and Sarrafzadeh [SS93] show that a maximally planar graph has a ◱-representation if and only if it admits a nesting assignment, i.e., it has a plane embedding that admits a nesting assignment. The necessity of this characterization is actually not very difficult and we present a proof in Theorem 2.2.4. The sufficiency is argued constructively by Sun and Sarrafzadeh [SS93], i.e., they describe an algorithm that given a nesting assignment constructs a ◱-representation in $\mathcal{O}(n^2)$ time, where n is the number of vertices in the graph. Here we give an algorithm for the same problem, having the following three main advantages:

- It is significantly simpler.
- It runs in $\mathcal{O}(n)$ time.
- A modified version of it constructs a ⬠-representation, rather than a ◱-representation.

Lemma 2.2.3. *Every maximally planar graph admitting a nesting assignment has a ◱-representation and a ⬠-representation. Moreover, given the assignment, either representation can be computed in $\mathcal{O}(n)$ time, where n denotes the number of vertices in G.*

Proof. We consider G to be given with the embedding and the nesting assignment. Recall that the vertices in the separation-tree $\mathcal{T}_G$ correspond to the non-empty triangles in G. We begin by imposing an additional condition on the nesting assignment.

Claim 1. A nesting assignment can be modified in linear time into a nesting assignment, in which every vertex is assigned to a (possibly empty) set of triangles, which appear *consecutively* in $\mathcal{T}_G$.

Proof of Claim 1. We can greedily reassign the triangles in the following way. Let v be any vertex that is assigned by a set S_v of at least two triangles that are not consecutive in $\mathcal{T}_G$. Since the assignment is nesting, any two triangles in S_v have the ancestor-descendant relation, that is, S_v is a subset of a descending path in $\mathcal{T}_G$. Let $P_v \supset S_v$ be the unique such path that has both endpoints in S_v. We assign every triangle on P_v to the vertex v, disregarding its former assignment. It is not difficult to see that the new assignment is nesting. Moreover, if a vertex was assigned by a set of consecutive triangles before the reassignment, it is so afterwards. Hence repeating this procedure at most once for every vertex finally gives a consecutive nesting assignment.

Choosing always the root as the vertex v in this procedure, and recursively treating the connected components in $\mathcal{T} \setminus P_v$, gives a linear-time implementation. △

Now, consider the embedded graph G with outer triangle Δ_0 and equipped with a consecutive nesting assignment. We construct both representations along the

separation-tree $\mathcal{T}_G$, i.e., we start by representing the level-0 subgraph $G[0]$ of G. Having represented the level-i subgraph $G[i]$ for $i \geq 0$, we integrate one after another every subgraph G_Δ of G with Δ having depth $i+1$ in $\mathcal{T}_G$. This way we have represented $G[i+1]$ and iterate this procedure. We maintain the following invariant on the ⌝-representation $\Gamma_i^{(1)}$ of $G[i]$, for $i \geq 0$.

- $\Gamma_i^{(1)}$ is a ⌝-representation of $G[i]$ with a *rectangular* hole for every non-empty triangle on level $i+1$ in $\mathcal{T}_G$.
- Every side of such hole is due to one participating vertex, except for the vertex that is assigned to the triangle which contributes two sides.
- Every vertex in $G[i]$ that is not assigned by a triangle in $G[i]$ is represented by a rectangle.

And similarly for the ⬠-representation $\Gamma^{(2)}$ of $G[i]$.

- $\Gamma_i^{(2)}$ is an ⬠-representation of $G[i]$ with a *triangular* hole for every non-empty triangle on level $i+1$ in $\mathcal{T}_G$.
- Every side of such hole is due to one participating vertex.
- Every vertex in $G[i]$ that is not assigned to a triangle in $G[i]$ is represented by a triangle or a convex quadrangle.

For convenience, we consider the outer triangle $\Delta_0 = \{v_1, v_2, v_n\}$ to be the graph $G[-1]$ and define the representations $\Gamma_{-1}^{(1)}$ and $\Gamma_{-1}^{(2)}$ for $G[-1]$, as depicted in the left of Figure 2.17. A rectangular, respectively triangular, hole for a non-empty triangle Δ satisfying the above requirements is called a *frame* and will be depicted as in the right of Figure 2.17.

Having representations $\Gamma_i^{(1)}$ and $\Gamma_i^{(2)}$ of $G[i]$ we consider any non-empty triangle Δ on level $i+1$ in $\mathcal{T}_G$ and label its vertices by v_1, v_2 and v_n in such a way that Δ is assigned to v_n. Then consider G_Δ, that is, the maximally planar subgraph of G on v_1, v_2, v_n and at least one fourth vertex from inside Δ, which contains no separating triangle. We build a □-representation of G_Δ, distinguishing two cases. Such a representation can be found in $\mathcal{O}(|V(G_\Delta)|)$.

In case no inner triangle in G_Δ is assigned to v_n we build any □-representation of $G_\Delta \setminus (v_1, v_n)$. As usual denote the fourth outer vertex besides v_1, v_2 and v_n by v_{n-1}. We place this □-representation (suitably stretched) into the frame for Δ as exemplified in Figure 2.18. For the ⬠-representation, we first apply an affine transformation to Γ_2 so that there is a right angle between v_1 and v_2. (Note that this affects neither the convexity nor the complexity of the polygons represented.) We then place the □-representation of G_Δ with the top right corner of $\mathcal{R}(v_{n-1})$ onto the boundary of $\mathcal{P}(v_n)$ as depicted in Figure 2.18. Afterwards we remove the two

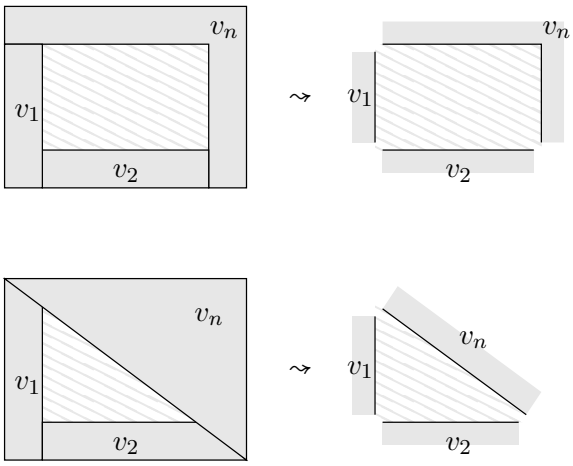

Figure 2.17: The L-representation $\Gamma^{(1)}_{-1}$ and 5-gon representation $\Gamma^{(2)}_{-1}$ of $G[-1] = \Delta_0 = \{v_1, v_2, v_n\}$ and the corresponding frames for Δ_0, which is assigned to v_n.

triangular holes by extending $\mathcal{R}(v_{n-1})$ upwards and the rectangles of all neighbors of v_n to the right.

In case some inner triangle $\{u, v, v_n\}$ is assigned to v_n, consider the counterclockwise order σ of neighbors around v_n starting with v_1, ending at v_2, and assume that u comes (immediately) before v, i.e., $\sigma = (\sigma_1 = v_1, \dots, \sigma_i = u, \sigma_{i+1} = v, \dots, \sigma_k = v_2)$. We then split the vertex v_n into two, one denoted by v_{n-1} and the other again by v_n. We connect v_{n-1} and v_n to $\sigma_1, \dots, \sigma_i$ and $\sigma_{i+1}, \dots, \sigma_k$, respectively. Next we add the edge (v_{n-1}, v_n) and another vertex v^*, with an edge to v_{n-1}, v_n, u, and v. Denote the resulting graph by G'_Δ. See Figure 2.19 for an example.

The graph G'_Δ is 4-connected because G_Δ was so. Indeed, one easily checks that the number of triangles that contain a fixed edge from $G_\Delta \cap G'_\Delta$ did not increase, and that any edge in $G'_\Delta \setminus G_\Delta$ is contained in at most two triangles as well. Hence G'_Δ has a $\square$-representation, which we place (again suitably stretched) into the frame for Δ as exemplified in Figure 2.20. For the $\square$-representation and $\pentagon$-representation we move the top right corner of $\mathcal{R}(v^*)$ onto the the reflex corner and the boundary of $\mathcal{P}(v_n)$, respectively. Afterwards we remove the rectangles $\mathcal{R}(v_{n-1})$ and $\mathcal{R}(v_n)$, and for the $\pentagon$-representation remove the holes by extending each rectangle of a neighbor of v_n. Finally, we remove the rectangle $\mathcal{R}(v^*)$, which leaves a rectangular, respectively triangular, hole whose sides are constituted by $\mathcal{P}(u)$, $\mathcal{P}(v)$, and $\mathcal{P}(v_n)$. In other words, the resulting representation features a frame for the non-empty triangle $\{u, v, v_n\}$.

One can easily check that in both cases, whether v_n is assigned by some triangle in G_Δ or not, we have extended the $\square$-representation $\Gamma^{(1)}_i$ and the $\pentagon$-representation $\Gamma^{(2)}_i$ of $G[i]$ by the vertices and edges in G_Δ.

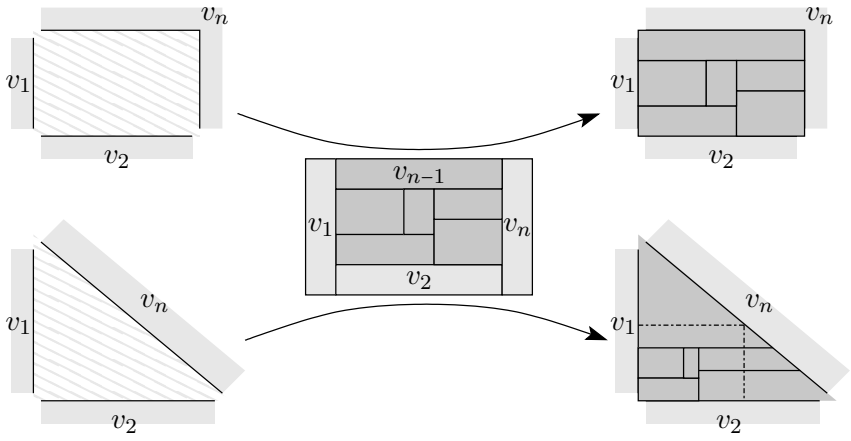

Figure 2.18: Placing a rectangle-representation of G_Δ into the frame for Δ.

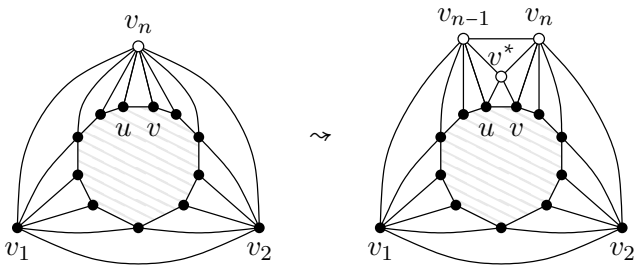

Figure 2.19: The maximally planar graph G_Δ with triangle $\Delta = \{u, v, v_n\}$ and the resulting graph G'_Δ after splitting v_n.

It remains to identify a frame for every inner facial triangle in G_Δ that is non-empty in G. We do this by a local modification of the representations $\Gamma_i^{(1)}$ and $\Gamma_i^{(2)}$. Note that for every inner facial triangle $\{u, v, w\}$ in G_Δ there is a unique point p_{uvw} in the plane that is shared by the borders of $\mathcal{P}(u)$, $\mathcal{P}(v)$, and $\mathcal{P}(w)$. Exactly two of these polygons have a corner at p_{uvw}. Assume the triangle $\{u, v, w\}$ is non-empty in the graph G, i.e., it is assigned to one of its vertices, say u. Since the nesting assignment is consecutive, we have that u is not assigned to any triangle in $G[i-1]$. Thus according to the invariant, $\mathcal{P}(u)$ is a rectangle in $\Gamma_i^{(1)}$ and a triangle or a convex quadrangle in $\Gamma_i^{(2)}$. We transform the representation, increasing the complexity of $\mathcal{P}(u)$, so that a frame for $\{u, v, w\}$ is formed. Figure 2.21 illustrates this local modification in case $\mathcal{P}(u)$ has a corner at p_{uvw} (top row) or not (middle row for $\Gamma_i^{(1)}$, bottom row for $\Gamma_i^{(2)}$). If $\mathcal{P}(u)$ goes straight at p_{uvw} let r be a ray starting at p_{uvw} and supporting the corresponding $\mathcal{P}(u)$-border. We introduce a new convex $\mathcal{P}(u)$-corner on r next to p_{uvw}. In case of the ⌝-representation $\Gamma_i^{(1)}$ we translate the entire representation on r slightly away from $\mathcal{P}(u)$, introducing a new concave $\mathcal{P}(u)$-corner and a rectangular

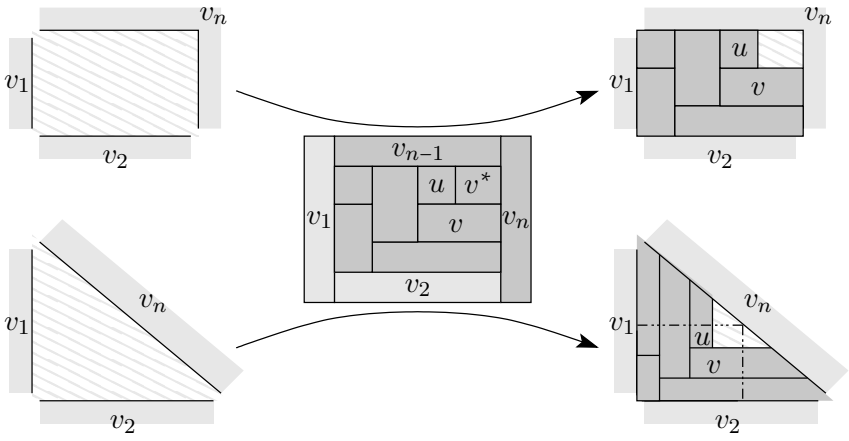

Figure 2.20: Placing a rectangle-representation of G'_Δ into the frame for Δ.

frame for the triangle $\{u, v, w\}$. In case of the $\pentagon$-representation $\Gamma_i^{(2)}$ we change the slope of r slightly so that all polygons on the same side of r as $\mathcal{P}(u)$ are truncated and all polygons on the other side are extended. This way a triangular frame for the triangle $\{u, v, w\}$ is introduced. Note that the runtime needed to represent G_Δ and to introduce the frames in G_Δ is $\mathcal{O}(|V(G_\Delta)|)$, for both the rectilinear and the pentagonal case.

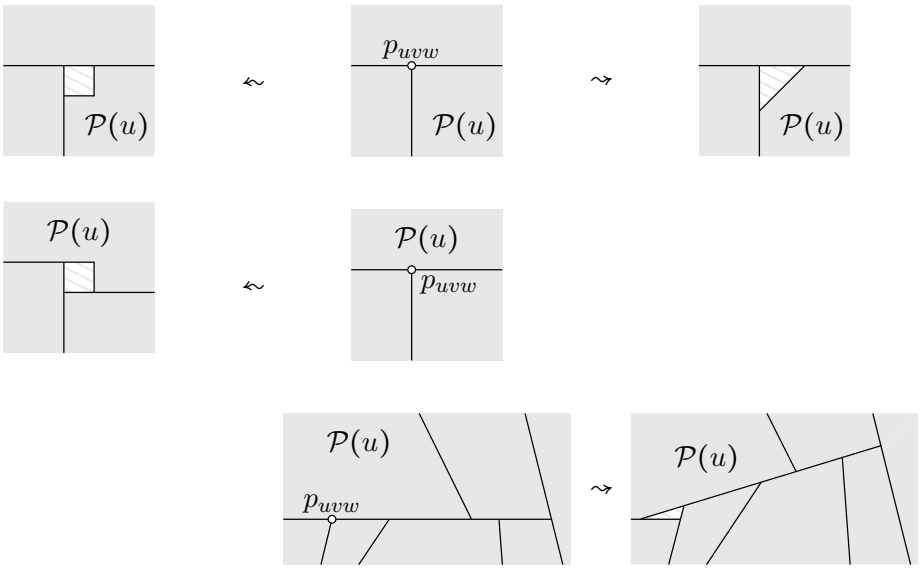

Figure 2.21: A local transformation reveals a frame for the triangle $\{u, v, w\}$.

Starting with $\Gamma_i^{(k)}$, $k = 1, 2$, we perform the above procedure for every non-empty triangle in $G[i]$. Note that a local modification is done only once for every vertex, since it is assigned to only one facial triangle in $G[i]$. It follows that the result is a

⌜-representation $\Gamma^{(1)}_{i+1}$ and a ⬠-representation $\Gamma^{(2)}_{i+1}$ of $G[i+1]$, which both satisfy our invariant. Iterating until the entire graph G is represented then proves the lemma. □

For completeness let us restate the following theorem.

Theorem 2.2.4 ([SS93])**.** *A maximally planar graph has a ⌜-representation if and only if it admits a nesting assignment.*

Proof. The "if"-part is proven in Lemma 2.2.3. It remains to prove the "only if"-part. Therefore consider a ⌜-representation and note that every non-empty triangle $\{u, v, w\}$ encloses a region R_{uvw} with at least four sides. By the pigeonhole principle at least one vertex of the triangle, say u, contributes to two sides of R_{uvw}. Consequently, $\mathcal{P}(u)$ is an L-shape and its concave corner is a corner of this region. It is not difficult to check that the assignment $\{u, v, w\} \to u$ is nesting. □

By Theorem 2.2.4, nesting assignments are necessary for ⌜-representations. However, they are not necessary for ⬠-representations. For example, consider the maximally planar graph G in Figure 2.22. It can be seen as a 6-vertex maximally planar graph with an additional vertex stacked into each inner face. This way, every inner face is non-empty and no two of them have the relation of containment. Since there are seven such faces and only six original vertices, there is no nesting assignment for G. This graph already has been presented as a graph with no nesting assignment [SS93]. On the other hand, the right of Figure 2.22 shows a ⬠-representation of G.

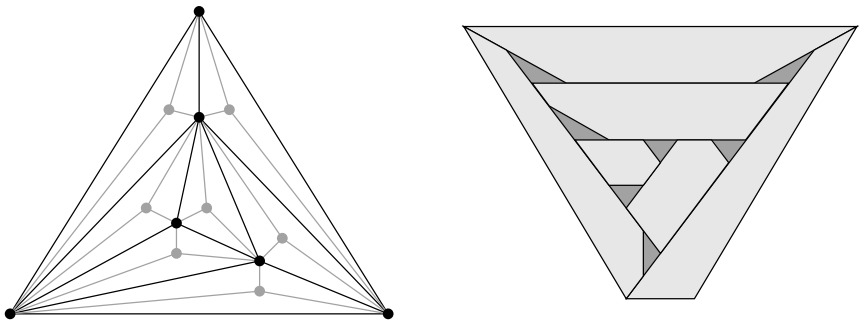

Figure 2.22: A maximally planar graph, that does not admit a nesting assignment and a representation of it with convex polygons of complexity at most five.

Of particular interest is the fact that Lemma 2.2.3 can be applied to Hamiltonian maximally planar graphs, which I was told by M. J. Alam. This especially shows that these graphs admit a ⬠-representation, which for instance was suspected by S. G. Kobourov.

Corollary 2.2.5. *Every Hamiltonian maximally planar graph has a* $\square$*-representation and a* $\pentagon$*-representation.*

Proof. Consider a Hamiltonian cycle $v_1, \ldots, v_n$ in the embedded graph G. We assign every non-empty triangle to one boundary vertex v_i, for which $v_{(i+1) \bmod n}$ lies inside. Then the set of triangles that is assigned to the same vertex v_i shares the same inner vertex $v_{(i+1) \bmod n}$. In other words, these triangles are pairwise nested and we can apply Lemma 2.2.3. □

2.3 Lower Bounds on the Complexity

In this section we present a method for proving lower bounds on the polygonal complexity of representations. We start with the crucial lemma (Lemma 2.3.1), in which we give lower bounds on the number of *convex* corners in a representation at hand. Afterwards, we apply Lemma 2.3.1 to obtain tight lower bounds on the polygonal complexity of certain graph classes, which means that we give a particular graph from the graph class of interest and argue on the polygonal complexity of *any* of its representations.

Lemma 2.3.1. *Let* Γ *be a representation of some connected plane graph* G*,* f *be a face of* G*,* $\mathcal{P}^*$ *the outer boundary of* Γ*, and* $\deg(\mathcal{P}^*)$ *its complexity.*

(i) If f *is an inner face represented as a hole, then* f *contains at least* $\deg(f)$ *convex corners of polygons in* Γ*.*

(ii) If f *is an inner face* not *represented as a hole, then* f *contains at least* $\deg(f)-1$ *convex corners of polygons in* Γ*.*

(iii) If f *is the outer face of* G*, then* f *contains at least* $\deg(f) + \deg(\mathcal{P}^*)$ *convex corners of polygons in* Γ*.*

Proof. All the statements of the lemma rely on the following easy observation.

Claim 1. If p is a corner of some polygon in Γ then p is the *convex* corner of all but possibly one polygon containing p.

If an inner face f is *not* represented as a hole, then it corresponds to a point p in the plane. In fact, p is the intersection of the polygons of all vertices incident to f, and a corner of at least one of it. Now (ii) follows from Claim 1.

Now suppose the face f is not just a point, i.e., it is an inner face not represented as a hole or the outer face. We trace the polygonal boundary of f in Γ, starting at a point on the boundary of only one polygon. Whenever we come across new polygons,

we encounter as many convex corners as new polygons, which already proves (i); see Figure 2.23. Moreover, at every point which is straight or concave w.r.t. f we encounter one additional convex corner. See Figure 2.23 again. Since $\deg(\mathcal{P}^*)$ points are concave w.r.t. the outer face, this finally proves (iii). □

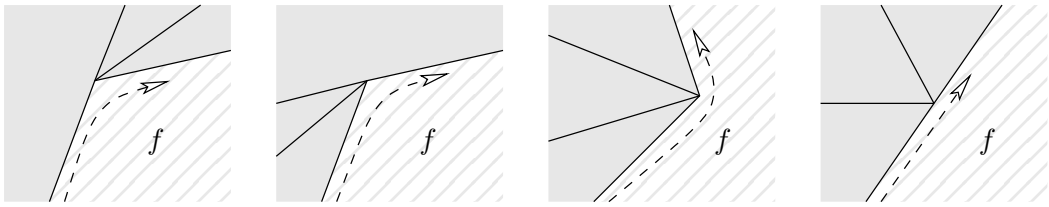

Figure 2.23: The possibilities of coming across new polygons when tracing the boundary of a face f, which is not represented as just a point in Γ.

Lemma 2.3.2. *Each of the following holds.*

(i) Every representation of any maximally planar graph on $n > 4$ vertices requires polygonal complexity at least 4.

(ii) There is a Hamiltonian maximally planar graph of which every representation requires polygonal complexity at least 5.

(iii) There is a planar 3-tree of which every representation requires polygonal complexity at least 6.

Proof. All the statements of the lemma follow from Lemma 2.3.1.

Let G be any maximally planar graph on $n > 4$ vertices. It has $2n - 5$ inner faces and one outer face, each of degree 3. Let Γ be any representation of G. Then by Lemma 2.3.1 there are at least two convex corners at every inner face and at least six convex corners at the outer face. Hence there are at least $2(2n - 5) + 6 = 4n - 4$ convex corners in total. The average number of convex corners per polygon is $(4n-4)/n$, which is strictly more than 3 if $n > 4$. This proves (i). We remark that the unique maximally planar graph on four vertices admits a representation by polygons of complexity 3; see Figure 2.24.

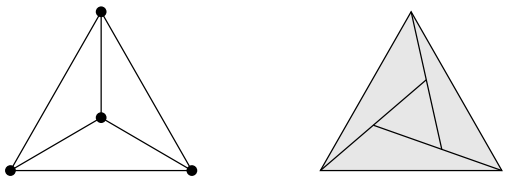

Figure 2.24: The unique maximally planar graph on four vertices and a representation by polygons of complexity 3.

Let G be any Hamiltonian graph on $n > 4$ vertices and $v_1, \ldots, v_n$ be a Hamiltonian cycle. It is not difficult to see that every edge on the Hamiltonian cycle can be assigned to one incident inner face, so that no face is assigned more than once. We insert a new vertex into every face that is assigned by some edge and connect the new vertex by an edge to each of the three vertices of the face. Rerouting every edge on the Hamiltonian cycle through the new vertex of the corresponding face gives a Hamiltonian cycle in the new maximally planar graph G'. See Figure 2.25 for an example. Let Γ' be any representation of G'. The polygons of the n vertices

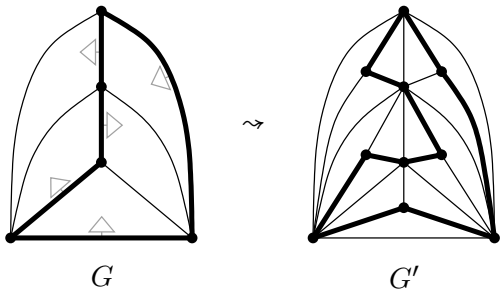

Figure 2.25: Making the Hamiltonian graph G' out of the Hamiltonian graph G. The one-to-one assignment of edges on the Hamiltonian cycle to incident faces is indicated by small arrows.

originally belonging to G form a representation Γ of G in which every inner face containing a G'-vertex is represented as a hole. There are exactly n such faces, $n-5$ further inner faces, and one outer face, each of degree 3. Applying Lemma 2.3.2 to Γ gives that the polygons corresponding to the original vertices contain at least $3n+2(n-5)+6 = 5n-4$ convex corners. The average number of convex corners among these polygons is $(5n-4)/n$, which is strictly more than 4 if $n > 4$. This proves (ii).

Finally let G be any planar 3-tree on $n > 9$ vertices. We obtain from G a planar 3-tree G' by inserting a new vertex into every inner face of G and connecting it by an edge to each of the incident vertices. Let Γ' be any representation of G' and Γ be the representation of G given by removing the polygon for every vertex in $G' \setminus G$. In Γ every inner face is represented as a hole. There are $2n-5$ such faces and one outer face, each of degree three. Applying Lemma 2.3.2 to Γ gives that the polygons corresponding to the original vertices contain at least $3(2n-5)+6 = 6n-9$ convex corners. The average number of convex corners among these polygons is $(6n-9)/n$, which is strictly more than 5 if $n > 9$. This proves (iii). □

Let us close this chapter by putting the obtained results into the picture of what is known about side contact representations of maximally planar graphs.

We start with non-rectilinear representations: By Lemma 2.3.2 a maximally planar graph admits a representation of complexity 3 if and only if it consists of at most 4 vertices. From Theorem 1.3.2 we get that every 4-connected maximally planar graph, i.e., one without separating triangles, admits a representation of complexity at most 4. In Lemma 2.2.3 and Corollary 2.2.5 we have identified the first non-trivial class of maximally planar graphs that are not 4-connected but admit a $\pentagon$-representation. Finally, it is known that *every* maximally planar graph admits a representation of complexity at most 6 [GHKK10]. However, no complete characterization of those maximally planar graphs that admit a representation of complexity at most 4 or 5 is known. When considering general planar graphs, a characterization for graphs admitting representations with any complexity $k \geq 3$ is open.

The situation with *rectilinear* representations is different: By Theorem 1.3.2 a near-triangulation (a maximally planar graph minus one outer edge) admits a $\square$-representation if and only if it is 4-connected, i.e., it does not contain separating triangles. By Theorem 2.2.4 a maximally planar graph admits a $\sqsupset$-representation if and only if it admits a nesting assignment. Finally, by Theorem 2.1.8 *every* maximally planar admits a $\sqsupset\sqsubset$-representation. However, we know almost nothing about rectilinear representations of general planar graphs.

Chapter 3

Prescribing Areas of Polygons – Cartograms

In this chapter we deal with a variant of the representation problem considered in the previous chapter. Given a vertex-weighted planar graph, we ask for a representation of it, in which the area of the polygon corresponding to a vertex is proportional to its weight. Most of the work presented here is joint work with M. J. Alam, S. G. Kobourov, and S. Felsner. Some results, in particular Theorem 3.1.7, Theorem 3.2.2, and Lemma 3.2.3, already appear in [ABF+11c].

Definition 3.0.3. For a planar graph $G = (V, E)$ and a weight function $w : V \to \mathbb{R}^+$, a side contact representation Γ of G is called a *cartogram* if the area of $\mathcal{P}(v)$ equals $w(v)$, for every $v \in V$.

Cartograms have practical applications in cartography, geography, and sociology, but also in VLSI layout, and floor planning. For instance, rectangular cartograms, i.e., those with an underlying $\square$-representation, are often used as schematized versions of geographic maps where the size of every region represents a geographic variable such as population, highway kilometers, or literacy. There is a rich and beautiful garden of cartograms used for visualizing proportional relations between different objects under the same measure. For example Figure 3.2 depicts the world map in which countries are distorted so that their areas represent the GDP (gross domestic product) in the year 2000. In Figure 3.1 the news of the day are visualized as rectangular boxes. The adjacencies somehow represent related news and their time stamps. Moreover, the size of every rectangle is chosen to represent the importance of the corresponding news. Let us mention the beautiful overview of Bettina Speckmann (http://www.win.tue.nl/~speckman/Cartograms/) on rectangular cartograms, i.e., cartograms with a rectangle per vertex. It provides many interesting examples, as well as some theoretical background and links. Figure 3.3 is taken from there.

Figure 3.1: A map of the news from August 1st, 2011. The area is supposed to represent importance.
Source: http://newsmap.jp/

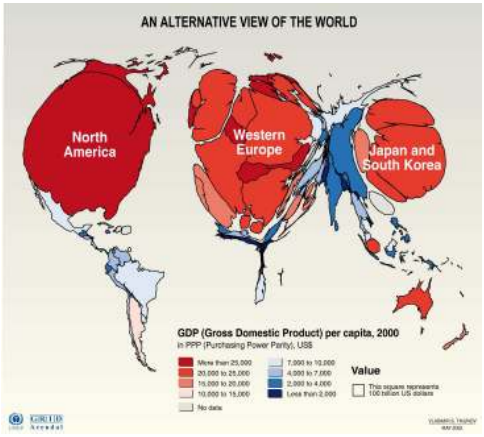

Figure 3.2: A cartogram of the world illustrating the Gross Domestic Product per capita in 2000.
Source: http://maps.grida.no/

Cartograms, in this case rectangular cartograms, were first considered by Raisz in 1934 [Rai34]. Heilmann *et al.* [HKPS04] and van Kreveld and Speckmann [vKS07] compute rectangular cartograms with possibly disturbed adjacencies and small errors on the weights, so-called *cartographic errors*. Within this thesis, we consider cartograms as defined in Definition 3.0.3, i.e., vertices being represented by general polygons and with neither disturbed adjacencies nor cartographic errors. Moreover, we address the "weighted version" of Question 2.0.5, which is the following.

Question 3.0.4. Given a graph class $\mathcal{G}$, what is the minimum number k, such that every graph $G \in \mathcal{G}$ equipped with a weight function $w : V(G) \to \mathbb{R}^+$ has a cartogram w.r.t. w with polygonal complexity at most k? What if we require the underlying representation to be hole-free and/or rectilinear?

The first natural question is, whether the weight function does affect the complexity at all, or whether the answer to Question 3.0.4 is actually finite. De Berg *et al.* [dBMS06] were the first to show that finite polygonal complexity is always enough, i.e., they prove that every vertex-weighted planar graph admits a cartogram of complexity at most 40. This number was later [KN07] reduced to 34 and then to 12 [BRV11]. Recently, this was further reduced to 10 [ABF+11a]. In this chapter we lower this once more to 8, which is clearly best-possible. Somehow surprisingly, this means that for general planar graphs, imposing the area of every polygon does not raise the worst-case complexity beyond 8. However, as we will see later, for many graphs, there is a weight function so that every cartogram requires a higher complexity than its best (unweighted) representation. In Table 3.1 the best known

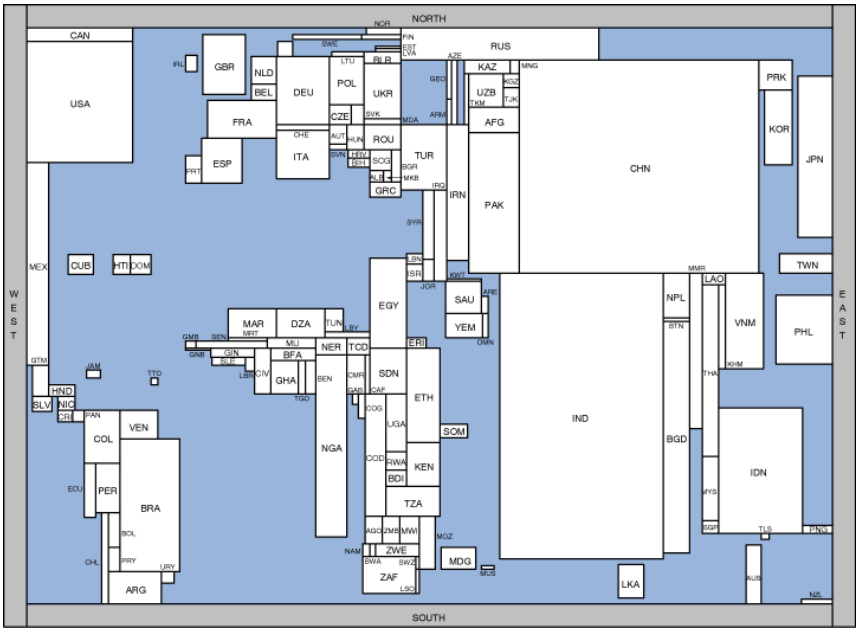

Figure 3.3: A cartogram of the world illustrating the population per country in 2006. **Source:** http://www.win.tue.nl/~speckman/Cartograms/

lower and upper bounds for Question 3.0.4 are listed. For comparison, we included the worst-case complexity in the unweighted case, i.e., Question 2.0.5, in the columns labeled 'NW'.

This chapter is organized as follows.

Section 3.1: We review the notion of area-universal layouts, as defined by Eppstein *et al.* [EMSV09] for rectangular layouts. We extend this notion to the non-rectangular, and even non-rectilinear case. We present our main result, saying that every maximally planar graph admits an area-universal rectilinear layout with polygonal complexity at most 8, c.f. Lemma 3.1.6. This implies the existence of a rectilinear cartogram with complexity 8, improving the previously best known value 10 [ABF$^+$11a]. Moreover, this is worst-case optimal.

Section 3.2: This section deals with rectilinear cartograms of Hamiltonian maximally planar graphs. We give an alternative area-universal layout with complexity at most 8 for those graphs which, in contrast to the general case, allows us to compute the actual cartogram in linear time. Moreover, we give a matching lower bound construction for this setting. As opposed to the unweighted case, this shows that there exist worst-case examples that are Hamiltonian. In Subsection 3.2.1 we define one-sided Hamiltonian cycles and present their close

Graph Class	Non-Rectilinear						Rectilinear					
	Holes			Hole-Free			Holes			Hole-Free		
	NW	LB	UB	NW	LB	UB	NW	LB	UB	NW	LB	UB
Maximally Outer-Planar	3	4	4	3	4	4 [ABF+11b]	4	6	6	6	6	6 [ABF+11a]
4-Connected Near-Triang.	4	4	7	4	4	8	4	6 [vKS07]	8	4	6	8
Hamiltonian Max. Planar	5	6 Lem. 3.3.2	7	5	6	8	6	8 Lem. 3.2.3	8	6	8	8 Thm.3.2.2
Planar 3-Tree	6	7 Lem. 3.3.2	7	6	7	7 Thm.3.3.4	8	8 Lem. 3.2.3	8	8	8	8
Max. Planar	6	7	7 Thm. 3.3.3	6	7	8	8	8	8	8	8	8 Thm. 3.1.7

Table 3.1: Summary of lower bounds and upper bounds for the polygonal complexity required for cartograms of all graphs within some graph class. (NW = no weights, LB = lower bound, UB = upper bound)

relation to Schnyder woods. We show how to compute rectilinear cartograms with complexity 6 for graphs admitting one-sided Hamiltonian cycles.

Section 3.3: We provide lower bounds on the complexity of (not necessarily rectilinear) cartograms for several graph classes. To this end, we combine an enforcing of concave corners with the counting argument presented in Section 2.3. We obtain the first lower bounds that are better than those from the unweighted case. If the cartogram is allowed to have holes, these bounds match existing upper bounds [ABF$^+$11b] for maximally planar graphs. In Subsection 3.3.1 we show that the lower bound of 7 obtained this way for hole-free cartograms of planar 3-tree is best-possible. We give a simple linear-time algorithm for constructing such a cartogram. This remains the only non-trivial case of a hole-free cartogram whose complexity is strictly less than in the rectilinear case.

Section 3.4: We discuss several open problems concerning rectilinear and non-rectilinear cartograms of 4-connected maximally planar graphs. We present two conjectures and point out possible directions to their solution.

3.1 Area-Universal Layouts

One approach for computing cartograms (or just prove their existence) is the following two-step procedure. First, define a representation Γ of the given graph G, which is independent of the given weight function $w : V(G) \to \mathbb{R}^+$. And second, prove that every possible weight function can be realized with a cartogram, whose underlying representation is combinatorially equivalent to Γ. Finally, describe a way of actually computing the cartogram based on the representation Γ and the weights w. Let us define this concept more formally.

Definition 3.1.1. A *segment* in a representation Γ, is a maximal line segment contained in the union of all polygon boundaries. The *layout* Λ of Γ is the set of its segments together with their touching/crossing relation and incidence order in Γ.

The incidence order of segments captures the clockwise order of those segments that are touching/crossing a fixed segment, as well as the order of those points on a segment that are shared by another segment. This way, every representation has only one layout, but there are many representation which have the same layout. If a representation Γ has layout Λ, we say that Γ *realizes* Λ, and that Γ is a *realization* of Λ. For example see Figure 3.4. The two representations in Figure 3.4 **a)** and **b)** realize the same layout. The representations in Figure 3.4 **c)** and **d)** realize different layouts, and both are different from the one in Figure 3.4 **a)**.

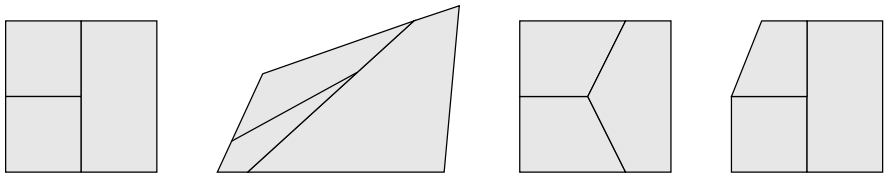

Figure 3.4: The two representations in **a)** and **b)** have the same layout. The layouts in **c)** and **d)** are different and different from the one in **a)**.

Every realization of a fixed layout Λ represents the same planar graph G. Hence, we call Λ a *layout for the graph* G. Moreover, the complexity of a polygon $\mathcal{P}(v)$ is the same in every realization of Λ. Lastly, the property of being hole-free is invariant under all representations of a fixed layout, while the property of being rectilinear is not. We proceed with the crucial definition, which first appeared for rectangular layouts only [EMSV09].

Definition 3.1.2. A layout Λ of a planar graph $G = (V, E)$ is *area-universal*, if for every weight function $w : V \to \mathbb{R}^+$ there is a cartogram Γ_w w.r.t. w of G, that realizes Λ.

Figure 3.5 shows two realizations of the same layout for K_4, the complete graph on four vertices, with different areas for each polygon. Indeed, one can argue that this layout is area-universal, i.e., whatever four positive weights are assigned to the vertices of K_4, there is a cartogram w.r.t. these weights, which essentially looks like the ones in Figure 3.5.

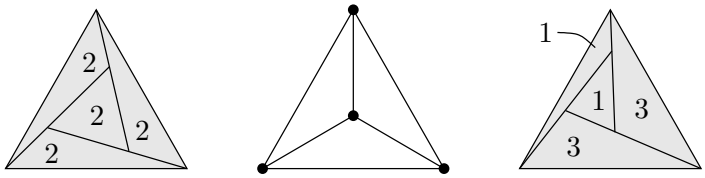

Figure 3.5: Two cartograms for the same planar graph but different weight function.

For the next theorem, we define a layout Λ to be *rectilinear area-universal* if for every weight function there is a *rectilinear* cartogram, that realizes Λ. In other words, we additionally prescribe the angle between any two touching (and crossing) segments. This way the degrees of freedom are restricted to the choice of a length for every sub-segment, i.e., the subset of a segment between two corners. Eppstein *et al.* [EMSV09] call a layout *one-sided* if every non-boundary segment is the side of a rectangle. (Note that in every $\square$-representation at least two of the boundary segments

are not a side of a rectangle.) Equivalently, for every segment s there are segments ending on s from at most one side. Figure 3.6 shows a graph G and a one-sided layout of G on the left. The layout of G in the right of Figure 3.6 is not one-sided, since the segment containing the side contact between $\mathcal{R}(v)$ and $\mathcal{R}(w)$ has an endpoint of one segment from either side. We remark that the left layout arises from the right one by a flip of an essential cycle of length 8 (c.f. Subsection 1.3).

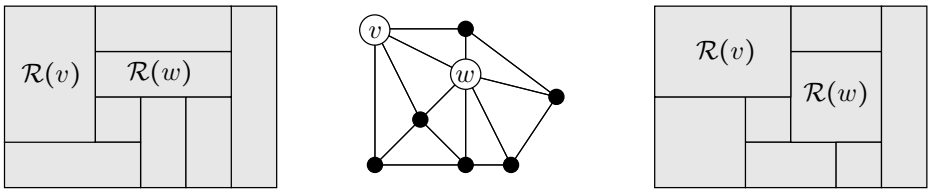

Figure 3.6: A graph with a layout of it that is one-sided on the left and one that is not on the right.

Theorem 3.1.3 ([EMSV09])**.** *A layout of a $\square$-representation is rectilinear area-universal if and only if it is one-sided.*

Theorem 3.1.3 does *not* imply that if some layout is not one-sided, then the graph G it is representing does not admit a rectilinear cartogram of complexity 4. There could be another layout for G which is one-sided. Or even if some G does not have any one-sided layout, there could still be one layout Λ_w for every possible weight function w, such that one realization of Λ_w respects these weights. However, such an argumentation does not follow the approach mentioned in the beginning and is possibly difficult to carry out.

Van Kreveld and Speckmann [vKS07] provide a simple example of a vertex-weighted graph G that does not admit a rectilinear cartogram of complexity 4.

Lemma 3.1.4 ([vKS07])**.** *Consider the graph $G = (V, E)$ in Figure 3.7. Define $w(a) = w(b) = D$ and $w(v) = \delta$ for $v \in V \setminus \{a, b\}$, where $D > \delta$. Then every rectilinear cartogram of G with weight function w requires at least one 6-sided polygon.*

Sketch of proof. W.l.o.g. we fix the layout for the outer four vertices. (This can always be done for rectangular layouts.) Now, G has only two layouts, depicted in the right of Figure 3.7, neither of which is one-sided. Now we argue that in both layouts the larger area of $\mathcal{R}(g)$ and $\mathcal{R}(h)$ is at least as large as the smaller area of $\mathcal{R}(a)$ and $\mathcal{R}(b)$. If in the left layout the area of $\mathcal{R}(g)$ is smaller than the area of $\mathcal{R}(a)$, then $\mathcal{R}(a)$ is wider than $\mathcal{R}(g)$ since both rectangles have the same height. Similarly, if the area of $\mathcal{R}(h)$ is smaller than the area of $\mathcal{R}(b)$, then $\mathcal{R}(b)$ is wider than $\mathcal{R}(h)$. Thus if both cases apply, then $\mathcal{R}(a)$ and $\mathcal{R}(b)$ would have a horizontal side contact,

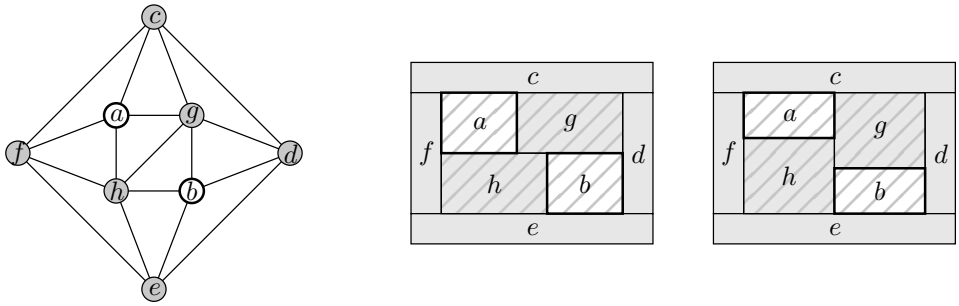

Figure 3.7: A graph G with grey and white vertices weighted δ and D, respectively. If $D > \delta$ then there is no cartogram for G with only rectangles.

which is not specified in the layout. The argumentation for the right layout is the same with the roles of a and b exchanged.

Hence if the area D of $\mathcal{R}(a)$ and $\mathcal{R}(b)$ is greater than the area δ of $\mathcal{R}(g)$ and $\mathcal{R}(h)$, there is no way to realize the given weights. □

Theorem 3.1.3 is important, not only because it gives a sufficient condition for the existence of a rectilinear cartogram with complexity 4, but also because of the next lemma. A layout Λ^* is a *refinement* of another layout Λ if every segment of Λ is contained in a segment of Λ^* and the restriction of Λ^* to the segments of Λ yields the layout Λ. In other words, every polygon $\mathcal{P}$ in Λ is subdivided into a set of polygons $\{\mathcal{P}_1, \dots, \mathcal{P}_k\}$ in Λ^*, while keeping the combinatorics of the boundaries of every $\mathcal{P} = \mathcal{P}_1 \cup \dots \cup \mathcal{P}_k$ intact. Figuratively speaking, a refinement Λ^* of Λ arises from Λ by adding some segments and/or extending existing segments.

Lemma 3.1.5. *A layout Λ is (rectilinear) area-universal if some refinement Λ^* of it is (rectilinear) area-universal.*

Proof. Let G and G^* denote the planar graph represented by Λ and Λ^*, respectively. Since Λ^* is a refinement of Λ, every polygon $\mathcal{P}(v)$ in Λ corresponds to a set of polygons $\mathcal{P}(v^1), \dots, \mathcal{P}(v^{k_v})$ in Λ^*. Every realization Γ^* of Λ^* gives a realization Γ of Λ, in which the union of $\mathcal{P}(v^1), \dots, \mathcal{P}(v^{k_v})$ is the polygon $\mathcal{P}(v)$. Given a weight function $w : V(G) \to \mathbb{R}^+$ be any weight function, we define $w^*(v^i)$ for $i = 1, \dots, k(v)$ arbitrarily, so that $w^*(v^1) + \dots + w^*(v^{k(v)}) = w(v)$. We do this for every polygon in Λ and obtain a weight function $w^* : V(G^*) \to \mathbb{R}^+$. Since Λ^* is one-sided, there is a realization Γ^* in which the area of every polygon equals the weight of the corresponding vertex, i.e., the area of $\mathcal{P}(v^i)$ equals $w^*(v^i)$ for every $v \in V(G)$ and $i = 1, \dots, k(v)$. Since $\mathcal{P}(v)$ is the union $\mathcal{P}(v^1), \dots, \mathcal{P}(v^{k_v})$ its area equals $w(v)$. Thus the induced realization Γ of Λ respects the weights, which proves the lemma. □

We use Lemma 3.1.5 to prove our main result in this section. Let us underline the following statement in the upcoming lemma. An area-universal layout can be found in linear time. It is independent of any weight function. However, we do not know how to compute the actual cartogram for given vertex-weights in polynomial time. We discuss this issue, which already appears for one-sided layouts, in Remark 3.1.8 below.

Lemma 3.1.6. *Every maximally planar graph has a rectilinear area-universal layout with complexity at most* 8*. Moreover, the layout can be computed in linear time.*

Proof. De Fraysseix *et al.* [dFOdMR94] prove that every maximally planar graph $G = (V, E)$ has a $\perp$-representation, i.e., every vertex $v_i \in V$ is represented by an upside-down figure T, denoted by $\perp(i)$, consisting of a horizontal segment h_i and a vertical segment b_i with the lower end of b_i lying on h_i. A $\perp$ has three ends, called the *upper end*, *left end*, and *right end*. Every edge (v_i, v_j) in G is represented by a contact between an end of $\perp(i)$ and an interior, i.e., non-end point, of $\perp(j)$, or vice versa. Moreover, if the intersection $\perp(i) \cap \perp(j)$ is non-empty, then it consists of only an end. Figure 3.8 **b)** shows a $\perp$-representation for the maximally planar graph in Figure 3.8 **a)**. For readability, the $\perp$'s do not quite touch in the figure.

Let Γ be a $\perp$-representation of G. Such a representation can be computed in linear time. We "fatten" each $\perp(i)$ so that each vertex is represented by a T-shaped polygon. We replace each horizontal segment h_i by an axis-aligned rectangle H_i which has the same width as h_i, and whose top (bottom) side is $\varepsilon/2$ above (below) h_i, for some $\varepsilon > 0$. Similarly, we replace each vertical segment v_i by an axis-aligned rectangle B_i which has the same height as b_i and whose left (right) side is $\varepsilon/2$ to the left (right) of b_i. Note that this process creates intersections of H_i with B_i and some B_j, B_k at the left and right end of $\perp(i)$, respectively. It may as well create an intersection of B_i with some H_l at the top end of $\perp(i)$. We remove these intersections by replacing H_i by $H_i - B_j - B_k$ and replacing B_i by $B_i - H_i - H_l$. The resulting layout would then be a contact representation Γ' of G where each vertex v_i of G is represented by the T-shaped polygon $H_i \cup B_i$. Figure 3.8 **c)** illustrates such a representation of the maximally planar graph in Figure 3.8 **a)**.

In a next step, we remove all of the unused area, to obtain a hole-free representation Γ'' of G with 8-sided polygons. We do this by assigning each (rectangular) hole to one of the polygons adjacent to it. We start by placing an axis-aligned rectangle of minimum size that encloses Γ'. This creates five new bounded holes out of the unbounded region outside Γ': (i) to the left of B_1 and above H_1, (ii) to the right of B_1 and between H_1 and H_2, (iii) to the right of B_2 and above H_2, (iv) to the left of B_n and above H_n, (v) to the right of B_n and above H_n. Call these L_1, R_1, R_2, L_n and R_n, respectively. Now we associate each T-shaped polygon, representing v_i

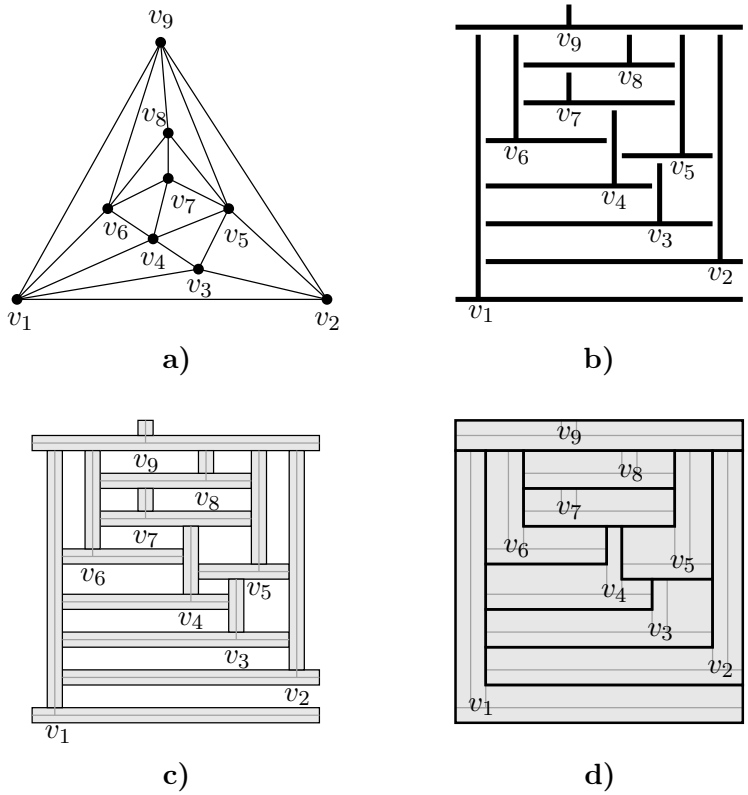

Figure 3.8: **a)** A maximally planar graph G. **b)** The $\perp$-representation Γ of G. **c)** The contact representation Γ' of G after the fattening. **d)** The hole-free representation Γ'' of G.

in Γ', with two holes: one above H_i and to the left of B_i, which we call the *left hole* of v_i; and the other to the top of H_i and to the right of B_i, which we call the *right hole* of v_i. Note that the left and the right hole of each v_i is bounded by H_i, B_i and some H_j and B_k, where (v_i, v_j) and (v_i, v_k) are edges in G. For $1 \leq i \leq n$, denote by L_i and R_i the left hole and the right hole of v_i.

We now can combine each T-shape region with its two associated holes to obtain an 8-sided rectilinear polygon. Specifically, for each vertex v_i, define $\mathcal{P}(v_i) = H_i \cup B_i \cup L_i \cup R_i$. It is easy to see that $\mathcal{P}(v_i)$ is an 8-sided rectilinear polygon since the left side of L_i is has the same x-coordinate as the left side of H_i and the right side of R_i has the same x-coordinate as the right side of H_i. Thus we have a hole-free representation Γ'', of G where each vertex v_i is represented by $\mathcal{P}(v_i)$. Figure 3.8 **d)** illustrates such a representation for the graph in Figure 3.8 **a)**. Note that Γ'' can be obtained from the $\perp$-representation in linear time.

Claim 1. Consider the set $\{H_i, B_i, L_i, R_i \mid i = 1, \ldots, n\}$ of $4n$ rectangles as a $\square$-

representation Γ^*; See Figure 3.9 for an example. Then the layout of Γ^* is one-sided.

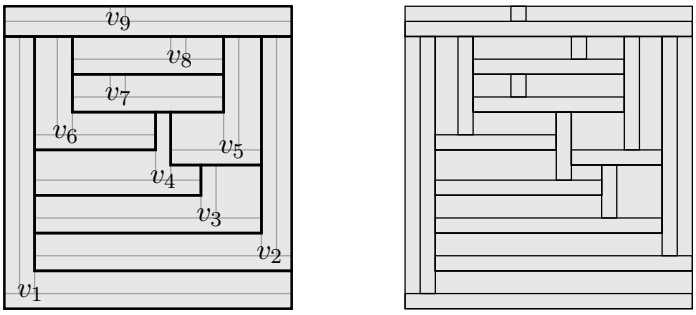

Figure 3.9: The representation Γ'' and the corresponding rectangle-representation Γ^*.

Proof of Claim 1. First note that the $\square$-representation Γ^*, which contains the rectangles L_i and R_i, $i = 1, \ldots, n$, has the same set of segments as the representation Γ' of G. Indeed, every horizontal and vertical segment arises from some rectangle H_i and B_i, respectively. (To be precise: This is not true for three of the four segments constituting the bounding rectangle. However those segments are always one-sided.) By the way rectangles H_i and B_i were defined, the segment for the top and bottom side of every H_i ends at the left and right side of this rectangle. Similarly, the left and right side of every B_i ends at the top and bottom side of this rectangle. Therefore, every segment (not on the outer boundary) is a side of one of the rectangles in Γ^*, i.e., the layout is one-sided. $\triangle$

Since the layout of Γ^* is a refinement of the layout Λ of Γ'', we conclude with Claim 1. and Lemma 3.1.5 that Λ is area-universal, which proves the lemma. $\square$

Theorem 3.1.7. *Let G be any maximally planar graph equipped with a weight function $w : V(G) \to \mathbb{R}^+$. Then G admits a cartogram Γ with rectilinear 8-sided polygons inscribed in any $H \times W$-rectangle, with area $\sum_{v \in V(G)} w(v)$. Moreover, the minimum feature size of Γ is at least $\frac{w_{min}}{2max(W,H)}$, which is worst-case optimal, and the number of supporting lines is at most $2|V(G)|$.*

Proof. The existence of some cartogram Γ with 8-sided polygons follows directly from Lemma 3.1.6. But note that, we have the freedom to choose how to divide the area assigned to any vertex v_i among the four rectangles associated with it. This flexibility makes it possible to achieve a large minimum feature size and a small set of supporting lines.

Specifically, let G be a maximally planar graph, $w : V(G) \to R^+$ some weight function, and Λ be the area-universal layout for G from Lemma 3.1.6. Assume we are given the height H and width W of the bounding rectangle. It clearly has to satisfy

$H \times W = \sum_{v \in V(G)} w(v)$. We define $w_{min} = min_{v \in V(G)} w(v)$. For each vertex v_i of G, we can assign zero areas to the rectangles L_i and R_i and split its weight $w(v_i)$ into two equal parts to H_i and B_i.

Now let Γ be a cartogram of G for the split weight function w^*. Since the height of each of the rectangles H_i and B_i is at most $max(W, H)$ and their weights are at least $w_{min}/2$, the width is at least $\frac{w_{min}}{2max(W,H)}$, and the same holds for its height. Thus the minimum feature size of Γ' is at least $\frac{w_{min}}{2max(W,H)}$ and it is worst-case optimal, as the polygon with the smallest weight might need to reach from the leftmost to the rightmost (or topmost to bottommost) polygon in the representation. We remark that if the bounding rectangle is not given, we may choose a square, i.e., $W = H = \sqrt{A}$, where $A = \sum_{v \in V(G)} w(v)$. Then the minimum feature size is $\frac{w_{min}}{2\sqrt{A}}$.

Furthermore it is not difficult to show that our assignment of zero areas to the rectangles L_i and R_i yields a cartogram Γ with at most $2|V(G)|$ supporting lines, instead of the bound of $3|V(G)|$ from the number of segments in the layout Λ. □

In Figure 3.10 we provide an example for a cartogram computed with Theorem 3.1.7. However, we did not set the weight of rectangles L_i and R_i to zero and hence obtain T-shape polygon with eight supporting lines. Figure 3.10 **a)** shows the adjacency graph of some European countries and Figure 3.10 **b)** an area-universal layout is shown. Since the graph is only internally triangulated three artificial outer vertices were first added and the corresponding polygons afterwards removed. As weights we choose the millions of tonnes CO_2 the corresponding country emitted in 2009. The final cartogram is shown in Figure 3.10 **c)**.

Remark 3.1.8. The proof of Lemma 3.1.6 provides a linear-time algorithm for constructing an area-universal layout of complexity at most 8. One step is a linear-time construction of a $\perp$-representation, which is a fairly simple algorithm [dFOdMR94]. However, Theorem 3.1.7 relies on Theorem 3.1.3, whose proof does not give a polynomial (not even exact) algorithm for computing the actual cartogram from the layout. The computation of a cartogram for a given one-sided layout and given weights can be accomplished using a result of Wimer *et al.* [WKC88], which in turn requires numerical iteration. It remains open to find a polynomial-time algorithm to compute cartograms with complexity at most 8. On the other hand, there is a linear-time algorithm [ABF+11a], which produces cartograms with complexity bounded by 10.

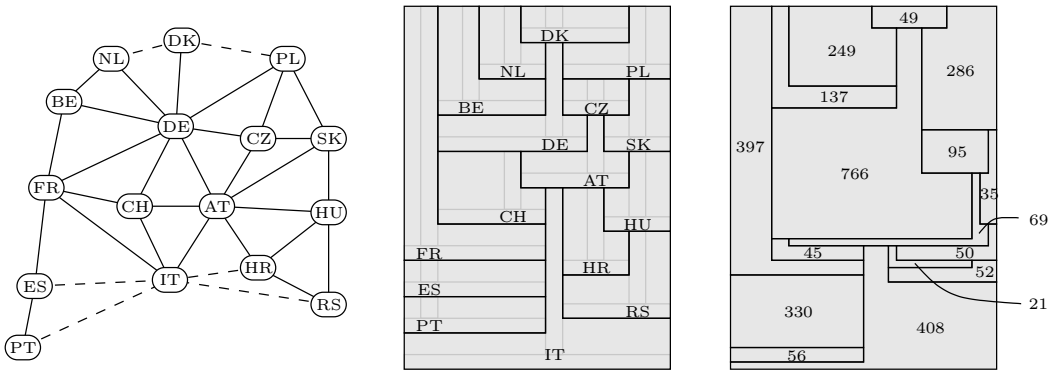

Figure 3.10: **a)** The adjacency graph G of some countries in central Europe. (For convenience, Italy and Denmark are bordering some more countries here.) **b)** An area-universal layout Λ for G with complexity at most 8. **c)** A cartogram of G w.r.t. the inscribed weights which realizes Λ. The numbers are the millions of tonnes of CO_2 that were emitted in the corresponding country in 2009.

3.2 Cartograms for Hamiltonian Maximally Planar Graphs

This section deals with cartograms for Hamiltonian maximally planar graphs. Since the Hamiltonicity problem for maximally planar graphs is NP-complete [Chv85], we assume here, that we are given such a graph together with a Hamiltonian cycle. If for instance, G is 4-connected, then a Hamiltonian cycle can be found in linear time [AKS84]. Having a Hamiltonian cycle in G at hand, we provide an alternative representation for G with complexity at most 8, which involves only polygons with the U-shape in the bottom-right corner in Figure 2.4 and degenerated versions of it. For simplicity, we call such a representation a $\sqcap$*-representation*. I was told the following lemma by M. J. Alam, S. Felsner, and S. G. Kobourov.

Lemma 3.2.1. *Every Hamiltonian maximally planar graph admits a* $\sqcap$*-representation.*

Proof. Let $v_1, \ldots, v_n$ be a Hamiltonian cycle of a maximally planar graph G. We choose any embedding of G with (v_1, v_n) on the outer face. The Hamiltonian cycle splits the plane graph G into two outer-planar graphs which we call the *left graph* G_l and *right graph* G_r. Edges on the Hamiltonian cycle belong to both graphs. The naming is with respect to a planar drawing of G in which the vertices $v_1, \ldots, v_n$ are placed in increasing order along a vertical line, and the edges are drawn with y-monotone curves with leftmost edge (v_1, v_n). See the left of Figure 3.11.

For each vertex v_i in the graph we define several indices as follows. First $b_l(i)$ is the index of the bottommost neighbor of v_i in G_l, i.e., $b_l(i)$ is the smallest index j such that (v_j, v_i) is an edge in G_l. Next $b_l^{-1}(i)$ is the (possibly empty) set of indices j such that $i = b_l(j)$. The other two indices, $b_r(i)$ and $b_r^{-1}(i)$, are defined analogously, using G_r in place of G_l. Let v_k be the third outer vertex besides v_1 and v_n. We define (exceptionally) the following indices of v_1, v_k and v_n: $b_l(v_1) = b_r(v_1) = 0$ and $t_l(v_1) = t_r(v_1) = t_r(v_k) = t_l(v_n) = t_r(v_n) = n + 1$.

The corresponding polygon of each vertex v_i in the graph is subdivided into three rectangles, which we call the *left leg*, *body*, and *right leg* of v_i. We present here a $\sqcap\!\sqcup$-representation with particular coordinates using the above indices. The proof that this representation is in fact is representing the graph G follows from the upcoming theorem and is omitted here. An example is depicted in Figure 3.11. We close this proof by listing the coordinates of the left leg, body, and right leg of each vertex v_i.

- The left leg of v_i is the (possibly zero-area) rectangle $[-t_l(i), -i] \times [b_l(i), i-1]$.
- The base of v_i is the rectangle $[-t_l(i), t_r(i)] \times [i-1, i]$.
- The right leg of v_i is the (possibly zero-area) rectangle $[i, t_r(i)] \times [b_r(i), i-1]$.

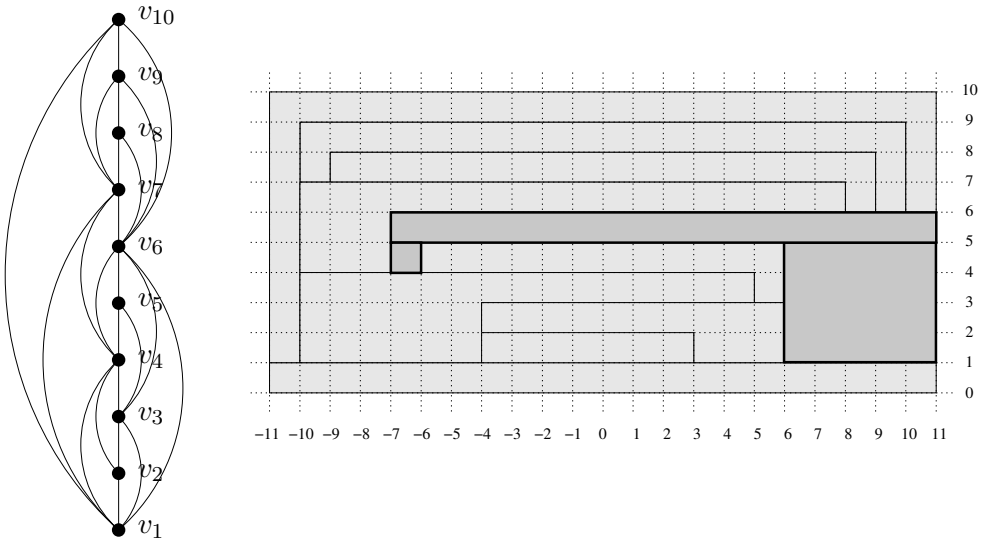

Figure 3.11: A Hamiltonian graph embedded with (v_1, v_n) on the outer face, $v_1, \ldots, v_n$ placed along a vertical line and a contact representation with rectilinear 8-sided polygons. The three rectangles for the vertex v_6 are highlighted.

□

Note that in the proof of Lemma 3.2.1 we require a Hamiltonian cycle that contains an outer edge. We ensure this by choosing the embedding of the graph. Hence, the

above result is the first within this thesis in which we can *not* prescribe the embedding that is inherited by the representation we construct.

From the next theorem follows that the ⊓⊔-layout Λ corresponding to the ⊓⊔-representation from Lemma 3.2.1 is area-universal. Indeed, we provide an algorithm to compute a cartogram Γ for G w.r.t. a given weight-function $w : V(G) \to \mathbb{R}^+$, so that Γ realizes Λ. The strength of this algorithm is its linear runtime. One can prove that the layout Λ as described in Lemma 3.2.1 is area-universal by finding a one-sided □-layout Λ^* which refines Λ and applying Lemma 3.1.5. However, as discussed in Remark 3.1.8 we do not know how to get an efficient or exact algorithm computing the actual cartogram out of this. Note that the □-layout consisting of left leg, body, and right leg for every vertex as shown in Figure 3.11 refines Λ, but is not one-sided. Indeed, every one-sided □-layout refining a ⊓⊔-layout from Lemma 3.2.1 in general requires four rectangles for some polygons.

Theorem 3.2.2. *Let G be a Hamiltonian maximally planar graph equipped with a weight function $w : V(G) \to \mathbb{R}^+$. Then a cartogram of G with rectilinear 8-sided polygons can be computed in linear time.*

Proof. Let $v_1, \ldots, v_n$ be a Hamiltonian cycle and G be embedded as in Lemma 3.2.1, i.e., with the vertices placed in increasing order along a vertical line, and the edges are drawn with y-monotone curves with leftmost edge (v_1, v_n). Suppose $\mathcal{R}$ is a rectangle of width W and height H where $W \times H = \sum_{v \in V} w(v)$. We compute a cartogram of G using the ⊓⊔-layout of the representation in Lemma 3.2.1. It is easy to see that the four indices $b_l(i)$, $b_l^{-1}(i)$, $b_r(i)$,$b_r^{-1}(i)$ can be computed for all the vertices v_i of G in linear time. We also define two special types of edges associated with every vertex v_i, $1 < i < n$: a *left i-chord* in G is an edge $(v_k, v_{b_l(k)})$ where $k > i > b_l(k)$, and a *right i-chord* in G is analogously defined. Additionally, we define $\lambda_i = w(v_i)/(2H + W)$, which will be the width of the legs of v_i.

Our algorithm is a line-sweep, where a horizontal line L sweeps the embedding of G from bottom to top, while the algorithm iteratively computes the cartogram of G inside $\mathcal{R}$. This algorithm does not rely on numerical iteration, as it computes the cartogram directly. We start the construction with the polygon for v_1, which is simply a rectangle with area $w(v_1)$ at the bottom of $\mathcal{R}$; call this rectangle $\mathcal{R}_1$. After the i-th step of our algorithm, the horizontal line L in the embedding of G is at vertex v_i, and the algorithm has constructed the polygons for all the vertices $v_1, \ldots, v_i$, and a part of the left leg and right leg of vertices v_k with $b_l(v_k) < i$ and $b_r(v_k) < i$, respectively. These polygons and partial legs completely fill up a rectangle $\mathcal{R}_i$ at the bottom of $\mathcal{R}$. The top side of $\mathcal{R}_i$ would contain the followings from left to right:

1) The top side of a partial left leg of width λ_j, for all the left i-chords $(v_j, v_{b_l(j)})$ ending on the top side of $\mathcal{P}(v_{b_l(j)})$, in the order of the intersection of these

edges with L;

2) The top side of $\mathcal{P}(v_i)$;
3) The top side of a partial right leg of width λ_j, for all the right i-chords $(v_j, v_{b_r(j)})$ ending on the top side of $\mathcal{P}(v_{b_r(j)})$, in the order of the intersection of these edges with L.

Suppose we are at the i-th step of our algorithm. We compute $\mathcal{R}_i$ from $\mathcal{R}_{i-1}$ as follows. If $b_l(i) = k < i-1$, then due to planarity, $b_l^{-1}(i-1)$ is empty and the rightmost left $(i-1)$-chord is $(v_i, b_l(v_i))$. Hence the rightmost partial left leg corresponds to v_i. We also have that every left i-chord is a left $(i-1)$-chord. Hence, we do not introduce any new partial left leg. Such a step is illustrated for the left side of v_i in Figure 3.12. Otherwise, if $b_l(i) = i-1$, then the polygon $\mathcal{P}(v_i)$ does not have a left leg. However, we introduce a partial left leg of width λ_j for every index $j \in b_l^{-1}(i-1)$. Note that the new partial legs do not cover the entire top side of $\mathcal{P}(v_{i-1})$, since $\sum \lambda_j = \frac{\sum w(v_j)}{2H+W} < \frac{W \cdot H}{H} = W$. This case is illustrated for the right side of v_i in Figure 3.12. We compute the right leg of $\mathcal{P}(v_i)$ and possibly new partial right legs in a similar fashion. We introduce the body of v_i horizontally spanning from the partial left leg of the rightmost left i-chord to the partial right leg of the leftmost right i-chord. We then set the height of the body appropriately, so that the polygon $\mathcal{P}(v_i)$ (the union of the left, the right leg, and the body) has area $w(v_i)$. Note that the sum of the weights of the left and right leg of v_i is at most $\lambda_i \cdot 2H$ and the width of the body of v_i is at most W. Thus, by the choice of λ_i, the height of the body of v_i is at least λ_i. We complete the i-th step by extending every partial left and right leg vertically, so that it lines up with the top side of $\mathcal{P}(v_i)$.

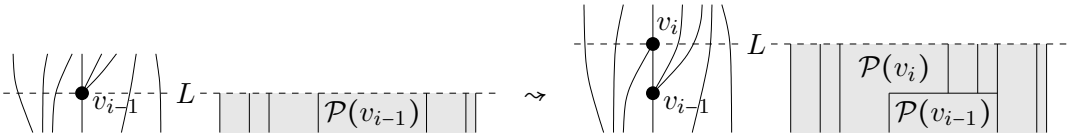

Figure 3.12: The i-th step of the algorithm with $b_l(i) = k < i-1$ and $b_r(i) = i-1$.

From this construction is it easy to see that, after the n-step of the algorithm, the polygon for each vertex v_i inside $\mathcal{R}$ has the correct area. We now show that the correct adjacencies are also preserved. Consider an edge $e = (v_i, v_j)$ of G, where $i > j$. If e is a Hamiltonian edge, i.e., $i = j+1$, then clearly the polygons for v_i and v_j are adjacent. Next consider the case that e is in G_l (the case that e is in G_r is analogous). If $j = b_l(i)$, then this adjacency is maintained by the left leg of v_i. Otherwise suppose $b_l(i) = k$; then (v_i, v_k) is the rightmost left j-chord and hence the left leg of v_j is adjacent to the left leg of v_i. Therefore, the algorithm has computed the desired cartogram of G. It is also not difficult to argue that this algorithm runs in linear time. □

If we only care about the complexity of a rectilinear cartogram, then Theorem 3.2.2 is strictly weaker than Theorem 3.1.7. It considers only a subclass of maximally planar graphs and still does provide the same bound on the sufficient complexity. On the other hand, we can show that complexity 8 is actually necessary even for cartogram for Hamiltonian maximally planar graphs. Note that without weights, every such graph has a rectilinear representation of complexity at most 6, c.f. Corollary 2.2.5.

Lemma 3.2.3. *Consider the Hamiltonian maximally planar graph $G = (V, E)$ in Figure 3.13. Define $w(a) = w(b) = D$ and $w(v) = \delta$ for $v \in V \setminus \{a, b\}$, where $D > \delta$. Then any rectilinear cartogram of G with weight function w requires at least one 8-sided polygon.*

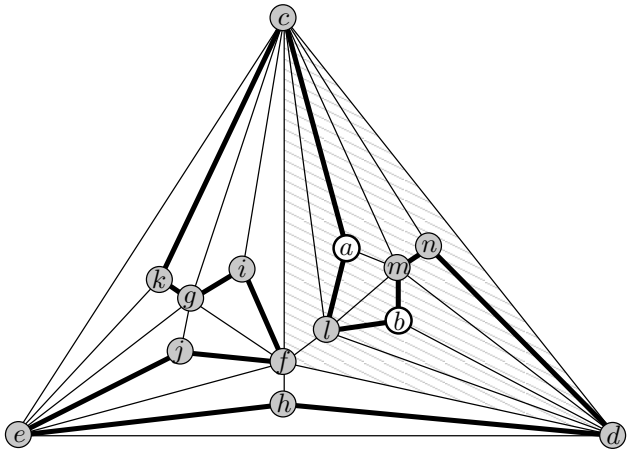

Figure 3.13: A Hamiltonian graph G with grey and white vertices weighted δ and D, respectively. If $D > \delta$ then there is no cartogram for G with only 4-sided and 6-sided rectilinear polygons.

Proof. Let Γ be a cartogram for G with respect to w with rectilinear polygons. Assume for a contradiction that all the polygons used in Γ have complexity at most 6. It is easy to see that if $\{u, v, w\}$ is some separating triangle in G, then $\mathcal{P}(u)$, $\mathcal{P}(v)$, and $\mathcal{P}(w)$ enclose a rectilinear region R_{uvw} in Γ, which has at least four sides. Consequently at least one corner of R_{uvw} is a concave corner of $\mathcal{P}(u)$, $\mathcal{P}(v)$ or $\mathcal{P}(w)$. We associate R_{uvw} with this polygon. Note that two regions can be associated to the same polygon only if they have the relation of containment. (Indeed, associating $\mathcal{R}_{uvw}$ with a polygon that has a concave corner there is a nesting assignment as defined in Definition 2.2.2, which by Theorem 2.2.4 is necessarily induced by Γ.)

First, consider the subgraph G' of G that is highlighted in Figure 3.13. The 4-vertex set $\{c, d, l, m\}$ is the disjoint union of three separating triangles $\{c, d, m\}$, $\{c, l, m\}$,

$\{d, l, m\}$ in G'. Since no two of these triangle are contained in each other, $\mathcal{P}(c)$ or $\mathcal{P}(d)$ (or both) is associated with one of regions R_{cdm}, R_{clm} and R_{dlm}.

Now consider the entire graph G and note that the 5-vertex set $\{c, d, e, f, g\}$ in G is the union of the five separating triangles $\{c, d, f\}$, $\{d, e, f\}$, $\{c, e, g\}$, $\{c, f, g\}$, and $\{e, f, g\}$ with disjoint interiors. It follows that each of the five separating triangles above is associated with the concave corner of the polygon for c, d, e, f, or g. Since $\mathcal{P}(c)$ or $\mathcal{P}(d)$ is associated within G', the same polygon is associated with R_{cdf} and the other of the two, as well as $\mathcal{P}(f)$, is associated with a region outside of G'. Hence R_{cdf} is a rectangle and by symmetry of the graph G', we may assume that $\mathcal{P}(c)$ is associated with it, i.e., constitutes two sides of R_{cdf}. There remain three possible assignments of concave corners of $\mathcal{P}(c)$, $\mathcal{P}(l)$ and $\mathcal{P}(m)$ to the regions R_{cdm}, R_{clm} and R_{dlm}:

1) $c \mapsto R_{clm}$, $l \mapsto R_{dlm}$, $m \mapsto R_{cdm}$ (first row of Figure 3.14)
2) $c \mapsto R_{cdm}$, $l \mapsto R_{dlm}$, $m \mapsto R_{clm}$ (second row of Figure 3.14)
3) $c \mapsto R_{cdm}$, $l \mapsto R_{clm}$, $m \mapsto R_{dlm}$ (symmetric to second row of Figure 3.14)

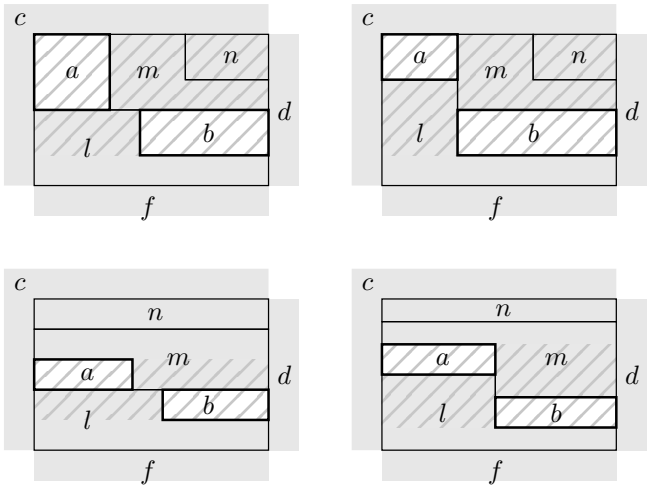

Figure 3.14: A detail of three hypothetical cartograms with rectilinear polygons of complexity at most 6.

Figure 3.14 shows the corresponding layout to each of the above assignments. (To be precise, each of $\mathcal{P}(a)$, $\mathcal{P}(b)$, $\mathcal{P}(n)$ may have L-shape if Γ is not hole-free, but this does not affect the upcoming argument.) Note that the shaded part in each layout in Figure 3.14 is equivalent to the shaded part in Figure 3.7 and hence by Lemma 3.1.4 none of the layouts occurs in the cartogram Γ. This is a contradiction. Thus at least one of the polygons in Γ contains 8 sides. □

Putting Theorem 3.2.2 and Lemma 3.2.3 together, we have that 8-sided polygons are always sufficient and sometimes necessary for a hole-free rectilinear cartogram of a Hamiltonian maximally planar graph.

3.2.1 One-Sided Hamiltonian Cycles

Again consider a Hamiltonian cycle $(v_1, \dots, v_n)$ of a maximally planar graph G, which is embedded with (v_1, v_n) on the outer face. Recall from the proof of Lemma 3.2.1 that the Hamiltonian cycle splits G into the left graph G_l and the right graph G_r. Moreover, for vertex v_i, $i = 1, \dots, n$, the indices $b_l(i)$ and $b_r(i)$ denote the smallest j such that (v_i, v_j) is an edge in G_l and G_r, respectively. Also recall that in the $\sqcap\!\sqcup$-representation from Lemma 3.2.1 a vertex v_i is assigned an 8-sided polygon only if neither its bottommost neighbor in G_l, nor its bottommost neighbor in G_r, is v_{i-1}. Indeed, the left leg of v_i, defined as $[-t_l(i), -i] \times [b_l(i), i-1]$, has zero area if $b_l(i) = i-1$. Similarly, the right leg of v_i, defined as $[i, t_r(i)] \times [b_r(i), i-1]$, has zero area if $b_r(i) = i-1$.

Here we define a special property of the embedded Hamiltonian cycle, called *one-sidedness*. We call a Hamiltonian cycle $(v_1, \dots, v_n)$ in an embedded maximally planar graph with (v_1, v_n) on the outer face *one-sided* if for every $i = 2, \dots, n$ we have $b_l(i) = i-1$ or $b_r(i) = i-1$. Figuratively speaking, a Hamiltonian cycle is one-sided, if for every vertex in the y-monotone drawing it has no edges going downwards on the left side or on the right side. The Hamiltonian cycle in the maximally planar graph in the left of Figure 3.15 is one-sided w.r.t. the outer edge (v, w), i.e., when choosing $v_1 = v$, $v_n = w$, and hence traversing the cycle in clockwise direction. In particular the figure shows a y-monotone drawing of the one-sided Hamiltonian cycle $(v, u, x, \dots, w)$. On the other hand, the right of Figure 3.15, which is a rotation of the left-hand picture, shows a y-monotone drawing of the cycle w.r.t. the outer edge (v, u), i.e., under the choice $v_1 = v$, $v_n = u$, and a thus a counterclockwise traversal of the cycle. Since the vertex x has edges going downwards on the left *and* on the right side of the cycle, we have that $(v, w, \dots, x, u)$ is not a one-sided Hamiltonian cycle.

Since a Hamiltonian cycle $(v_1, \dots, v_n)$ is one-sided if and only if $b_l(i) = i-1$ or $b_r(i) = i-1$ for every $i = 1, \dots, n$ the following lemma is a direct consequence of Lemma 3.2.1 and Theorem 3.2.2.

Lemma 3.2.4. *Let $G = (V, E)$ be a maximally planar graph with a one-sided Hamiltonian cycle and let $w : V \to \mathbb{R}^+$ be a weight function. Then a rectilinear cartogram with 6-sided polygons can be computed in linear time.*

Note that not every Hamiltonian maximally planar graph admits a one-sided Hamiltonian cycle; e.g., the graph in Figure 3.13 does not even admit a rectilinear cartogram with complexity 6.

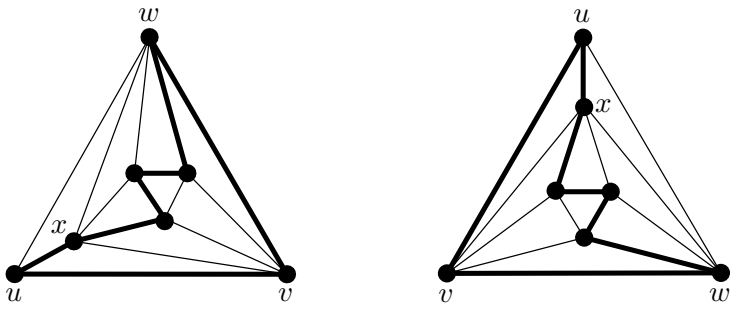

Figure 3.15: A maximally planar graph with a Hamiltonian cycle that is one-sided w.r.t. $(v_1 = v, v_n = w)$ (left) but not w.r.t. $(v_1 = v, v_n = u)$ (right).

It turns out that one-sided Hamiltonian cycles can be defined in many equivalent ways, some of which are intimately related to Schnyder woods. Given a Hamiltonian cycle $(v_1, \dots, v_n)$, we fix a plane embedding of G with outer triangle $\{v_1, v_k, v_n\}$. Then for every $i = 2, \dots, n$ we consider the plane subgraph $G_i \subseteq G$ induced by $\{v_1, v_2, \dots, v_i\}$ with the embedding inherited from G. Moreover, define $w_i := v_{n-i+1}$ and denote by $\tilde{G}_i$ the subgraph of G induced by $\{w_1, w_2, \dots, w_i\}$, i.e., $\tilde{G}_i = G \setminus G_{n-i}$ and $\tilde{G}_{n-i} = G \setminus G_i$.

Lemma 3.2.5. *Let $(v_1, \dots, v_n)$ be a Hamiltonian cycle in an embedded maximally planar graph G with (v_1, v_n) on the outer triangle. Then the following are equivalent:*

(a) $b_l(i) = i - 1$ or $b_r(i) = i - 1$ for $i = 2, \dots, n$, i.e., the cycle is one-sided.
(b) The edge (v_{i-1}, v_i) is an outer edge of G_i for $i = 2, \dots, n$.
(c) v_{n-1} is an outer vertex and the vertex v_i has at least two neighbors with a larger index for $i = 1, \dots, n-2$.
(d) $(w_1, \dots, w_n)$ is a canonical order for G.

Proof. (a) $\iff$ (b): For $i = 2, \dots, n$ we argue that (a) holds for i if and only if (b) holds for i. Indeed, (v_{i-1}, v_i) is an inner edge in G_i if and only if there are boundary edges (v_i, v_j) and (v_i, v_k) with $j, k < i - 1$ in G_l and G_r, respectively. But this is equivalent to $b_l(i) = j < i - 1$ and $b_r(i) = k < i - 1$.

(b) $\iff$ (c): Since v_n is an outer vertex and $G_n = G$, (b) holds for $i = n$ if and only if v_{n-1} is an outer vertex. For $i = 2, \dots, n-1$ we argue that (b) holds for i if and only if (c) holds for $i - 1$. Assume (v_{i-1}, v_i) is an inner edge in G_i. Let v_i^l, respectively v_i^r, denote the third vertex in the inner facial triangle containing the edge (v_{i-1}, v_i) in G_l, respectively G_r. (Both triangles exist, since (v_{i-1}, v_i) is an inner edge in G.) Now v_i^l and v_i^r have a smaller index than v_{i-1}, which holds if and only if the index of *every* neighbor of v_{i-1}, different from v_i, is smaller than $i - 1$.

(c) $\Longrightarrow$ (d): By (c) $\{w_1, w_2, w_n\} = \{v_n, v_{n-1}, v_1\}$ is the outer triangle of G, and moreover, $\tilde{G}_3$, which is induced by v_n, v_{n-1}, v_{n-2}, is a triangle. Hence the outer boundary of $\tilde{G}_3$ is a simple cycle C_3 containing the edge (w_1, w_2). In other words, the first condition of a canonical order is met for $i = 4$ (c.f. Definition 1.1.3). Assuming (c) and the first condition for fixed $i \in \{4, \dots, n-1\}$, we show that the second and first condition holds for i and $i+1$, respectively. In the end, the second condition holds for $i = n$ since w_n is an outer vertex.

First note that w_i is in the outer face of $\tilde{G}_{i-1}$ since w_n lies in the outer face, the path $w_i, \dots, w_n$ is disjoint from vertices in $\tilde{G}_{i-1}$ and the embedding is planar. By (c) w_i has at least two neighbors in $\tilde{G}_{i-1}$. If the neighbors would not form a subpath of the path $C_{i-1} \setminus (w_1, w_2)$, there would be a non-triangular inner face in $\tilde{G}_i$, which contains a vertex w_j with $j > i$ in its interior. But then the path $w_j, \dots, w_n$, which is disjoint from $\tilde{G}_i$, would start at an inner vertex of $\tilde{G}_i$ and end in the outer face of $\tilde{G}_i$. This again contradicts planarity. Thus the second condition of a canonical order is satisfied for i. Moreover, $\tilde{G}_i$ has a simple outer cycle C_i containing the edge (w_1, w_2). In other words, the first condition holds for $i+1$.

(d) $\Longrightarrow$ (c): Since it is a canonical order, (w_1, w_2) is an outer edge. In particular, $w_2 = v_{n-1}$ is an outer vertex. Clearly, v_1 has at least two neighbors and every neighbor has a larger index, i.e., (c) holds for $i = 1$. Moreover, by the second condition of a canonical order every vertex $v_i = w_{n-i+1}$, for $i = 2, \dots, n-2$, has at least two neighbors in $\tilde{G}_{n-i} = G \setminus G_i$, which is the subgraph induced by $v_{i+1}, \dots, v_n$. □

Lemma 3.2.6. *If G is a maximally planar graph with outer triangle $\{w_1, w_2, w_n\}$, and $(w_1, w_2, \dots, w_n)$ is a vertex ordering, then the following are equivalent:*

(a) $(w_1, w_2, \dots, w_n)$ is a one-sided Hamiltonian cycle.
(b) $(w_1, w_2, \dots, w_n)$ is a canonical order, which is a Hamiltonian path.
(c) The graph $G_i = G[w_1, \dots, w_i]$ contains no valley vertex for $i = 3, \dots, n$.
(d) In every Schnyder wood (T_1, T_2, T_3) defined by $(w_1, \dots, w_n)$ every inner vertex is a leaf in T_1 or T_2.

Proof. (a) $\Longleftrightarrow$ (b) follows directly from Lemma 3.2.5 (a) and (d).

(b) $\Longleftrightarrow$ (c): In *every* canonical order w_i is a hill vertex of G_i for $i = 3, \dots, n-1$, and hence G_3 contains no valley vertex. Moreover, by Lemma 1.1.7 there is always one more hill vertex than there are valley vertices in G_i, for $i = 3, \dots, n-1$. Now again by Lemma 1.1.7 w_i is connected to w_{i-1} if and only if w_{i-1} is no hill vertex in G_i, for $i = 4, \dots, n$.

(c) $\Longleftrightarrow$ (d): By Lemma 1.1.7 every valley vertex in G_i is neither a leaf in T_1 nor a leaf in T_2. On the other hand, every vertex that is neither in L_1 nor in L_2 is a valley vertex in some G_i. □

We remark that a canonical order defined by a Schnyder wood is a Hamiltonian path if and only if it is the *unique* canonical order defined by this Schnyder wood.

Figure 3.16 shows an example of a maximally planar graph with a one-sided Hamiltonian cycle, the corresponding canonical order, and Schnyder wood. Note that since the canonical order $(w_1, \dots, w_n)$ is the *reverse* of the vertex ordering $(v_1, \dots, v_n)$ that we used for the Hamiltonian cycle, every vertex has no incident edges of the left or the right going *up*, rather than down.

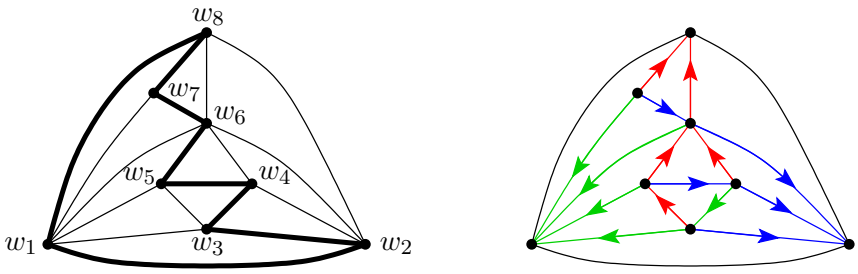

Figure 3.16: A maximally planar graph with a one-sided Hamiltonian cycle (left) and the corresponding Schnyder wood (right).

Once we have a one-sided Hamiltonian cycle, we can build a 6-sided cartogram via Lemma 3.2.4 in linear time. Alternatively, we could obtain from it a Schnyder wood, rooted such that every vertex is a leaf in T_1 or T_2. Taking this as the underlying Schnyder wood in Theorem 3.1.7 results in the very same cartogram. However, we prefer to use Lemma 3.2.4 since it additionally states that a cartogram can be computed in linear time.

Remark 3.2.7. Let G be a maximally planar graph and $(w_1, \dots, w_n)$ be a canonical order of G which is a Hamiltonian path, or equivalently $(w_n, \dots, w_1)$ be a one-sided Hamiltonian cycle. We get yet another hole-free rectilinear for G layout with complexity 6 by observing that

- if (T_1, T_2, T_3) is a Schnyder wood such that every vertex is in L_1 or L_2, then we have $|L_1 \cap L_2 \cap in_3(v)| \leq 1$ for every vertex v.

The above is true, since by Lemma 3.2.6 the graph G_i contains no valley vertex and thus only one hill vertex, for $i = 3, \dots, n-1$. Because $|L_1 \cap L_2 \cap in_3(v_{i+1})|$ is the number of hill vertices of G_i in $in_3(v_{i+1})$, this quantity is at most one. Thus, by Theorem 2.1.7 there exists a non-rotated $\sqsupset$-representation of G, whose layout we call Γ_2. Actually, the layout Λ_2 differs significantly from the one in Lemma 3.2.4, which we call Λ_1. For instance, Λ_2 is non-rotated, while generally Λ_2 is not. But even the overall complexities (see Subsection 2.1.1) in Λ_1 and Λ_2 differ. Let $(w_1, \dots, w_n)$

be a canonical order which is a Hamiltonian path, or equivalently $(w_n, \ldots, w_1)$ be a one-sided Hamiltonian cycle. It is not difficult to see that $\mathcal{P}(w_i)$ has an L-shape in Λ_2 if and only if w_{i-1} is a leaf in both, T_1 *and* T_2, which in turn holds if and only if $\mathcal{P}(w_{i-1})$ is a rectangle in Λ_1. Thus it follows that the number of L-shaped polygons in Λ_2 equals the number of rectangles in Λ_1 minus 1. (We subtract one, since $\mathcal{P}(w_n)$ and $\mathcal{P}(w_1)$ is always a rectangle in Λ_1 and Λ_2, respectively.) Figure 3.17 **b)** and Figure 3.17 **c)** show the layout Λ_1 and Λ_2 for the one-sided Hamiltonian cycle from Figure 3.17 **a)**, respectively.

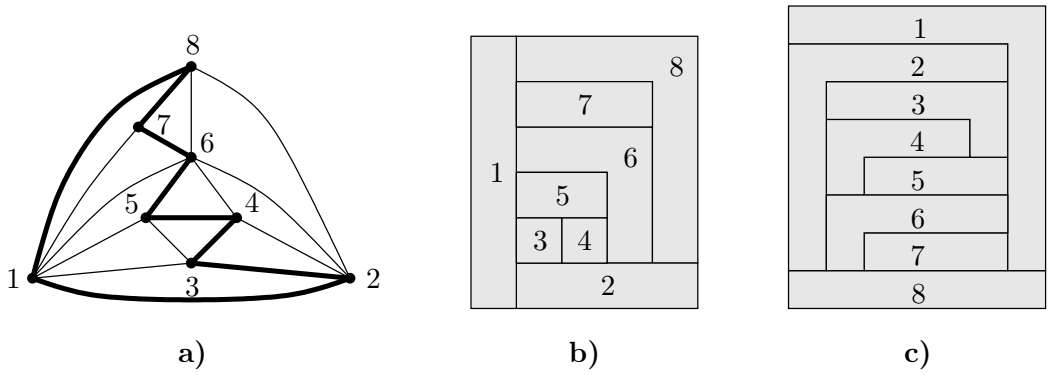

Figure 3.17: **a)** A maximally planar graph with a one-sided Hamiltonian cycle. **b)** The layout Λ_1 as given by Lemma 3.2.4. **c)** The layout Λ_2 as given by Theorem 2.1.7.

Unfortunately, in general the layout Λ_2 is not (rectilinear) area-universal; it may even be degenerated to a $\square$-layout that is not one-sided. Hence we do not know how to derive a cartogram from it.

Let us close this section with the following corollary of Lemma 3.2.4.

Corollary 3.2.8. *A cartogram of a maximally outer-planar graph using* 6*-gons can be computed in linear time.*

Proof. Choose an outer-planar embedding of G. Let $(v_1, \ldots, v_n)$ be the Hamiltonian cycle in G consisting of all outer edges. Add a vertex v_{n+1} into the outer face connect it by an edge to every v_i, so that (v_1, v_n) remains an outer edge. Then $(v_1, \ldots, v_n, v_{n+1})$ is a Hamiltonian cycle of the resulting embedded maximally planar graph G'. Indeed, this cycle is one-sided since for $i = 2, \ldots, n$ the vertex v_i has only three neighbors in G_l, namely v_{i-1}, v_{i+1} and v_{n+1}, which implies $b_l(i) = i - 1$ here. The new vertex v_{n+1} has only one neighbor in G_r, namely v_n, i.e., $b_r(n+1) = n$. With Theorem 3.2.2 we get a cartogram of G' with 6-sided rectilinear polygons. (We give v_{n+1} any positive weight.)

Finally note that removing the polygon for v_{n+1} from the cartogram leaves a cartogram of the original graph G, still of complexity 6. Moreover, this representation fits inside a rectangular area since the removed polygon is 6-sided and occupies two consecutive sides of the outer boundary. □

A linear-time algorithm for constructing a cartogram of a maximally outer-planar graph with 6-sided rectilinear polygons was known before [ABF$^+$11a]. However, the construction based on Theorem 3.2.2 is simpler. Moreover, a lower bound of 6 on the complexity of such a cartogram with the additional constraint that the outer boundary is a rectangle is known [Rin87]. (If we do not insist on the additional property, then maximally outer-planar graphs admit cartograms that use only rectangles [ABF$^+$11a].) Putting things together, we have that 6-sided polygons are always sufficient and sometimes necessary for a hole-free rectilinear cartogram of a maximally outer-planar graph.

3.3 Lower Bounds on the Complexity

In this section we consider general (not necessarily rectilinear or hole-free) cartograms of maximally planar graphs. We provide new lower bounds on the polygonal complexity of such cartograms for the class of Hamiltonian maximally planar graphs and planar 3-tree. The proofs consist of Lemma 2.3.1 for the previous chapter, which counts the number of *convex* corners in a representation, and the upcoming Lemma 3.3.1, which gives a way of forcing additional *concave* corners. The following is a very recent result [ABF$^+$11b].

Lemma 3.3.1 ([ABF$^+$11b])**.** *Consider the maximally planar graph $G^* = (V, E)$ in Figure 3.18* ***a)****. Define $w(a) = w(b) = w(c) = D$ and $w(v) = \delta$ for $v \in V \setminus \{a, b, c\}$, where $D \gg \delta$. Then every cartogram of G^* with weight function w requires at least one polygon with a concave corner.*

We are going to use the graph G^* from Lemma 3.3.1 as a building block for our examples below. In particular, we first define a maximally planar graph G, identify a subset of its inner faces and put the graph G^* into each such face, i.e., the three outer vertices of G^* are identified with the three vertices in G of the inner face. The path in Figure 3.18 **a)** will be used to show, that the graph G may remain Hamiltonian after the insertion of G^*.

Lemma 3.3.2. *Each of the following holds.*

(i) There is a Hamiltonian maximally planar graph equipped with a weight function, such that every cartogram requires polygonal complexity at least 6*.*

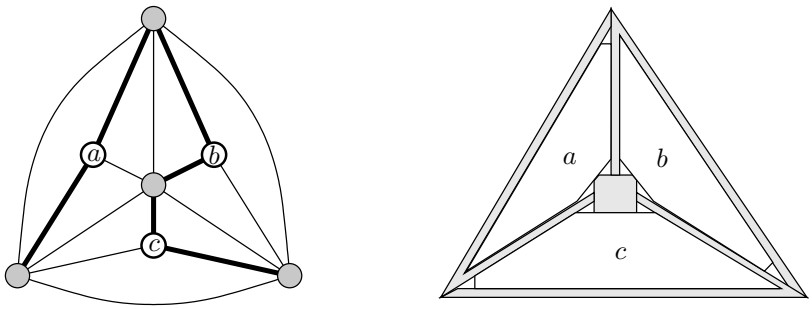

a) **b)**

Figure 3.18: **a)** A graph G^* with grey and white vertices weighted δ and D, respectively. If $D \gg \delta$ then there is no cartogram for G^* with only convex polygons. **b)** A cartogram for G^* with complexity 6 in which every inner face is represented as a hole.

(ii) There is a planar 3*-tree equipped with a weight function, such that every cartogram requires polygonal complexity at least* 7*.*

Proof. Both statements of the lemma follow from Lemma 2.3.1 and Lemma 3.3.1.

Consider the maximally planar Hamiltonian graph G in Figure 3.19. Insert the graph G^* from Figure 3.18 **a)** into each shaded triangle, i.e., identify the outer triangle of a copy of G^* with each triangle shaded in Figure 3.19. The resulting graph G' is still maximally planar and Hamiltonian. Denote the number of vertices in G' by n.

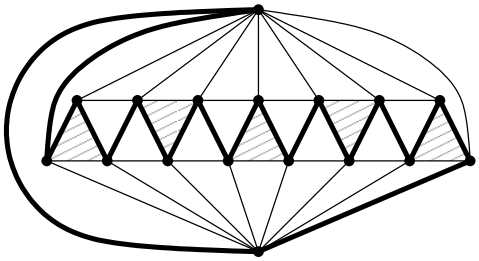

Figure 3.19: A maximally planar Hamiltonian graph G with 17 vertices and five shaded triangles. If every shaded triangle is replaced with the graph G^* in Figure 3.18 **a)**, the resulting graph remains Hamiltonian.

Next, assign every edge of any Hamiltonian cycle in G' to one incident inner face, so that no face is assigned more than once. Insert a new vertex into every assigned face and connect it by an edge to the three incident vertices of G'. Call the resulting graph G'' and set $w(v) = D$ for the vertices a, b, and c in every copy of G^*, and $w(v) = \delta$ for all other vertices, where $D \gg \delta$. Consider any cartogram of G'' with respect to w.

The polygons of the n vertices originally belonging to G' form a cartogram of G' with respect to $w|_{V(G')}$. Moreover, every inner face containing a G''-vertex is represented as a hole. There are exactly n such faces, $n-5$ further inner faces, and one outer face, each of degree three. Applying Lemma 2.3.2 to this cartogram gives that the polygons corresponding to the original vertices contain at least $3n+2(n-5)+6 = 5n-4$ *convex* corners. Additionally, by Lemma 3.3.1 in each copy of G^* there is at least one reflex corner. Hence, the total number of corners among the n polygons corresponding to vertices of G' is at least $5n+1$. Thus, at least one such polygon has complexity 6. This proves (i).

In order to prove (ii), we proceed similarly. Let G be any planar 3-tree with 10 inner faces with pairwise disjoint vertex sets. For example take G to be the graph in Figure 3.20. Put a copy of G^* into each of the 10 faces. Denote the resulting graph by G' and its number of vertices by n. Now, insert a new vertex into every inner face connecting it to the three incident vertices in G'. Set $w(v) = D$ for the vertices a, b, and c in every copy of G^*, and $w(v) = \delta$ for all other vertices, where $D \gg \delta$. Every cartogram of G'' induces a cartogram of G', in which every inner face is represented as a hole. There are $2n-5$ such faces, and one outer face, each of degree three. Applying Lemma 2.3.2 to this cartogram gives that the polygons of vertices in G' contain at least $3(2n-5)+6 = 6n-9$ convex corners. Additionally by Lemma 3.3.1 in each copy of G^* there is at least one reflex corner. Hence, the total number of corners among the n polygons corresponding to vertices of G' is at least $6n+1$. Thus, at least one such polygon has complexity 7. This proves (ii). □

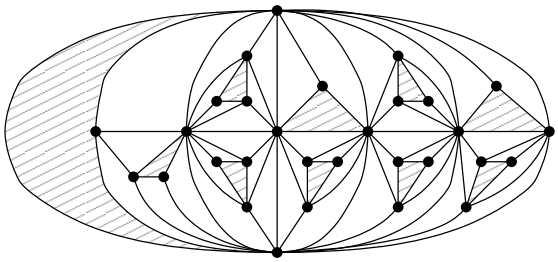

Figure 3.20: A planar 3-tree G with 10 inner faces with pairwise disjoint vertex sets. If every shaded triangle is replaced with the graph G^* in Figure 3.18 **a)**, the resulting graph remains a planar 3-tree.

3.3.1 Hole-Free Cartograms for Planar 3-Trees

Theorem 3.1.7 and Lemma 3.3.2 together show that for a hole-free cartogram of a maximally planar graph 8-sided polygons are always sufficient, while 7-sided polygons

are sometimes necessary. See for instance the summary in Table 3.1. On the other hand, 7-sided polygons are always sufficient (and sometimes necessary) if we do not insist on hole-freeness.

Theorem 3.3.3. *Seven-sided polygons are always sufficient and sometimes necessary for a cartogram of a maximally planar graph with a given weight function on the vertices.*

Proof. It is known [ABF$^+$11b] that every planar graph admits a *point* contact representation Γ_p with 4-sided polygons that realizes any set of given weights, i.e., a point contact cartogram with complexity 4. For example, Figure 3.21 **b)** shows such a cartogram for the planar graph in Figure 3.21 **a)**. Each polygon in Γ_p has three convex and one concave corner, and every (point) contact between two polygons consists of a (convex) corner of one polygon and a side of the other polygon. By a local modification we transform every point contact into a side contact while making two convex corners out of one. Figure 3.21 **c)** shows the side contact representation resulting from the point contact representation in Figure 3.21 **b)**.

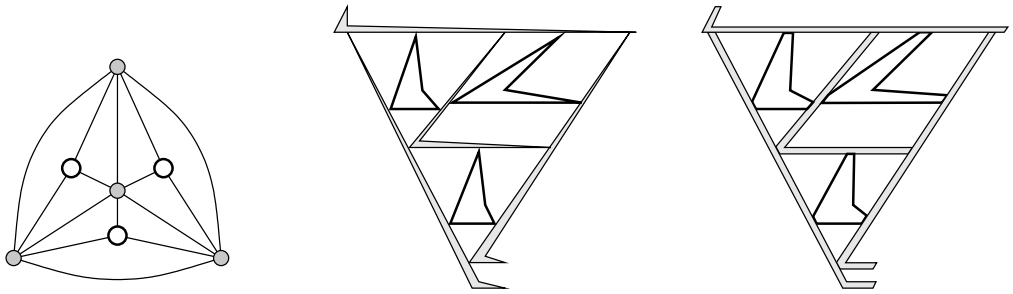

a) **b)** **c)**

Figure 3.21: **a)** A maximally planar graph G in which white and grey vertices have weights D and δ, respectively, where $D \gg \delta$. **b)** A point contact cartogram of G w.r.t. these weights with complexity 4. **c)** A side contact cartogram of G w.r.t. these weights with complexity 7.

The increase in area for every polygon can be easily chosen to be proportional to the weight of the corresponding vertex. Hence, we have proven that 7-sided polygons are always sufficient to obtain a cartogram for any vertex-weighted planar graph. By Lemma 3.3.2 (ii) 7-sided polygons are sometimes necessary, even for planar 3-trees. □

The cartogram from Theorem 3.3.3 contains many and especially big holes. It is a natural question whether or not 7-sided polygons are always sufficient, even if we

insist on hole-freeness. We next answer this question in the affirmative in case the maximally planar graph is a 3-tree.

Theorem 3.3.4. *Let G be any planar 3-tree equipped with a weight function $w : V(G) \to \mathbb{R}^+$. Then a cartogram of G with polygons of complexity at most 7 can be computed in linear time.*

Proof. Throughout this proof, a *chord* of a polygon $\mathcal{P}$ is a straight segment inside $\mathcal{P}$ with endpoints p and q on the boundary of $\mathcal{P}$. A chord pq divides $\mathcal{P}$ into two parts. Consider the part of $\mathcal{P}$ that lies on the left-hand side when traversing the chord from p to q. We refer to the area of that part by the *area left of* pq. Similarly, the area of the other part is called the *area right of* pq. In the construction presented below we always define a chord in a polygon $\mathcal{P}$ such that one of its endpoints is a corner of $\mathcal{P}$ and such that the areas left and right of the chord equal some prescribed values. We then use the following operation on chords, which we call *sliding*, to assure that none of its endpoints is a corner of $\mathcal{P}$, while keeping the area left and right of the chord invariant. More precisely, let pq be a chord of a polygon $\mathcal{P}$, and let a direction, either clockwise or counterclockwise, be given. Then sliding pq in the given direction, say w.l.o.g. counterclockwise, is done by identifying a point p' on the same side of $\mathcal{P}$ as p that is very close to p when traversing the boundary of $\mathcal{P}$ in counterclockwise direction. The area left of $p'q$ is greater than the area left of pq. Similarly, we identify a point q' on the same side of $\mathcal{P}$ as q that is very close to q when traversing the boundary of $\mathcal{P}$ in counterclockwise direction. We choose q' such that the area left of $p'q'$ equals the area left of pq. This is possible, provided we have chosen p' close enough to p. The sliding operation is illustrated with an example in Figure 3.22.

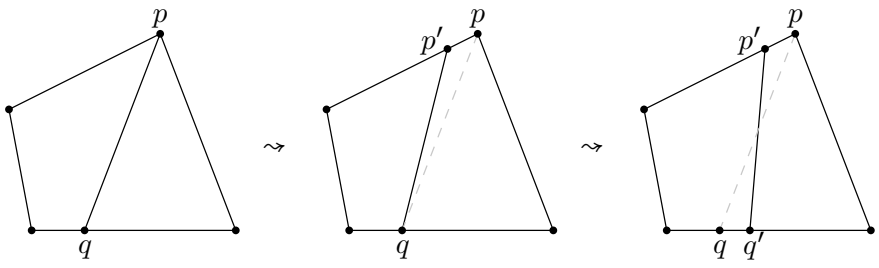

Figure 3.22: Illustration of the sliding operation for chord pq in counterclockwise direction.

Now we build a cartogram of the 3-tree G along its building sequence (c.f. Subsection 1.1.2) starting with any triangle ABC whose area equals $\sum_{v \in V(G)} w(v)$. We maintain the following invariant: The polygons corresponding to the three vertices of a non-empty triangle Δ in G enclose a triangular or convex quadrangular region

R_Δ whose area equals $\sum_{v \in V(G_\Delta) \setminus \Delta} w(v)$. Moreover, every side of R_Δ is constituted by the polygon of only one vertex $v \in \Delta$.

We start by computing the polygons for the three outer vertices v_1, v_2, v_3 of G. Therefore consider the chord of the triangle ABC starting at A and ending at the point p_1 on BC, such that the area left of Ap_1, that is, the area of the triangle ABp_1, equals $w(v_1)$; See the left of Figure 3.23. Next consider the chord of the triangle ACp_1 starting at p_1 and ending at the point p_2 on AC, such that the area of the triangle p_1p_2C equals $w(v_2)$. We apply a sliding operation to the chord p_1p_2 in counterclockwise direction. In the figures we draw a chord before and after a sliding in dashed grey and solid black, respectively.

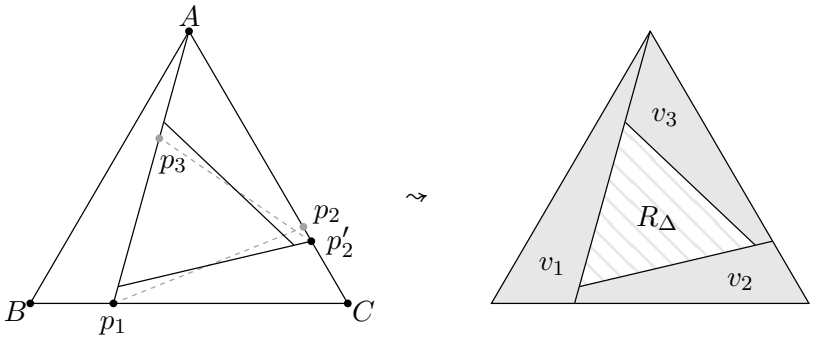

Figure 3.23: Putting the polygons $\mathcal{P}(v_1)$, $\mathcal{P}(v_2)$, and $\mathcal{P}(v_3)$ into the triangle ABC, such that they enclose a triangular region R_Δ of area $\sum_{v \neq v_1, v_2, v_3} w(v)$.

Let p_2' denote the endpoint of the chord on AC after the sliding. Now consider the chord starting at p_2' and ending at the point p_3 on Ap_1, such that the area of the triangle $Ap_2'p_3$ equals $w(v_3)$. We slide that chord in clockwise direction and assign v_i to the polygonal region whose area equals $w(v_i)$, $i = 1, 2, 3$. This leaves a triangular region R_Δ of area $\sum_{v \neq v_1, v_2, v_3} w(v)$. In the remainder of the proof we show how to insert a polygon for a vertex x into a triangular or convex quadrangular region R_Δ.

Case 1: R_Δ is a triangle. We compute the polygon for the vertex x in G_Δ that is adjacent to all three vertices a, b, and c in Δ. Denote the triangle $\{a, b, x\}$ in the graph G by $\Delta(abx)$. Define $w_{abx} = \sum_{v \in G_{\Delta(abx)} \setminus \Delta(abx)} w(v)$ if $\Delta(abx)$ is non-empty, and $w_{abx} = 0$ otherwise. Define $\Delta(axc)$, $\Delta(xbc)$, w_{axc} and w_{xbc} similarly. Moreover, label the corners of R_Δ by A, B, and C, so that A and B lies opposite to the side constituted by $\mathcal{P}(a)$ and $\mathcal{P}(b)$, respectively. See the left of Figure 3.24.

For convenience let us assume that each of w_{abx}, w_{axc}, and w_{xbc} is non-zero. (The case that some of these weights is zero is actually easier.) First, we split off a triangular region of area w_{axc} with a side contact to $\mathcal{P}(a)$ and $\mathcal{P}(c)$. To this end, consider the chord starting at A and ending at the point p_1 on BC, such that the area of the

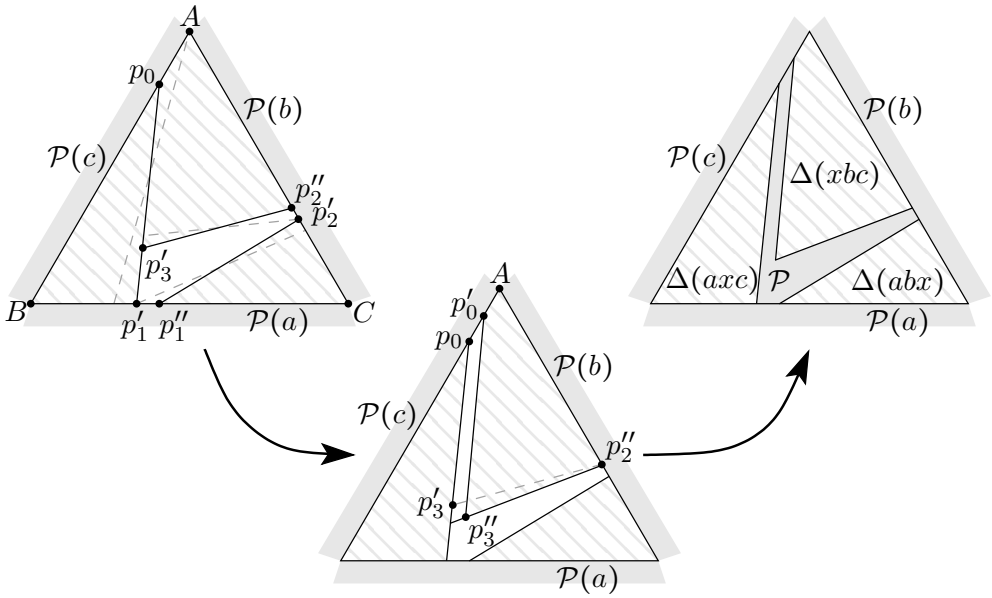

Figure 3.24: Putting polygon $\mathcal{P}(x)$ into triangle R_Δ, such that it leaves a triangular region for $\Delta(axc)$ and $\Delta(abx)$ and a convex quadrangular region for $\Delta(xbc)$, each of the correct area.

triangle ABp_1 equals w_{abx}. Then slide this chord in counterclockwise direction and denote its new endpoint on AB and BC by p_0 and p_1'. Similarly, we obtain a triangle $p_1''p_2'C$ of area w_{abx} by starting with a chord between p_1' and the side AC and sliding it in counterclockwise direction; See the left of Figure 3.24.

Starting with a chord between p_2' and the side p_0p_1' and applying a counterclockwise sliding, we split off a convex quadrangular region with corners A, p_0, p_3', and p_2'' of area w_{xbc}. Next we choose two points p_3'' and p_0' very close to p_3' and p_0, so that the (still convex quadrangular) region with corners by A, p_0', p_3', and p_2'' has area w_{xbc}, too. See the middle of Figure 3.24 for an example. We define $\mathcal{P}(x)$ to be the polygon of complexity 7 left over after removing the triangular and quadrangular regions from R_Δ, as illustrated in the right of Figure 3.24.

Case 2: R_Δ is a convex quadrangle. Let the vertices of Δ and the corners of R_Δ be denoted by a, b, c, and A, B, C, D, respectively. Moreover let $\mathcal{P}(a)$ be the polygon that constitutes two sides of R_Δ. We proceed in a similar way as in Case 1, i.e., we split off triangular or convex quadrangular regions from R_Δ by choosing chords with one fixed endpoint and applying a counterclockwise sliding operation to it. In Figure 3.25 we illustrate the following three sub-cases with one example each. In Case 2a and Case 2b (these are depicted in the left and the right in Figure 3.25, respectively) w_{axc} is less than the area of the triangle ABC. Here, we split off a

triangular region $p_0p_1'B$ of area w_{axc} as in Case 1. If, Case 2c, the area of ABC is at most w_{axc}, we split off a convex quadrangular region by starting with a chord between A and CD, rather than between A and BC. In either case we denote the corners of the split off region distinct from B and C by p_0 and p_1', just as we did in Case 1. The situation is illustrated in Figure 3.25.

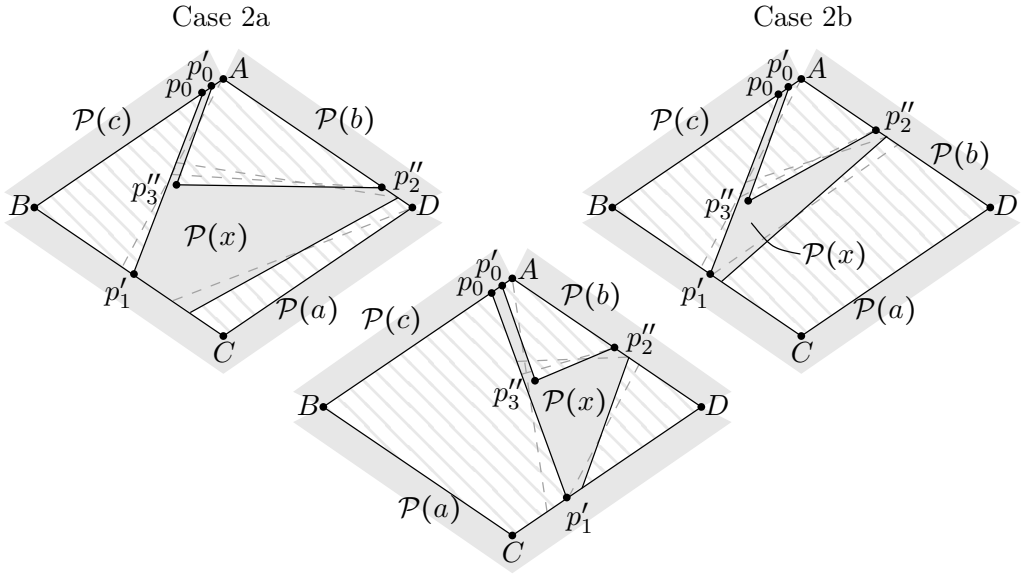

Figure 3.25: Putting polygon $\mathcal{P}(x)$ into convex quadrangle R_Δ, such that it leaves one triangular and two convex quadrangular regions.

If, Case 2a, the area of the triangle $p_1'CD$ is less than w_{abx}, we start with a chord between D and $p_1'C$, obtaining a convex quadrangular region of area w_{abx} after a sliding. In Case 2b (The area of the triangle $p_1'CD$ is at least w_{abx}.) and Case 2c, we obtain this region by starting with a chord between p_1' and AD. The region we split off is convex quadrangular in Case 2b and triangular in Case 2c. Finally, in all cases we compute the points p_2'', p_3'', and p_0' in the very same way as in Case 1. This gives another convex quadrangular region of weight w_{xbc} and the 7-sided polygon $\mathcal{P}(x)$. Again, the special cases, where some (or all) of w_{abx}, w_{axc}, and w_{xbc} are zero are similar and even less complicated. □

3.4 Tackling 4-Connected Maximally Planar Graphs

In this section we discuss cartograms of 4-connected maximally planar graphs. As evident from Table 3.1 for this graph class we have the biggest gaps between the

best lower and upper bounds. For instance, by Theorem 3.3.3 and Theorem 3.1.7 we have that 7-sided and 8-sided polygons are sufficient for a non-rectilinear cartogram with and without holes, respectively. On the other hand, we do not even know a 4-connected maximally planar graph with a weight function that requires 5-sided polygons in any cartogram, even if holes are not allowed.

In the preceding section (Section 3.3) we provide lower bounds on the polygonal complexity of (not necessarily rectilinear) cartograms. In order to enforce many corners, we make use of two little "tricks", or building blocks. At first, we put vertices into inner faces of a maximally planar graph we started with. Disregarding the polygon corresponding to one of these vertices in a representation of the resulting graph, gives a representation of the original graph, where the inner face is represented as a hole. We used this technique in Lemma 2.3.2 and Lemma 3.3.2.

The second trick is to put the graph $G^\star$ from Lemma 3.3.1 into inner faces of a starting graph. This enforces concave corners in every cartogram of the resulting graph. Noting that both operations necessarily produce separating triangles, we may ask for an alternative technique with the same result. Interestingly, we do not know an alternative construction for either of the tricks that does *not* produce separating triangles. We do not even know a planar 4-connected graph with a set of vertex weights, so that every cartogram requires a concave corner. We suspect, that such a graph does not exist.

Conjecture 3.4.1. Every vertex-weighted 4-connected maximally planar graph admits a cartogram with solely convex polygons.

For 4-connected near-triangulations, Conjecture 3.4.1 is implied by the much stronger Conjecture 3.4.4 below. Again, we usually drop the word 4-connected and simply talk about near-triangulations, although we consider only 4-connected near-triangulations in this section. But first, let us recall area-universal rectangular layouts. Due to Eppstein *et al.* [EMSV09] a $\square$-layout Λ is rectilinear area-universal if and only if it is one-sided, c.f. Theorem 3.1.3. Moreover, there are (4-connected) near-triangulations, which do not have a one-sided $\square$-layout, e.g., the graph G Figure 3.26 **a)**, c.f. Lemma 3.1.4. In Figure 3.26 **b)** shows a $\square$-layout Λ of G and Figure 3.26 **c)** shows a cartogram Γ of G realizing Λ. One can show that the layout Λ of G is indeed area-universal, i.e., for every weight function $w : V(G) \to \mathbb{R}^+$ there is a cartogram Γ of G that realizes Λ. Note the difference between *rectilinear* area-universality and area-universality in general.

We want to find necessary conditions for a $\square$-layout being area-universal. We impose any cartogram to fit into a rectangular box, even though this may turn an area-universal layout into one that is not. The shape of the box does not matter because going from one shape to another can be done by a linear transformation,

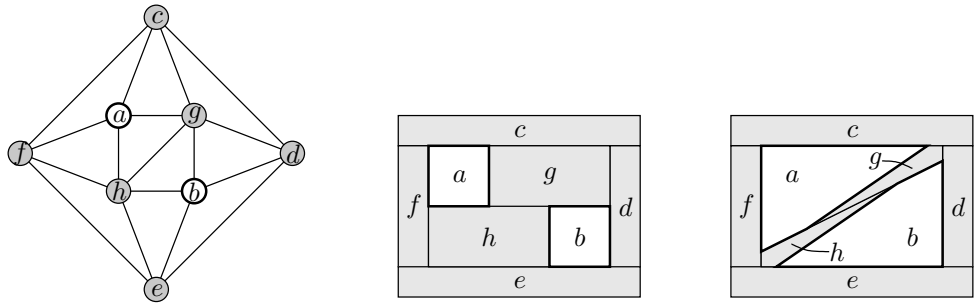

Figure 3.26: **a)** A graph G with grey and white vertices weighted δ and D, respectively. If $D > \delta$ then there is no cartogram for G with only 4-sided rectilinear polygons. **b)** A rectangular layout Λ of G. **c)** A cartogram Γ of G realizing Λ.

which preserves the layout and all areas. Moreover, we ignore the four outer vertices of G, since we can easily define a rectangle of the correct area for each of them in a way that is compatible with the given layout. The removal of these four rectangular polygons leaves a rectangular hole again. In other words, for now we consider G to be an inner triangulation of a simple outer cycle of length at least 4.

It is not the case that *every* $\square$-layout is area-universal. For example consider the $\square$-layout Λ in the left of Figure 3.27 and the weight function w of the represented graph G as marked in the figure, where $D \gg \delta$. For the sake of contradiction, assume that Γ is a cartogram of G w.r.t. w that realizes Λ. W.l.o.g. assume that $3D + 3\delta = 1$, i.e., the total weight of vertices is one, and that the bounding rectangle of Γ is a unit square. Let s denote the segment in Λ that is not one-sided. The weights of the three vertices above s sum up to more than $\frac{1}{3}$. Hence the distances between the square's top-side and the two endpoints of s sum up to more than $\frac{2}{3}$. If the distance on the left is more than $\frac{1}{3}$, then the bottom-right corner of $\mathcal{P}(v_1)$ is very close to left endpoint p of s, since $w(v_1) = \delta$, which is very small. In consequence the top-right corner of $\mathcal{P}(v_4)$ is even closer to p, which means that $\mathcal{P}(v_4)$ is basically a triangle with height less than $\frac{2}{3}$ and base length less than 1. The area of $\mathcal{P}(v_4)$ is $\frac{1}{3} - \delta \approx \frac{1}{3}$ and hence its height is approximately $\frac{2}{3}$ and its base length is approximately 1. In other words, the left endpoint p_1 (and in consequence the right endpoint p_2 of s as well) is in distance $\approx \frac{1}{3}$ to the square's top side. Hence the bottom-left corner of $\mathcal{P}(v_3)$ is very close to p_2 and the major part of the upside and downside of s is constituted by $\mathcal{P}(v_2)$ and $P(v_5)$, respectively. But then $P(v_5)$ has too big area – a contradiction.

If we do not insist on a rectangular outer boundary, then there is a cartogram Γ for G w.r.t. w that realizes Λ. Such a cartogram is shown in the right of Figure 3.27. However, with a reasoning similar to the one above one can show the following: First,

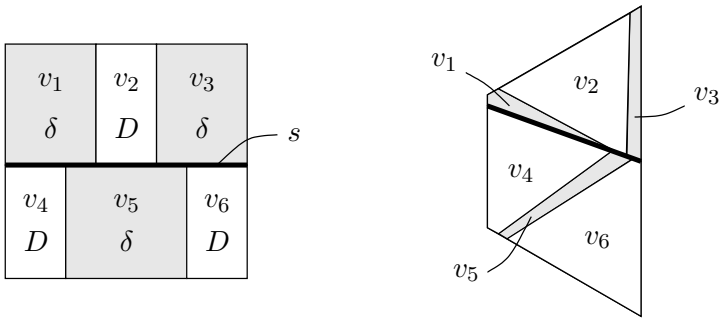

Figure 3.27: Left: A rectangular layout Λ with vertex weights ($D \gg \delta$) that is not realized by any cartogram with rectangular outer boundary. Right: A cartogram that realizes Λ.

in every cartogram Γ that realizes Λ the length of each of the four outer segments of Λ is bounded from below by some function of D. And second, the shortest side of each of the three polygons $\mathcal{P}(v_2)$, $\mathcal{P}(v_4)$, and $\mathcal{P}(v_6)$ is bounded above by some function of δ. For example, in the cartogram in the right of Figure 3.27 the bottom side of $\mathcal{P}(v_2)$ and $\mathcal{P}(v_4)$, as well as the top side of $\mathcal{P}(v_6)$, is "very short". Hence, replacing of these rectangles in Λ, say $\mathcal{R}(v_2)$, by another copy of the same layout, then results in a layout which is not area-universal even in the looser sense. So, not every $\square$-layout is area-universal. Let us define a weakening of one-sided rectangular layouts as follows.

Definition 3.4.2. A $\square$-layout is *two-sided* if every segment s has an interior point that divides s into two one-sided segments.

In other words, a segment s is two-sided if it is the union of two sides of two rectangles, one from either side. For example, the layout in Figure 3.26 **b)** is two-sided, but not one-sided. Indeed, the middle horizontal segment is the union of the bottom side of $\mathcal{R}(g)$ and the top side of $\mathcal{R}(h)$, or equivalently, this segment can be split into two one-sided segments by any point in $\mathcal{R}(g) \cap \mathcal{R}(h)$. Recall that the graph in Figure 3.26 **a)** has no one-sided $\square$-layout.

Next we argue that *every* near-triangulation has a two-sided $\square$-layout. To this end, consider the details $I - IV$ of a $\square$-layout Λ that are depicted in Figure 3.28. Evidently, the bold segment in each details is not one-sided. On the other hand, if a segment is not one-sided, then this is witnessed by a detail like one of $I - IV$. Now consider a horizontal segment, that is two-sided. It may belong to a detail I or III, but not both. Similarly, a vertical two-sided segment does not belong to both a detail II and IV. We summarize the following.

Observation 3.4.3. Each of the following holds.

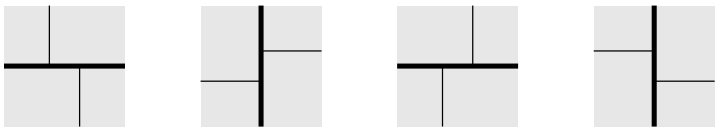

Figure 3.28: Four details of a segment that is not one-sided.

(i) A layout is one-sided if and only if it contains neither of I, II, III, and IV.
(ii) A layout is two-sided if and only if no segment belongs to I and III, or II and IV.
(iii) The layout corresponding to the minimal $\square$-representation contains neither I nor II.
(iv) In the Schnyder wood inherited from the layout every vertex is a leaf in T_1 or T_2 if and only if the layout contains neither I nor IV.

Part (iii) follows directly from the definition of the minimal $\square$-representation, while (iv) maybe needs some explanation. Indeed, applying the definition of the Schnyder wood inherited from a $\square$-representation (see Definition 2.1.2) gives that the rectangle in the lower left of details I and IV has a child in T_1 (the rectangle in the the lower right) and a child in T_2 (the rectangle in the upper left). The if-part of (iv) is equally easy to see. Note that (iv) together with Lemma 3.2.6 gives that the inherited Schnyder wood defines a canonical order which is a Hamiltonian path. Indeed, we can easily read of this ordering. Felsner *et al.* [FFNO11] define the *equatorial line* of a *separating decomposition* and associate a $\square$-representation with it, such that the equatorial line corresponds to a Hamiltonian path in the represented graph. It is straightforward to show, that if we start with the separating decomposition of a layout Λ which contains neither I nor IV, then the resulting $\square$-representation realizes Λ, and that the Hamiltonian path is a canonical order defined by the inherited Schnyder wood. See Figure 3.29 for an illustrating example.

Not every near-triangulation admits a $\square$-layout that contains neither I nor IV, but by Observation 3.4.3 (iii) every such graph admits a $\square$-layout that contains neither I nor II, and hence a two-sided layout. We suspect that two-sided layouts are area-universal, i.e., for every weight function w and two-sided layout Λ there is a cartogram w.r.t. w within a rectangular outer boundary that realizes Λ.

Conjecture 3.4.4. Every two-sided $\square$-layout is area-universal.

Conjecture 3.4.4 implies Conjecture 3.4.1 for 4-connected near-triangulations because a cartogram Γ that realizes a rectangular layout consists solely of *convex* quadrangles. With regard to *rectilinear* cartograms of 4-connected maximally planar

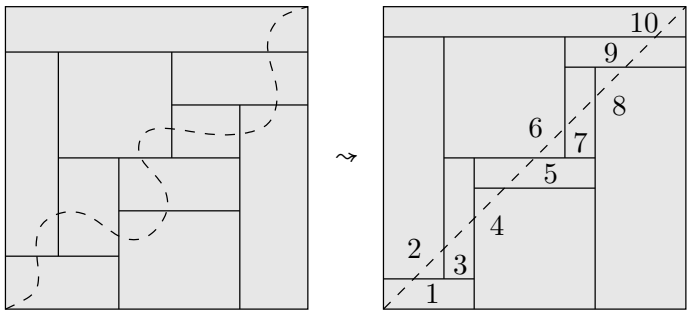

Figure 3.29: A two-sided layout which contains neither I nor IV, its equatorial line, and a canonical order which is a Hamiltonian path.

graphs, we know equally little. It is known that every maximally planar graph admits a rectilinear cartogram of complexity 8 (Theorem 3.1.7) and that there is a vertex weighted 4-connected near-triangulation, of which every rectilinear cartogram requires at least one 6-gon (Lemma 3.1.4). However, we think that the theory of one-sided Hamiltonian cycles, which we discussed in Subsection 3.2.1 may give the right answer here.

Conjecture 3.4.5. Every 4-connected maximally planar graph has a one-sided Hamiltonian cycle. In particular, every such graph has an rectilinear area-universal layout with 6-sided rectilinear polygons.

We remark that Whitney proved [Whi31] that every 4-connected maximally planar graph has a Hamiltonian cycle. However, the cycle he constructs is not one-sided.

Chapter 4

Edge-Intersection Graphs of Grid Paths

In this chapter we investigate edge-intersection graphs of paths in the plane grid, EPG graphs for short. The results presented here are mostly joint work with D. Heldt and K. Knauer, some of which are submitted for publication [HKU10]. EPG graphs were recently introduce by Golumbic, Lipshteyn and Stern [GLS07]. This concept arises from VLSI grid layout problems [BS90], and is a natural generalization of *edge-intersection graphs of paths on degree* 4 *trees* as considered by the same authors [GLS08]. Again, we start with the formal definition.

Definition 4.0.6. An *EPG representation* of a simple graph $G = (V, E)$ is a set $\Gamma = \{P(v) \mid v \in V\}$ of finite paths in the plane square grid, one for each vertex, with the following property:

- Two vertices v and w are connected by an edge in G if and only if the corresponding paths $P(v)$ and $P(w)$ have at least one grid edge in common.

A graph is called an *EPG graph* if it admits an EPG representation.

Figure 4.1 shows $K_{3,3} \setminus e$, i.e., the complete bipartite graph on $3 + 3$ vertices after the removal of one edge, and an EPG representation Γ of it. For better readability of our figures, we draw two or more grid path that share the same grid edge very close to each other, instead of letting them overlap. Moreover, we omit drawing the underlying square grid. Indeed, we consider a grid path to be a finite chain of horizontal and vertical *segments*, alternating in their orientation and joint at their endpoints. The two vertices corresponding to two paths $P(v)$ and $P(w)$ are then adjacent, if and only if $P(v) \cap P(w)$ contains a piece of non-zero length. Note that

$P(v) \cap P(w)$ may contain many isolated points although (v, w) is not an edge. For instance, $P(b) \cap P(d)$, as well as $P(a) \cap P(e)$, consists of just a single point and indeed (b, d), as well as (a, e), is *not* an edge in the graph. Furthermore, a grid path may self-intersect several times; It may even contain the same grid edge more than once. For convenience, we sometimes refer to a vertex and the path that represents this vertex with the same name.

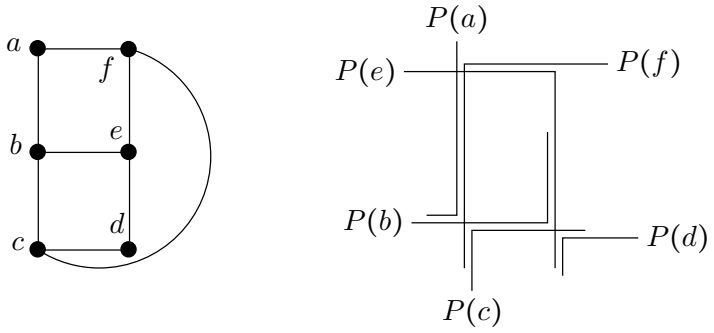

Figure 4.1: The graph $K_{3,3} \setminus e$ and an EPG representation Γ of it.

It is known [GLS07] that every graph is an EPG graph and finding some EPG representation of it is very simple. However such representations can have too many bends (or are too complicated). Therefore it is of interest to consider representations with few bends. A *bend of a grid path* is a point where two consecutive segments are joint, i.e., where the grid path changes its orientation from vertical to horizontal or vice versa, and the number of bends of a path is one less than the number of its segments.

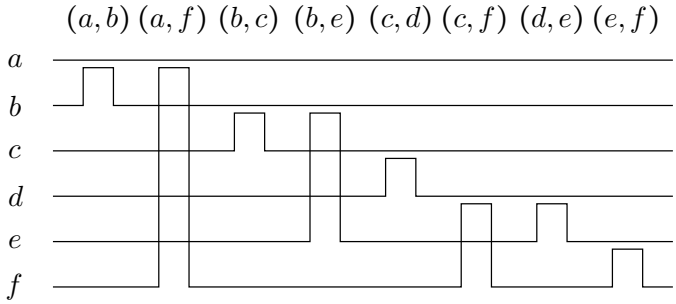

Figure 4.2: Illustration of how to compute any EPG-representation for any graph using the example of $K_{3,3} \setminus e$. Rows and columns correspond to vertices and edges in the graph, respectively. The path for a vertex v runs along the associated with row, except within the columns that correspond to edges of v with vertices whose rows lie above, where the path makes a detour and runs along the row of the other vertex.

Definition 4.0.7. An EPG representation Γ is a *k-bend representation*, for some $k \geq 0$, if every path in Γ has at most k bends. A graph is called a *k-bend graph* if

there is a k-bend representation of it. The *bend-number* of G (written $b(G)$) is the minimum k, such that G is a k-bend graph.

A closer look reveals that every path in the right of Figure 4.1 has at most one bend, and hence it is a 1-bend representation. In consequence, $b(K_{3,3} \setminus e) \leq 1$. Graphs with bend-number at most 1 are called *single-bend graphs* and have been studied independently [GLS07, Rie09]. If $b(G) = 0$ for a graph G, then it has a 0-bend representation, i.e., a representation with horizontal and vertical segments in the plane in which edges correspond to non-trivial intersections of equally oriented segments. If now all grid lines are put in any order on a single real line, then every vertex is represented by an interval and edges correspond to overlapping intervals, i.e., G is an interval graph. The converse is true as well, since every interval representation of an interval graph can be seen as a 0-bend representation using only one grid line. Together we have, $b(G) = 0$ if and only if G is an interval graph.

It has been shown [AS09, BS09] that the bend-number of $K_{m,n}$ is at least $2m - 2$ for sufficiently large n (c.f. Theorem 4.1.7), in particular the bend-number of a graph can be arbitrarily large. We will provide many more examples for this fact in this chapter. Hence it is interesting to determine graphs or graph classes with bounded bend-number. Within the scope of this thesis, we consider the following question:

Question 4.0.8. Given a graph class $\mathcal{G}$, what is the minimum number k, such that every graph $G \in \mathcal{G}$ has a k-bend representation? In other words, what is the maximum bend-number among graphs in $\mathcal{G}$?

Question 4.0.8 has already been considered for some graph classes, including complete bipartite graphs, forests, planar and outer-planar graphs, graphs of fixed tree-width, path-width, degeneracy, or maximum degree. In particular, Asinowski and Suk [AS09] give upper and lower bounds on the bend-number of complete bipartite graphs. Biedl and Stern [BS09] show $b(G) \leq 5$ for planar G and $b(G) \leq 3$ for outer-planar G. They also give upper bounds on $b(G)$ in terms of tree-width, path-width, coloring number and maximum degree of G.

Just as for Question 2.0.5 and Question 3.0.4 in the preceding chapters, a complete answer to Question 4.0.8 for a particular graph class $\mathcal{G}$, would consist of two parts. First, for *every* graph $G \in \mathcal{G}$ we have to find a k-bend representation of G. We call such a proof an "upper bound" for the graph class. And second, for *some* particular graph $G^* \in \mathcal{G}$ we have to show that *every* EPG representation of G^* contains at least one path with k or more bends. Consequently, we call this a "lower bound".

The table below summarizes all upper and lower bounds to Question 4.0.8 that we know of for a set of some graph classes. The columns in Table 4.1 labeled LB and

UB contain the lower and upper bounds, respectively. Note that we have matching upper and lower bounds in several cases.

Remark 4.0.9. Most of the literature, including [AS09, BS09, GLS07], is considering B_k, the *class of k-bend graphs*. Clearly $b(G) \leq k$ just paraphrases $G \in B_k$. However, we prefer to use $b(G)$ rather than B_k.

Graph Class $\mathcal{G}$	Maximum Bend-Number in $\mathcal{G}$			
	LB		UB	
Tree	1	[GLS07]	1	[GLS07]
Outer-Planar	2	[BS09]	2	Thm. 4.2.1
Planar + Bipartite	2	[BS09]	2	[BS09]
Planar + 4-Connected	2		3	Thm. 4.2.4
Planar 3-Tree	3	Lem. 4.2.3	3	Thm. 4.2.2
Planar	3		4	Thm. 4.2.6
$K_{m,n}$, $(m \leq n)$	$\lceil \frac{mn+\sqrt{mn}}{m+n} \rceil$	Lem. 4.1.1	$\lceil \frac{n}{2} \rceil$ $2m-2$	[AS09] [AS09]
Line Graph*	2		2	[BS09]
Maximum Degree Δ	$\lceil \frac{\Delta}{2} \rceil$	Lem. 4.3.9	Δ	Thm. 4.3.8
Tree-Width k	$2k-2$	Lem. 4.3.4	$2k-1$	Cor. 4.3.6
Degeneracy k	$2k-1$	Lem. 4.3.2	$2k-1$	Thm. 4.3.1

Table 4.1: Summary of lower bounds and upper bounds for the maximum bend-number among all graphs within some graph classes. (LB = lower bound, UB = upper bound, * = without subgraphs)

This chapter is organized as follows:

Section 4.1: We provide a first important lemma that will be used later in this chapter, the Lower-Bound-Lemma. Considering complete bipartite graphs, we prove a lower bound on $b(K_{m,m})$, which equals the known upper bound [AS09]. In particular, this proves that for every $k \geq 0$ there is a G with $b(G) = k$, which has been asked [GLS07] and partially answered [AS09] before. Moreover, we derive upper and lower bounds on the bend-number of $K_{m,n}$ for $m \neq n$. We determine $b(K_{m,n})$ in a couple of cases: if n is a specific quadratic function of

m and if n is bigger than some degree 4 polynomial in m, which improves the bounds in previous results [AS09, BS09].

Section 4.2: We deal with planar and outer-planar graphs. We prove a conjecture of Biedl and Stern [BS09], that the bend-number of an outer-planar graph is at most 2, matching the known lower bound. We also provide an upper bound of 3 for the bend-number of subgraphs of planar 3-trees, together with a matching lower bound. This lower bound improves the previously best known value of 2 for general planar graphs. For general planar graphs, we improve the upper bound of 5 [BS09] down to 4. Hence the maximum bend-number among all planar graphs is either 3 or 4. Beforehand, the best known lower and upper bounds were 2 and 5, respectively.

Section 4.3: We provide upper and lower bounds on the bend-number in term of the graph's degeneracy, tree-width, and maximum degree. In case of degeneracy (Subsection 4.3.1) we improve the known upper bound [BS09] and show that it is tight even for bipartite graphs. With regard to the tree-width (Subsection 4.3.2) the presented upper and lower bound differ by an additive 1. Moreover, we provide a matching upper bound for tree-width 2. In Subsection 4.3.3 we improve the upper bound of Biedl and Stern [BS09] for the bend-number in terms of the graph's maximum degree. However, this still leaves a factor of 2 to the best-known lower bound.

Section 4.4: We show that recognizing single-bend graphs is NP-complete, which answers a question that has been raised several times [GLS07, BS09, Rie09]. This is the first NP-completeness result in the field.

Section 4.5: We review the following three graph parameters: the interval-number, the local track-number and the track-number. Figuratively speaking, each of them is measuring "how far" the graph is from being an interval graph. We argue that these concepts are related to the bend-number and compare the maximum bend-number within a graph class with the corresponding maximum interval-number, local track-number, and track-number.

4.1 The Bend-Number of Complete Bipartite Graphs

This section deals with the bend-number of complete bipartite graphs. Throughout this section we denote the bipartition classes of $K_{m,n}$ by $A = \{a_1, \ldots, a_m\}$ and $B = \{b_1, \ldots, b_n\}$ and always assume $2 \leq m \leq n$. For convenience, we depict vertices of A by blue paths and vertices of B by red paths. See Figure 4.3 for an example.

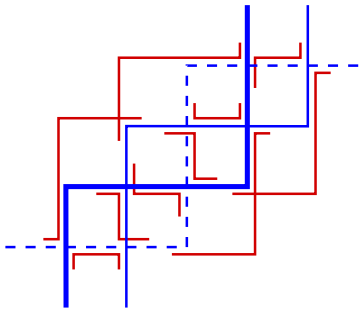

Figure 4.3: A 2-bend representation of $K_{3,10}$.

We start this section with the Lower-Bound-Lemma (Lemma 4.1.1), which is our "standard tool" to derive lower bounds, even though it applies only to special EPG representations. We define the *depth* of an EPG representation Γ to be the maximum number of grid paths in Γ that share a grid edge. For abbreviation we call EPG representations of depth 2, *simple EPG representations*. Clearly, every EPG representation of a triangle-free graph is simple. However, even when constructing representations for general graphs, it is convenient to restrict to depth 2. Indeed, *almost all* EPG representations in the literature so far (except the 2-bend representation of line graphs [BS09]) are simple. Also, many upper bounds in this thesis rely on simple representations. In general, a simple EPG representation may require many more bends than a non-simple one: A good example is K_n. But when restricted to simple representations we quickly obtain a lower bound on the required number of bends.

Lemma 4.1.1 (**Lower-Bound-Lemma**). *Let L denote the set of supporting grid lines of a simple k-bend representation Γ of $G = (V, E)$. Then we have*

$$|E| + |L| \leq (k+1)|V|.$$

Moreover, we can assume w.l.o.g.

$$|L| \geq \sqrt{k|V|}.$$

Proof. W.l.o.g. consider a simple k-bend representation in which every vertex path has exactly k bends. Look at the rightmost or up-most grid edge of each of the $k+1$ segments of each vertex v. If this grid edge is shared by another vertex w (there can be only one in a simple representation), we assign v to the edge $\{v, w\}$ in the graph. This way

- every vertex is assigned to at most $k+1$ edges,
- every edge is assigned at least once, and

- at the rightmost (up-most) grid edge of every line in L either no assignment is done or an edge is assigned twice.

Hence, we have $|E| \leq (k+1)|V| - |L|$. Moreover, since the representation is simple, at most four bends can share a grid point. Thus, at least $\frac{1}{4}k|V|$ grid points support a bend. The minimum number of grid lines crossing in that many grid points is at least $2\sqrt{\frac{1}{4}k|V|} = \sqrt{k|V|}$, which proves the lemma. □

Remark 4.1.2. The Lower-Bound-Lemma can be generalized from simple k-bend representations (depth 2) to arbitrary fixed depth $d \geq 2$. For instance, if the maximum size of a clique in the graph is bounded by d, then every EPG representation of it has depth $\leq d$. Clearly, the lower bounds on the bend-number we obtain this way get worse, but may still be best-possible in some cases.

Although the Lower-Bound-Lemma only depends on the number of vertices and edges, it turns out to be very powerful in several cases. For example, it is tight for $K_{m,n}$ in some particular cases. At first, let us consider the case $m = n$. Asinowski and Suk [AS09] have shown that $b(K_{m,n}) \leq \lceil \frac{\max\{m,n\}}{2} \rceil = \lceil \frac{n}{2} \rceil$. Figure 4.4 illustrates their construction using $K_{5,6}$ as an example. The authors also ask, whether this representation is best-possible in the case $m = n$. We can confirm this by applying the Lower-Bound-Lemma (Lemma 4.1.1). At the same time this solves a conjecture of Golumbic *et al.* [GLS07], asking whether for every $k \geq 0$ there is a graph G with $b(G) = k$.

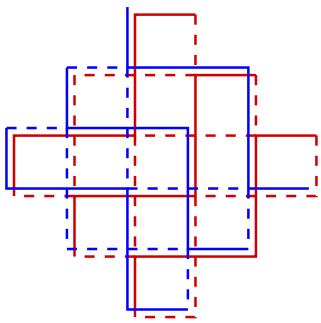

Figure 4.4: A 3-bend representation of $K_{5,6}$. The smaller bipartition class A is blue.

Theorem 4.1.3. *For all $2 \leq m \leq n$, such that $m = n$ or $m + 1 = n$ is even, we have*

$$b(K_{m,n}) = \lceil \frac{m}{2} \rceil.$$

Proof. We prove that $K_{m,2\lceil \frac{m}{2} \rceil}$ has no k-bend representation with $k < \lceil \frac{m}{2} \rceil$. Assuming the existence of such a representation, the Lower-Bound-Lemma (Lemma 4.1.1)

gives $2m\lceil\frac{m}{2}\rceil + |L| \leq \lceil\frac{m}{2}\rceil(m + 2\lceil\frac{m}{2}\rceil)$. If m is even we obtain $m^2 + |L| \leq m^2$ which is a contradiction since $|L| \geq 1$. For odd m we calculate $|L| \leq \frac{1}{2}(m+1)$. But by the Lower-Bound-Lemma we can assume $|L| \geq \sqrt{\frac{1}{2}(m-1)(2m+1)}$ which leads to a contradiction for $m \geq 2$.

The upper bound follows from $b(K_{m,n}) \leq \lceil\frac{\max\{m,n\}}{2}\rceil$ [AS09]. □

When considering $K_{m,n}$ with increasing n compared to m, the upper bound $\lceil\frac{\max\{m,n\}}{2}\rceil$ of Asinowski and Suk gets worse. Being interested in the behavior of $b(K_{m,n})$, we now determine the exact value for a certain $n \in \theta(m^2)$.

Theorem 4.1.4. *Let $m \geq 3$. For even m we have $b(K_{m,(m+1)(m-2)}) = m-1$ and for odd m we have $b(K_{m,m(m-2)}) = m-1$.*

Proof. Let m be even. We use a braid-like path P with $m-1$ bends and $m/2-1$ crossings as a template for every path in A. We represent each of the m vertices in A by a copy of P translated by a very small amount along the diagonal. See the left of Figure 4.5 for an illustration. The vertices of B are represented by small staircases with $m-1$ bends, each interlaced around a bend or a crossing of P. At every bend of P, except the braid's turning point, we interlace m red staircases and at every crossing another two staircases. This is illustrated in the right of Figure 4.5. In total we have $(m-2)\cdot m + (m/2-1)\cdot 2 = (m-2)(m+1)$ vertices in B.

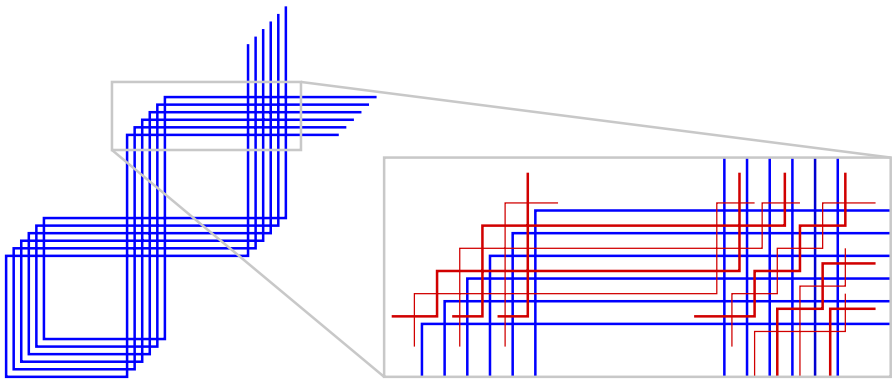

Figure 4.5: Representing $K_{m,(m+1)(m-2)}$ (m even) with $m-1$ bends: Vertices in A (blue) are represented by braid-like paths. Vertices in B (red) are represented by staircases interlaced at the bends and crossings.

For odd m, first represent $K_{m-1,m(m-3)}$ as described in the even case with $m-2$ bends for each path. Then add the missing vertex $a_m \in A$ by a snake-like $(m-1)$-bend path like the dashed one on the left of Figure 4.6. The vertical end of the so far existing paths in $B \setminus \{b^*\}$ can be extended to reach a_m and endowed with another

bend to connect to it. The special vertex b^* is extended horizontally to reach a_m as depicted in Figure 4.6.

Finally, every path $a_i \in A \setminus \{a_m\}$ receives a last bend at its vertical or horizontal end, depending on the parity of i. We obtain several new crossings where an additional set of m paths may be threaded in; see the right of Figure 4.6. This way B contains $m(m-2)$ vertices.

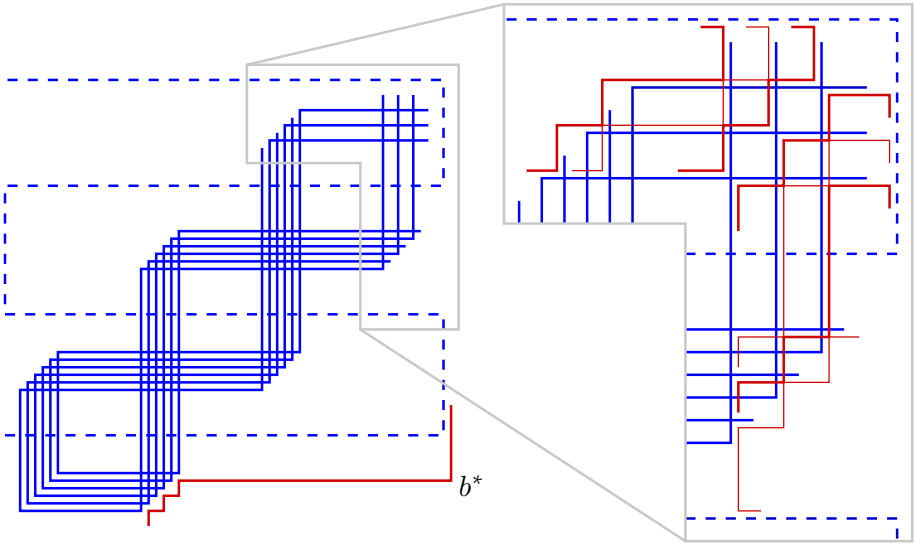

Figure 4.6: Representing $K_{m,m(m-2)}$ (m odd) with $m-1$ bends: Vertices in $A \setminus \{a_m\}$ (blue) are represented by braid-like paths and a_m (dashed) by a snake-like path. The first $m(m-3)$ vertices in B (red) are interlaced as in the even case and extended vertically to reach a_m. Only b^* is extended as depicted. Another m vertices in B are threaded in as shown on the right.

The above constructions show that both bend-numbers are at most $m-1$. Equality follows from a straightforward application of the Lower-Bound-Lemma. □

Theorem 4.1.4 only holds if $m \geq 3$. The bend-number of $K_{2,n}$ has been determined for all n [AS09]: $b(K_{2,n}) = 2$ if and only if $n \geq 5$ and $b(K_{2,n}) = 0$ if and only if $n = 0, 1$.

Now we investigate the extremal case where n gets very large compared to m. Asinowski and Suk [AS09] showed that $b(K_{m,n}) = 2m-2$ for every $n \geq N$ with $N \in \Omega(m^m)$; see Figure 4.9 for a representation. Later on, Biedl and Stern [BS09] improved this to $N = 4m^4 - 8m^3 + 2m^2 + 2m + 1$. In Theorem 4.1.7 we lower this once more to $N = m^4 - 2m^3 + 5m^2 - 4m + 1$ and show in Theorem 4.1.8 that this leaves at most a quadratic gap to the true value, disproving the conjecture of Biedl and Stern [BS09], that the true value is $\mathcal{O}(m^2)$. We begin with bounding the number of

crossings of two paths which have a given odd number of bends. This bound may be of interest on its own.

Lemma 4.1.5. *Two $(2m-1)$-bend paths cross in at most $m(m+1)$ points. This is tight.*

Proof. Consider two given $(2m-1)$-bend paths v and w. Both have exactly m horizontal and m vertical segments. We color the vertical segments of v and the horizontal segments of w black and the remaining segments grey. Along each path we index the segments, starting with its black end, i.e., v_1 and w_1 are black and $v = (v_1, \dots, v_{2m})$ and $w = (w_1, \dots, w_{2m})$. Now every crossing is monochromatic, either black with odd indices or grey with even indices. We partition the pairs of segments that have the same color but come from different paths into four sets. Set $\mathcal{B}$ contains all black pairs that *do* cross and $\overline{\mathcal{B}}$ all black pairs that *do not* cross. Similarly $\mathcal{G}$ and $\overline{\mathcal{G}}$ are defined for grey segments.

Consider a black crossing $\{v_i, w_j\} \in \mathcal{B}$ and the grid line ℓ containing v_i. Each of $v_{i-1}, v_{i+1}, w_{j-1}, w_{j+1}$, if existent, is grey and lies completely on one side of ℓ. Moreover w_{j-1} and w_{j+1} lie on different sides since w_j crosses ℓ. Now consider v_{i-1} (or v_{i+1}) and the w-segment on the other side of ℓ. This pair is in $\overline{\mathcal{G}}$ since they lie on different sides of ℓ and thus do not cross. This way, we associate up to two grey non-crossings with every black crossing, even if there are more. See Figure 4.7 for an example.

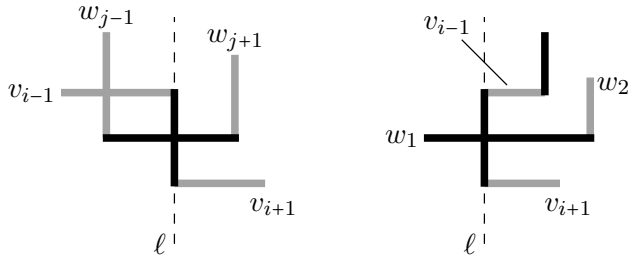

Figure 4.7: A black crossing is associated with every pair of grey segments from different sides of ℓ: The black crossing on the left is associated with $\{v_{i-1}, w_{j+1}\}$ and $\{v_{i+1}, w_{j-1}\}$. The black crossing on the right is associated with no grey non-crossing.

Next we partition the black crossing pairs $\mathcal{B}$ in two ways. Above we have associated every such pair with two, one or no grey non-crossing. For $i = 0, 1, 2$ we put a black crossing into $\mathcal{B}_i$ if it is associated with i grey non-crossings. Secondly, we write $\mathcal{B}(v_1)$ for the set of black crossings v_1 participates in and do the same with w_1. We denote $\mathcal{B}_i(w_1) := \mathcal{B}_i \cap \mathcal{B}(w_1)$ for $i = 0, 1, 2$. Then $\mathcal{B}_0 = \mathcal{B}_0(w_1)$ and $\mathcal{B}_1 = (\mathcal{B}(v_1) \setminus \mathcal{B}(w_1)) \cup \mathcal{B}_1(w_1)$. Note that every grey non-crossing is associated with at most two black crossings and hence we have $|\mathcal{B}_1| + 2|\mathcal{B}_2| \leq 2|\overline{\mathcal{G}}|$. This leads to:

$$2|\mathcal{B}| \leq 2|\overline{\mathcal{G}}| + |\mathcal{B}(v_1) \setminus \mathcal{B}(w_1)| + |\mathcal{B}_1(w_1)| + 2|\mathcal{B}_0(w_1)|$$

Now observe the following: When tracing the path v between any two black segments contributing to a $\mathcal{B}_0(w_1)$-crossing there is a black segment of v that either participates in a $\mathcal{B}_2(w_1)$-crossing or does not cross w_1 at all. Hence, because v has only m black segments, we have $2|\mathcal{B}_0(w_1)|-1+|\mathcal{B}_1(w_1)| \leq m$, as well as, $|\mathcal{B}(v_1)\setminus\mathcal{B}(w_1)| \leq m-1$.

Plugging both into the above inequality, we calculate $2|\mathcal{B}| \leq 2|\overline{\mathcal{G}}| + 2m$. Thus, $|\mathcal{B}| - |\overline{\mathcal{G}}| \leq m$. Now adding m^2 on both sides we obtain: $|\mathcal{B}| + |\mathcal{G}| \leq m^2 + m$, where $\mathcal{G}$ denotes the set of grey crossings. □

Figure 4.8 shows that two $(2m-1)$-bend paths can indeed have $m(m+1)$ crossings.

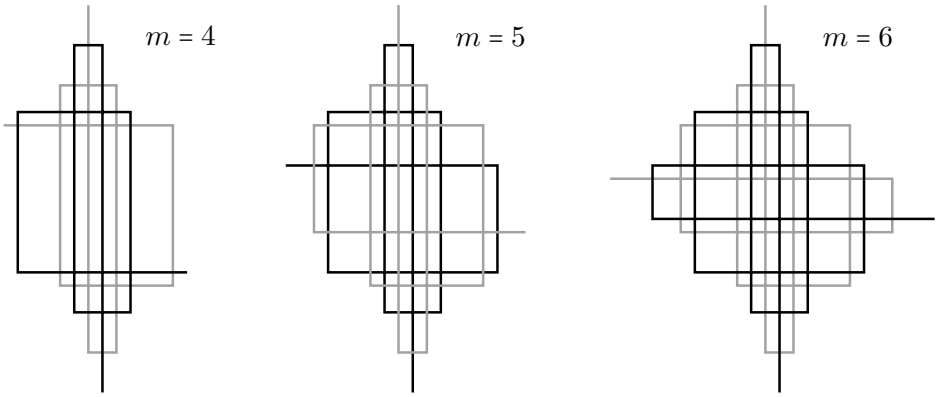

Figure 4.8: Two $(2m-1)$-bend paths can cross in $m(m+1)$ points.

Part (i) of the following lemma is due to Biedl and Stern [BS09].

Lemma 4.1.6. *Consider a k-bend representation of $K_{m,n}$ and a subset B' of B, such that every vertex in B' establishes two of its incidences with either a single segment or two consecutive segments. Then*

(i) $|B'| \leq 2\binom{m(k+1)}{2}$ *and*
(ii) $|B'| \leq (k+1)(\frac{k+3}{2}\binom{m}{2} + 2m)$ *if k is odd.*

Proof. Let s and s' be segments of distinct vertices in A and $b \in B'$ a path intersecting s and s' with either the same or consecutive segments. In the first case, the corresponding segment of b must contain an endpoint of each s and s'. Since B is an independent set, b is the only vertex that intersects s and s' with the same segment. In the second case b must have a bend where the grid lines through s and s' intersect. Here, beside b at most one other vertex in B can intersect s and s' with consecutive segments.

Since there are at most $\binom{m(k+1)}{2}$ pairs of segments of vertices in A, there are at most twice as many vertices in B'. This concludes part (i).

For a more careful analysis, in the case of odd k, observe that *two* paths in B can intersect s and s' with consecutive segments only if s and s' *cross*. By Lemma 4.1.5, two k-bend paths can cross in at most $(k+1)(k+3)/4$ points if k is odd. Hence there are at most $\binom{m}{2}(k+1)(k+3)/4$ crossings between vertices in A.

If s and s' are perpendicular but *do not cross*, at most *one* vertex $b \in B$ can intersect both with consecutive segments and at least one such segment of b must contain an endpoint of s or s'. Since there are at most $2m(k+1)$ endpoints of segments of vertices in A, we conclude that $|B'| \leq 2\binom{m}{2}(k+1)(k+3)/4 + 2m(k+1) = (k+1)(\binom{m}{2}(k+3)/2 + 2m)$. □

Theorem 4.1.7. *We have $b(K_{m,n}) = 2m-2$ for all $n > m^4 - 2m^3 + 5m^2 - 4m$.*

Proof. If $m = 1$, then $K_{m,n}$ is a star and thus an interval graph, i.e., $b(K_{1,n}) = 0$ for all $n > 0$.

For $m > 1$ suppose $b(K_{m,n}) \leq 2m-3$. Then applying Lemma 4.1.6–(ii) with $k = 2m-3$ yields, that at most $N := (2m-2)(\binom{m}{2}m + 2m) = m^4 - 2m^3 + 5m^2 - 4m$ vertices in B can establish two of its incidences with the same or consecutive segments. Hence if $|B| > N$, there must be a vertex $b \in B$ with at least one "empty" segment between any two "non-empty" segments. Moreover every segment of b establishes at most one incidence. Since b has degree m, we conclude that b must have at least $2m-1$ segments.

Figure 4.9 shows that $b(K_{m,n}) \leq 2m-2$, regardless of n. □

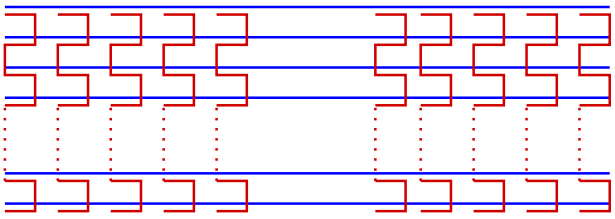

Figure 4.9: A representation verifying $b(K_{m,n}) \leq 2m-2$.

Biedl and Stern [BS09] conjectured that Theorem 4.1.7 is already true for all $n \geq N$ with $N \in \mathcal{O}(m^2)$. We disprove this and show that Theorem 4.1.7 is not far from being tight.

Theorem 4.1.8. *If $n \leq m^4 - 2m^3 + \frac{5}{2}m^2 - 2m - 4$ then $b(K_{m,n}) \leq 2m-3$. Note that this leaves only a quadratic discrepancy to the bound of the preceding theorem.*

Proof. We divide A equally into A_1 and A_2. We use the two $(2m-3)$-bend paths P_1 and P_2 from the tight example in Figure 4.8 as templates for vertices in A_1 and A_2, respectively. Note that each P_i has $\binom{m-1}{2}$ crossings and P_1 and P_2 cross $m(m-1)$ times.

For $i = 1, 2$ every $a \in A_i$ runs within a small distance along P_i. We ensure that at every bend of P_i every pair of paths in A_i crosses. This way, every such pair crosses $2\binom{m-1}{2} + 2m - 3 = m(m-1) - 1$ times. A pair of vertices, one from A_1 and the other from A_2, crosses $m(m-1)$ times. Hence the total number of crossings between vertices in A is given by $\lfloor \frac{m}{2} \rfloor \lceil \frac{m}{2} \rceil m(m-1) + \left(\binom{\lceil m/2 \rceil}{2} + \binom{\lfloor m/2 \rfloor}{2}\right)(m(m-1)-1)$. At every crossing we can interlace two vertices of B.

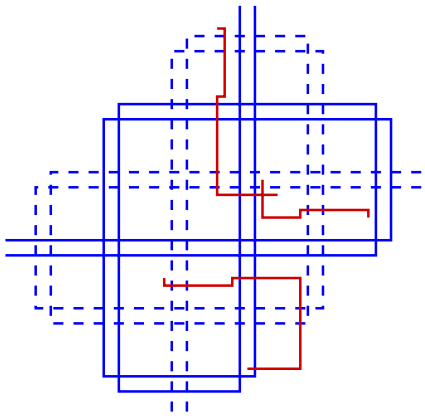

Figure 4.10: A $(2m-3)$-bend representation of $K_{m,n}$, where $n \leq m^4 - 2m^3 + \frac{5}{2}m^2 - 2m - 4$: Every vertex in A (blue) runs within a small distance along one of two paths with the maximum number of crossings. Every vertex in B (red) is interlaced around one crossing or two endpoints of blue segments.

Moreover every endpoint of a segment from A (except the ends of the paths and the eight endpoints of the topmost, rightmost, bottommost, and leftmost segment) can be used to interlace vertices of B. There are $2m(2m-3) - 8$ suitable endpoints. Interlacing one vertex uses two of them.

Figure 4.10 suggests how to insert one vertex $b \in B$ with two suitable endpoints and two vertices $b, b' \in B$ at one crossing. By doing this we can insert $n = \lfloor m^4 - 2m^3 + \frac{5}{2}m^2 - 2m - 4 \rfloor$ vertices into B. ☐

By Theorem 4.1.8 we have $b(K_{m,n}) \leq 2m-3$ if $n \leq m^4 - 2m^3 + \frac{5}{2}m^2 - 2m - 4$. On the other hand, by Theorem 4.1.7 we have $b(K_{m,n}) = 2m-2$ if $n > m^4 - 2m^3 + 5m^2 - 4m$. The bend-number of $K_{m,n}$ for every $n \in \{m^4 - 2m^3 + \frac{5}{2}m^2 - 2m - 3, \ldots, m^4 - 2m^3 + 5m^2 - 4m\}$ is either $2m-3$ or $2m-2$. But determining the exact value remains open.

Question 4.1.9. What is the maximum n for which $b(K_{m,n}) \leq 2m - 3$?

The representations we have given for complete bipartite graphs can naturally be extended to different values of n, though they might not be optimal anymore. For instance, the first and the second representation in Theorem 4.1.4 yields that $b(K_{5,24}) \leq 5$ and $b(K_{5,25}) \leq 6$, respectively. Moreover Figure 4.3 certifies $b(K_{3,10}) \leq 2$, i.e., $b(K_{3,5}) \leq 2$. In Figure 4.11 we sketch a region which contains the graph of $b(K_{5,n})$.

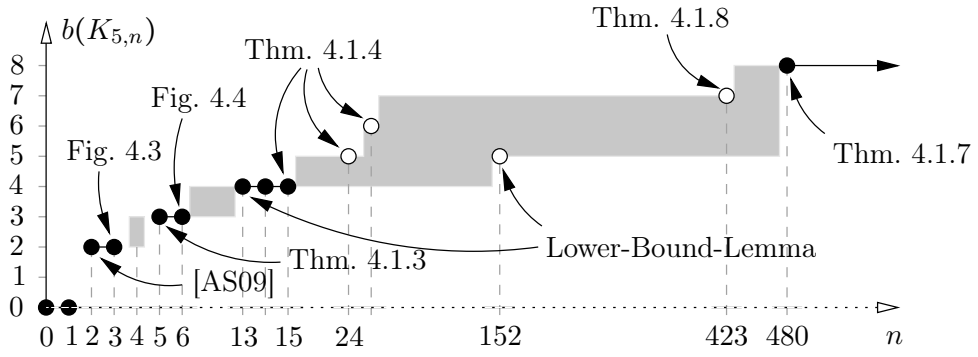

Figure 4.11: Upper and lower bounds for $b(K_{5,n})$. The filled circles and the straight line are values of $b(K_{5,n})$.

With Theorems 4.1.3, 4.1.4, and 4.1.7 we have determined the bend-number of $K_{m,n}$ for n being a fixed polynomial in m of degree 1, 2, and 4, respectively. From these results and the Lower-Bound-Lemma we get a first approximation of the function $b(K_{m,.})$. In the light of these results it seems interesting to determine a further exact value of this function for n being a polynomial in m of degree 3.

Question 4.1.10. What is the behavior of $b(K_{m,n})$ for $n \in \theta(m^3)$?

4.2 The Bend-Number of Planar and Outer-Planar Graphs

This section deals with planar and outer-planar graphs. In particular, we show that every outer-planar graph has a 2-bend representation. This confirms a conjecture of Biedl and Stern [BS09], who presented a 3-bend representation for every such graph and moreover an outer-planar graph that does not admit a 1-bend representation. This graph is shown in Figure 4.12. We conclude that the maximum bend-number among all outer-planar graphs is 2. We also improve the previously best known result for planar graphs, i.e., we construct a 4-bend representation for every planar

graph. On the other hand we provide a planar graph, which does not have a 2-bend representation. Hence, the maximum bend-number among all planar graphs is either 3 or 4.

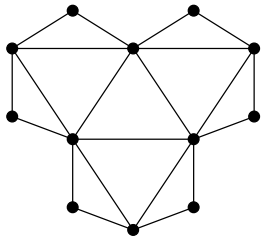

Figure 4.12: An outer-planar graph (a so-called snowflake graph) with bend-number 2. This example is due to Biedl and Stern [BS09].

Many proofs in this and the next section make use of a certain building sequence (c.f. Section 1.1) of the graph. When arguing about an EPG representation, we need the concept of displayed parts. Recall that a k-bend path consists of $k+1$ segments, each of which is either horizontal and vertical. A *sub-segment* is a connected subset of one segment and a set of sub-segments in an EPG representation is called a *part*. A part Γ' of a representation Γ is said to be *displayed* if no sub-segment in $\Gamma \setminus \Gamma'$ overlaps a sub-segment in Γ'. In particular, if we restrict a representation Γ to any set of grid edges, then this restriction Γ' is a displayed part of Γ. Indeed, every displayed part of Γ arises this way. For example, Figure 4.13 shows an EPG representation with a blue part and a red part. We highlight a part by underlining the corresponding grid edges in grey. The blue part is displayed, while the red part is not, because the corresponding grid edges contain some red sub-segment that is not red. Note that segments not belonging to some displayed part may well cross sub-segments of the part.

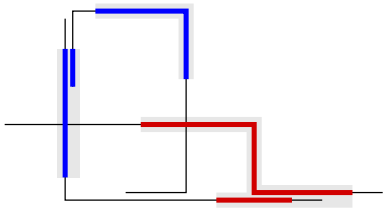

Figure 4.13: An EPG representation with two parts highlighted in blue and in red, respectively. The blue part is displayed, while the red part is not.

We say that two horizontal sub-segments in an EPG representation *see each other* if there is a vertical grid line crossing both sub-segments. Similarly, two vertical sub-segments see each other if there is a horizontal grid line crossing both.

We now present a 2-bend representation Γ for every outer-planar graph G. Note that the representation is indeed simple.

Theorem 4.2.1. *For every outer-planar graph G we have $b(G) \leq 2$.*

Proof. Outer-planar graphs have tree-width at most 2. We construct a 2-bend representation of the outer-planar graph G along a building sequence of G associated with its tree-width (c.f. Subsection 1.1.2). In particular, we consider a maximally outer-planar $\tilde{G}$ which is a super-graph of G, i.e., $\tilde{G}$ is a 2-tree. Let $(v_1, \ldots, v_n)$ be a vertex ordering of $\tilde{G}$, such that $\deg_{G_i}(v_i) \leq 2$ for every $i = 1, \ldots, n$. We maintain that for every outer edge (u, v) of G_i there is a displayed part as in one of the cases that are illustrated in the top row of Figure 4.14. That is, u and v have displayed sub-segments with the same orientation that see each other (right in Figure 4.14) or with distinct orientation that share a grid point (left in Figure 4.14). Moreover, we require at all time that all the displayed parts are pairwise disjoint.

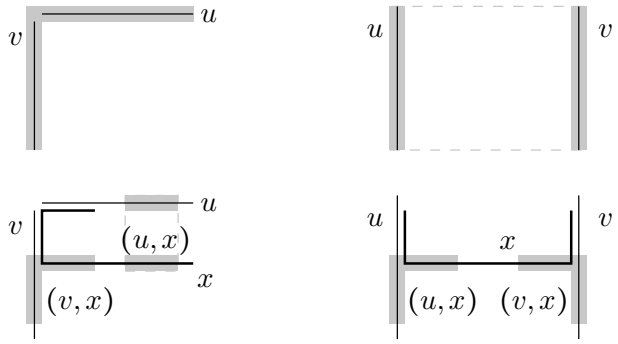

Figure 4.14: Building a 2-bend representation of an outer-planar graph, a vertex x is attached to the edge $(u, v) \in E(\tilde{G})$: The displayed parts for (u, v) and $(u, x), (v, x)$ are highlighted in the upper and lower row respectively.

The second row shows how to introduce a vertex x and maintain the invariant for the new edges (u, x) and (v, x) in $\tilde{G}$. If (u, x) or (v, x) is not an edge in G, the corresponding segment of x is simply omitted. To omit only the edge (v, x) in the left of Figure 4.14 just interchange the roles of u and v. □

Let us illustrate the construction in the proof of Theorem 4.2.1 with an example. Figure 4.15 **a)** shows an outer-planar graph G, as well as a maximally outer-planar super-graph $\tilde{G}$ of G and a vertex ordering $(v_1, \ldots, v_6)$ with $\deg_{G_i}(v_i) \leq 2$ for $i = i, \ldots, 6$. In Figure 4.15 **b)** a 2-bend representation Γ of the subgraph G_3 on v_1, v_2 and v_3 is given. The three displayed parts for the three edges (v_1, v_2), (v_1, v_3), and (v_2, v_3) in $\tilde{G}$ are highlighted. Note that (v_1, v_2) is not an edge in G, but this edge could be established in Γ by adding a bend to the path of v_3 in the lower left

corner. The construction of a 2-bend representation of G along its building sequence $G_3 \subset \cdots \subset G_6 = G$ is illustrated in Figure 4.15 **c)** and **d)**.

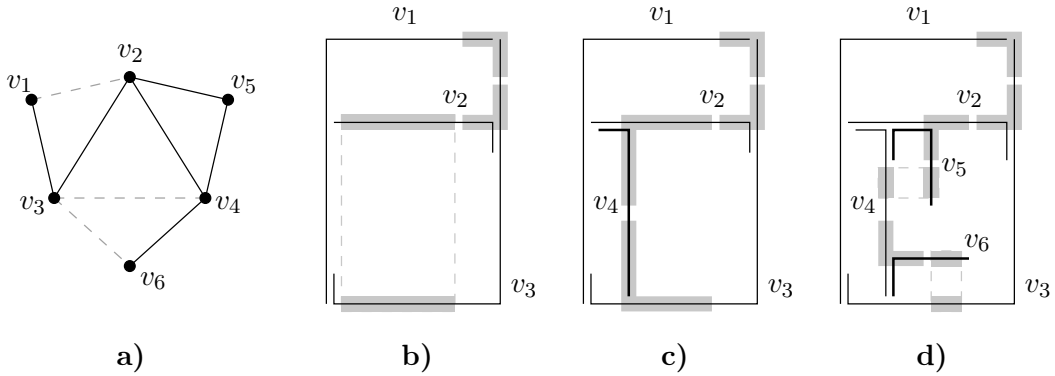

Figure 4.15: **a)** An outer-planar graph G with a vertex ordering corresponding to a maximally outer-planar super-graph $\tilde{G}$ of G. **b**) 2-bend representation of $G_3 = G[v_1, v_2, v_3]$. **c**) 2-bend representation of G_4. **d**) 2-bend representation of $G = G_6$. In each partial representation the displayed parts are highlighted.

Let us consider two further special classes of planar graphs, namely (subgraphs of) planar 3-trees and (subgraphs of) 4-connected maximally planar graphs. For graphs of either class we present 3-bend representations.

Theorem 4.2.2. *For every planar graph G with $tw(G) \leq 3$ we have $b(G) \leq 3$.*

Proof. Recall from Subsection 1.1.2 that there is a planar 3-tree $\tilde{G}$, which is a super-graph of G. There is a vertex ordering $(v_1, \ldots, v_n)$, such that $\tilde{G}_3$ is a triangle, and every vertex v_i is connected to the three vertices u, v, w of a facial triangle in $\tilde{G}_{i-1}$, for $i \geq 4$. The triangle $\{u, v, w\}$ is then no longer facial in $\tilde{G}_i$ and hence no second vertex may be attached to it. As in the proof for Theorem 4.2.1, we build a 3-bend representation of G concurrently with the building sequence $G_3 \subset \cdots \subset G_n$ of G w.r.t. the vertex ordering $(v_1, \ldots, v_n)$. We maintain the following invariant on the 3-bend representation Γ_i of G_i, for $i \geq 3$:

(I) Every path in Γ_i has a horizontal *and* a vertical displayed sub-segment, and
(II) every facial triangle $\{u, v, w\}$ of $\tilde{G}_i$ has two vertices, say u and v, such that
   a) there is a displayed part consisting of a sub-segment of $P(u)$ edge-intersecting a sub-segment of $P(v)$; see Figure 4.16 **a)**,
   b) or a displayed *entire* segment of $P(v)$ and a displayed sub-segment of $P(u)$ cross; see Figure 4.16 **a)**.

Moreover, we require that all the displayed parts above are pairwise disjoint and that the entire segment in every cross in b) *can not* see displayed parts from (I). However, we need this assumption only in the case in Figure 4.16 **c)**.

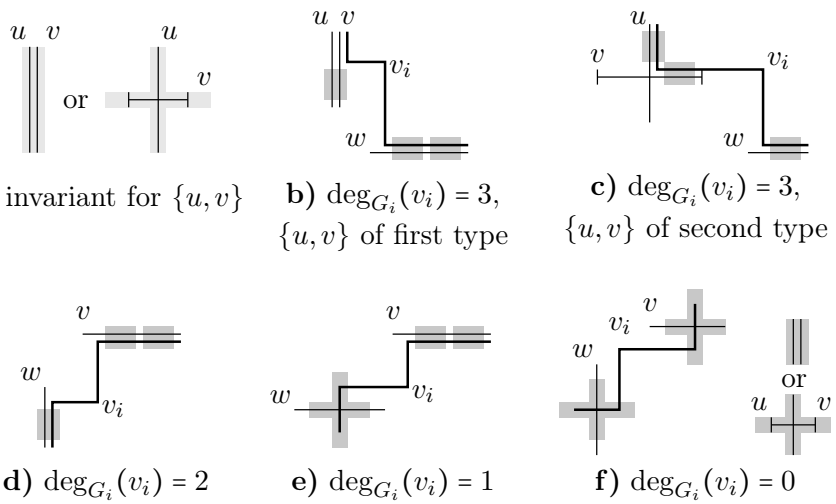

Figure 4.16: Building a 3-bend representation of a planar graph with tree-width 3, a vertex v_i is attached to the facial triangle $\{u, v, w\}$ in $\tilde{G}_{i-1}$: In **a)** the two types of invariant for $\{u, v\}$ are shown. In **b)**–**f)** it is shown how to insert the new vertex v_i (drawn bold) depending on its degree in G_i and the invariant of $\{u, v\}$. The invariants for the three new facial triangles $\{u, v, v_i\}$, $\{u, v_i, w\}$, and $\{v_i, v, w\}$ in $\tilde{G}_i$ are highlighted.

First, we represent the subgraph G_3 of G as follows. Figure 4.17 shows a 3-bend representation of a triangle $\{v_1, v_2, v_3\}$, which satisfies invariants (I) and (II) as indicated by the grey sets of grid lines. If G_3 is not a triangle, i.e., one or more of the three edges is not present in G_3, the representation can easily be modified to represent G_3. To be precise, if some edge at v_3 is not present, one segment of the path of v_3 is omitted. If the edge (v_1, v_2) is not in G_3, then one segment of the path for v_2 is shortened a bit. The result is the representation Γ_3 of G_3, which still satisfies the invariants.

For $i \geq 4$, the path for vertex v_i is introduced to Γ_{i-1} according to the degree of v_i in G_i and the type of invariant for the facial triangle $\{u, v, w\}$ in $\tilde{G}_{i-1}$ that v_i is connected to. Figure 4.16 **b)**–**f)** shows all five cases and how to introduce v_i, which is illustrated by the bold path. Consider in particular the case that v_i has an edge with each of u, v, and w, and moreover the triangle $\{u, v, w\}$ has a cross; see Figure 4.16 **c)**. Here we use that the v-segment in the cross can not see the displayed sub-segment of w. Otherwise the new path corresponding to v_i would not have a horizontal displayed sub-segment.

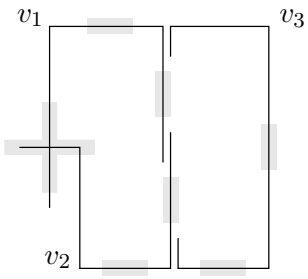

Figure 4.17: A 3-bend representation of a triangle $\{v_1, v_2, v_3\}$, which satisfies the invariant.

In the figure the displayed parts from invariant (II) for the new facial triangles $\{u, v, v_i\}$, $\{u, v_i, w\}$ and $\{v_i, v, w\}$ in $\tilde{G}_i$ are highlighted in dark grey. Additionally, every path, including the new path for v_i, has a horizontal and a vertical displayed sub-segment and these can be chosen so that they do not see any entire segment of a cross.

For example, consider the case that v_i is not adjacent to u, v, or w in G_i; see Figure 4.16 **f)**. Here the new facial triangles $\{u, v_i, w\}$ and $\{v_i, v, w\}$ have a cross of an entire segment of $P(v_i)$ and a displayed sub-segment of $P(w)$ and $P(v)$, respectively. The new facial triangle $\{u, v, v_i\}$ has the displayed part of $\{u, v\}$, which existed for the triangle $\{u, v, w\}$ before. □

The next proposition shows that Theorem 4.2.2 is tight. Moreover this confirms what Biedl and Stern strongly suspected [BS09].

Lemma 4.2.3. *There is a planar graph G with $tw(G) = 3$ and $b(G) = 3$.*

Proof. The graph G is depicted in the left of Figure 4.18 and is constructed the following way. Two vertices u and v together with 50 white vertices form an induced $K_{2,50}$ subgraph, i.e., a complete bipartite graph on 2+50 vertices. Any two consecutive white vertices together with 50 black vertices form another induced $K_{2,50}$. Finally, between any two consecutive black vertices we place a copy of the graph H that is the 29-vertex graph depicted in the right of Figure 4.18.

It is known [AS09] that $b(K_{2,n}) = 2$ if $n \geq 5$. Hence $b(G) \geq 2$. For the sake of contradiction let us assume that $b(G) = 2$ and consider a 2-bend representation of G together with the induced 2-bend representations of all induced $K_{2,50}$ subgraphs. Furthermore, it is known [BS09] that in any EPG representation of $K_{2,n}$ at most 12 vertices of the n-bipartition class establish their two edges with the same or consecutive segments. Moreover the 2-bend paths of u and v together have at most 12 endpoints of segments. Thus at most another 12 white vertices contain an endpoint

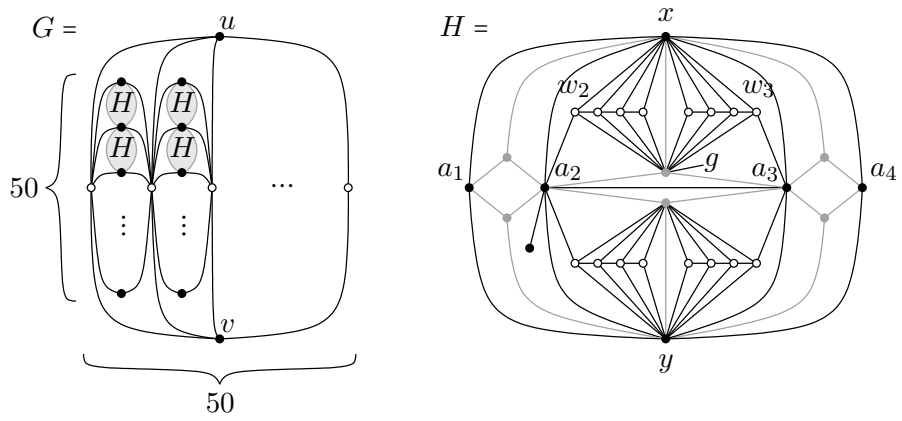

Figure 4.18: A planar graph G with $tw(G) = 3$ and $b(G) = 3$, as well as the graph H used inside G.

of u or v in the interior of a segment. Hence there are two consecutive white vertices, say u' and v' that establish their edges with u and v with two distinct segments of the same direction, each completely contained in the corresponding segment of u and v.

Now the same argument can be applied to the $K_{2,50}$ subgraph induced by u', v', and 50 black vertices. Thus there are two consecutive black vertices, which we denote by x and y, such that the following holds: There is a segment $s_{u'}$ of u' and a segment $s_{v'}$ of v', both of the same direction. Either of x and y has two segments, one completely contained in $s_{u'}$ and the other completely contained in $s_{v'}$. In particular, the paths of u', v', x, and y are positioned as the black bold paths in Figure 4.19.

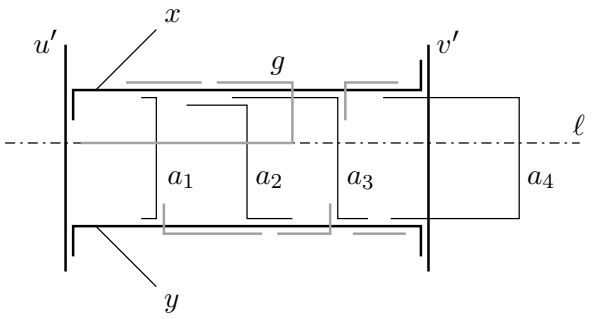

Figure 4.19: A detail of a hypothetical 2-bend representation of G.

Now consider the copy of H that is placed between x and y. W.l.o.g. we assume that the displayed segment of x and y is horizontal. The graph H contains, aside from x and y, a set A of four black vertices denoted by a_1, a_2, a_3, and a_4, and a set B of six grey vertices; see the right of Figure 4.18. Every vertex in $A \cup B$ is adjacent

to x or y, but not to u' and not to v'. Hence each path of such vertex overlaps with a horizontal segment of x or y. From the adjacencies between vertices in $A \cup B$ it follows that the vertical segments of a_1, a_2, a_3, and a_4 appear in order – w.l.o.g. from left to right. Moreover the paths of a_2 and a_3 do not overlap vertically since a_2 has a private neighbor. W.l.o.g. assume that the paths of a_2 and a_3 overlap within the horizontal x-segment; see Figure 4.19. Let us denote the grey vertex adjacent to a_2, a_3, and x by g.

It follows that the horizontal segment of g within the x-segment is completely covered by a_2 and a_3 and hence lies strictly between a_1 and a_4. In consequence, the six white vertices that are adjacent to x and g, but not to a_2 or a_3, intersect g on its *second horizontal* segment. Let us denote the set of eight white vertices, which are common neighbors of g and x by W and the grid line supporting at least six edges between g and W by ℓ. Evidently, the segment of g on ℓ crosses only one of $\{a_1, a_4\}$, but not both, say it does *not* cross a_4; see Figure 4.19.

The vertex in W that is adjacent to a_2, respectively a_3, is denoted by w_2, respectively w_3. One can check that for $i = 1, 2$ the edge between w_i and its white neighbor is established on ℓ and moreover that w_2 lies to the left of w_3 on ℓ. This implies that the white vertices in the triangle $\{x, g, w_3\}$ intersect the horizontal x-segment to the right of a_4. In order to establish their adjacency with g all these three paths contain the part of ℓ between a_3 and a_4, making them pairwise adjacent, which in G they are not – a contradiction. □

The next theorem show that another large class of planar graphs are 3-bend graphs, namely subgraphs of 4-connected maximally planar graphs, i.e., the graphs admitting a planar embedding without separating triangles. Note that this includes all 4-connected planar graphs, which is a class disjoint to the class of planar graphs from Theorem 4.2.2 (as long as there are at least five vertices).

Theorem 4.2.4. *For every subgraph G of a 4-connected maximally planar graph we have $b(G) \leq 3$.*

Proof. If G is a *proper* subgraph of a 4-connected maximally planar graph, it has a representation which is a part of a □-representation (Theorem 1.3.2). The boundary of each rectangle can be seen as a closed path with three bends starting and ending at the same corner of the rectangle. This immediately gives a simple 3-bend representation of G.

If G is a 4-connected maximally planar graph itself, take a □-representation of $G \setminus (v, w)$. It looks as depicted in Figure 4.20 (c.f. Definition 1.3.1). We replace the outer rectangles by the black paths (as shown in the figure) and each inner rectangle again by its boundary path. □

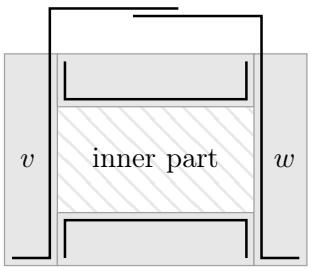

Figure 4.20: A simple 3-bend representation of a 4-connected maximally planar graph: This shows how to represent the missing edge (v, w).

Remark 4.2.5. Recall that an EPG-representation is called simple, if no three segments do pairwise overlap. It is easy to see that Theorem 4.2.4 is best possible in the sense that not every planar graph has a *simple* 2-bend representation. For instance, in every maximally planar graph $G = (V, E)$ we have $|E| = 3|V| - 6$. Hence by the Lower-Bound-Lemma every *simple* 2-bend representation of G gets by with at most six grid lines, which is obviously not possible when the graph is large. Indeed, one can argue that no maximally planar graph with nine or more vertices admits a simple 2-bend representation. On the other hand, Figure 4.21 shows a planar 3-tree on eight vertices with a simple 2-bend representation. Moreover, Euler's Formula together with the Lower-Bound-Lemma immediately shows that no maximally graph embedded on a higher-genus surface has a simple 2-bend representation.

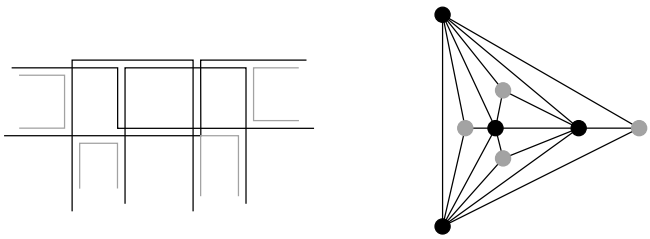

Figure 4.21: A maximally planar graph on eight vertices with a simple 2-bend representation.

Next we show that the bend-number of every planar graph is at most 4, improving a recent result of Biedl and Stern [BS09].

Theorem 4.2.6. *For every planar graph G we have $b(G) \leq 4$.*

Proof. First, we add vertices to the given graph until we obtain a super-graph G that is maximally planar. If we find a 4-bend representation for G, removing the paths from it that correspond to artificial vertices, yields a 4-bend representation of the original graph.

We build up G w.r.t. its separation-tree $\mathcal{T}_G$, as defined in Subsection 1.1.5, i.e., for each i we construct a 4-bend representation Γ_i for the level-i subgraph $G[i]$ of G. As in the proof of Theorem 2.2.4 in Chapter 2, we define the subgraph $G[-1]$ of G to be just the outer triangle $\Delta_0 = \{u_0, v_0, w_0\}$. We then start with the representation Γ_{-1} of $G[-1]$ in Figure 4.22 **c)**. Having a 4-bend representation Γ_i of $G[i]$ for $i \geq -1$, we incorporate one after another the graph G_Δ into Γ_i for every triangle Δ with depth $i+1$ in $\mathcal{T}_G$. As soon as each graph G_Δ is incorporated, the result is a 4-bend representation Γ_{i+1} of $G[i+1]$. We maintain the following invariant for Γ_i.

(I) Every inner facial triangle $\{u, v, w\}$ in $G[i]$ has two vertices, say u and v, such that

  a) there is one displayed sub-segment of $P(u)$ and one of $P(v)$ lying orthogonal to each other,

  b) there is a displayed part consisting of a sub-segment of $P(u)$ overlapping a sub-segment of $P(v)$ that lies on same the grid line as the displayed sub-segment of $P(u)$, and

  c) there is a displayed sub-segment of $P(w)$ that sees the sub-segment of $P(u)$ (type A) or the sub-segment of $P(v)$ (type B), such that the displayed sub-segment that is orthogonal to $P(w)$ lies between the two.

The invariant is illustrated in Figure 4.22 **a)** and **b)**. Figure 4.22 **c)** shows that the representation Γ_{-1} of $G[-1]$ satisfies the invariant of type B. If the invariant holds for a triangle $\Delta = \{u, v, w\}$ in $G[i]$, we say that Δ is *good*. Hence Δ_0 is good in Γ_{-1}.

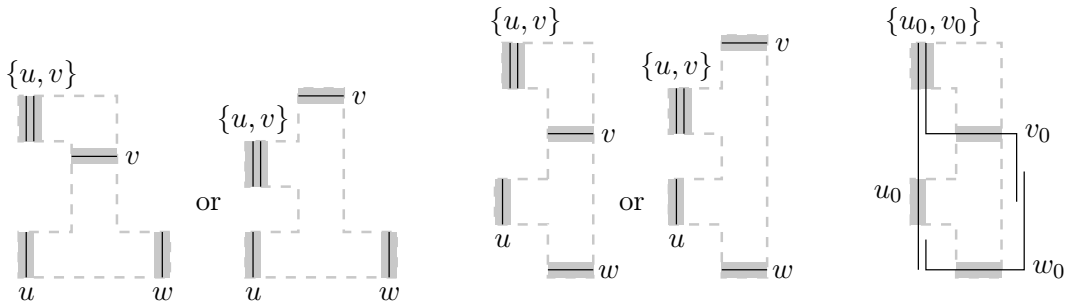

Figure 4.22: **a)** The invariant of type A for a triangle $\{u, v, w\}$ in $G[i]$. **b)** The invariant of type B for a triangle $\{u, v, w\}$ in $G[i]$. **c)** The 4-bend representation of $G[-1] = \{u_0, v_0, w_0\}$ and its invariant of type B.

Let $i \geq -1$, Γ_i be a 4-bend representation of $G[i]$ that satisfies the above invariant, and Δ be a triangle of G of depth $i+1$ in $\mathcal{T}_G$. We now show how to incorporate a 4-bend representation of G_Δ into Γ_i in such a way that the resulting representation

again satisfies the invariant. In doing so, we do not change paths that are already represented in Γ_i. Moreover, every new path will lie completely within the region of the good triangle Δ that is indicated by dotted lines in Figure 4.22 **a)** and **b)**.

We summarize the crucial claim.

Claim 1. Let Γ_i be a 4-bend representation of $G[i]$ and Δ be a triangle of G of depth $i+1$ in $\mathcal{T}_G$. Assume Δ is good, i.e., there is a region of type A or B associated with Δ as illustrated in Figure 4.22 **a)** or **b)**. Then Γ_i can be extended to a 4-bend representation of $G[i] \cup G_\Delta$, such that:

(a) every inner facial triangle in G_Δ is good,
(b) the corresponding regions are pairwise disjoint, and
(c) every such region is contained in the region corresponding to Δ.

Proof of Claim 1. Denote the vertices of Δ by u, v, and w, such that $P(u)$ and $P(v)$ have the common displayed part in the invariant. Define $G'_\Delta = G_\Delta \setminus \{u, w\}$ if the invariant for Δ is of type A, and $G'_\Delta = G_\Delta \setminus \{v, w\}$ if the invariant for Δ is of type B. The graph G'_Δ has a $\square$-representation $\Gamma_\mathcal{R}$ due to Theorem 1.3.2. We apply two transformations to $\Gamma_\mathcal{R}$ as illustrated in the top row and the bottom row in Figure 4.23, respectively. The first transformation is done for every *inner* rectangle $\mathcal{R}$. Afterwards, every such $\mathcal{R}$ is marked with a vertical *marking segment* suspended between the rightmost rectangle on top of $\mathcal{R}$ and the leftmost rectangle at the bottom of $\mathcal{R}$. The second transformation is done for every pair of a rectangle $\mathcal{R}$ and its rightmost neighboring rectangle $\mathcal{R}'$ on top of $\mathcal{R}$. Afterwards the marking segment of $\mathcal{R}$ lies *to the right* of the marking segment of $\mathcal{R}'$.

We denote the resulting $\square$-representation of G'_Δ again by $\Gamma_\mathcal{R}$. We position $\Gamma_\mathcal{R}$ on top of the region for Δ such that each of the following holds. Let z denote the fourth outer vertex in G'_Δ, x denote the vertex in G'_Δ that is adjacent to u and v, and y denote the vertex adjacent to v and w in type A, and to u and w in type B. Note that two or all of the vertices $\{x, y, z\}$ may coincide.

(i) The displayed sub-segment of $P(u)$, $P(v)$, and $P(w)$ runs along the inner side of the rectangle $\mathcal{R}(u)$, $\mathcal{R}(v)$, and $\mathcal{R}(w)$, respectively.
(ii) The rectangle $\mathcal{R}(z)$ lies completely inside the region for Δ.
(iii) The marking segment of $\mathcal{R}(x)$ and $\mathcal{R}(y)$ lies completely inside the region for Δ.
(iv) The displayed sub-segments of $P(u)$, $P(v)$ see every rectangle in $\Gamma_\mathcal{R}$.

It is not difficult to see that the $\square$-representation $\Gamma_\mathcal{R}$ can be stretched along some cuts, such that each of the above holds, and the order of marking segments is still maintained. See Figure 4.24 for an illustration.

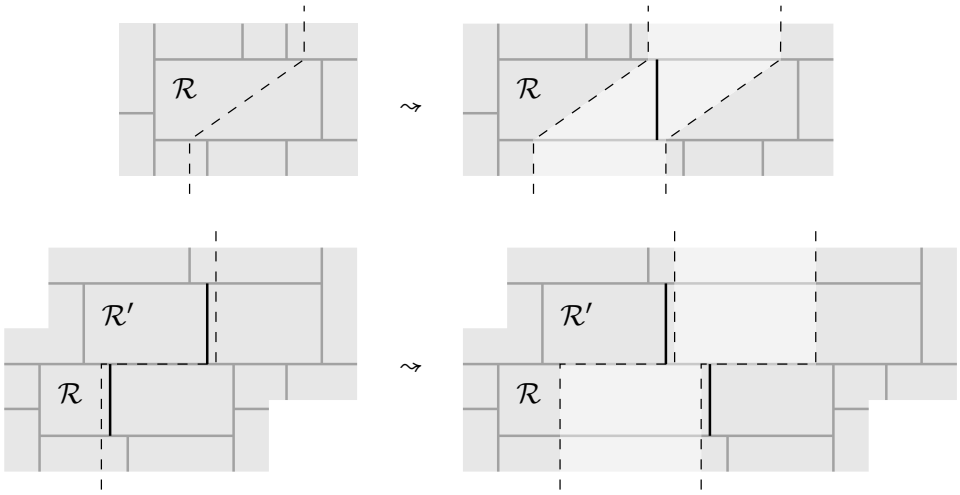

Figure 4.23: **Top row:** How to transform the rectangle representation at an inner rectangle $\mathcal{R}$ and mark $\mathcal{R}$ with a vertical segment. **Bottom row:** How to transform the rectangle representation at a pair $\mathcal{R}$, $\mathcal{R}'$, with $\mathcal{R}'$ being the rightmost rectangle on top of $\mathcal{R}$.

Now we introduce a snake-like 4-bend path for every vertex $i \neq u, v, w$ in G'_Δ within its corresponding rectangle in $\Gamma_\mathcal{R}$. The path $P(i)$ starts at the bottom-left corner of $\mathcal{R}(i)$, goes up to the top-left corner, where it bends and goes right up to the segment that marked $\mathcal{R}(i)$. Then $\mathcal{P}(i)$ follows along the marking segment down to the bottom side of the rectangle, where it bends and goes to the right up to the bottom-right corner and then up to the top-right corner. Afterwards, every path $P(i)$ is shortened by just a little bit, such that no edge-intersection is lost. See Figure 4.25 for an illustration.

It is not difficult to check that these snake-like paths together with the already constructed paths $P(u)$, $P(v)$, and $P(w)$, give a 4-bend representation Γ_{EPG} of G'_Δ. Note that there is an edge-intersection between $P(i)$ and $P(j)$ in case $\mathcal{R}(j)$ is the rightmost rectangle on top of $\mathcal{R}(i)$ in $\Gamma_\mathcal{R}$, because the marking segment of $\mathcal{R}(i)$ lies to the right of the marking segment of $\mathcal{R}(j)$. Note further, that for every edge incident to v there is a corresponding edge-intersection because the marking segments of $\mathcal{R}(x)$ and $\mathcal{R}(y)$ lie within the region of Δ.

We claim that every inner facial triangle $\{u', v', w'\}$ in G'_Δ is good, i.e., we find a certifying region for $\{u', v', w'\}$ similar to the one from Figure 4.22 **a)** or **b)**. We distinguish two cases. First, assume that at most one of $\{u', v', w'\}$ is a vertex of Δ. Then the point $p_{u'v'w'}$ in the $\square$-representation $\Gamma_\mathcal{R}$, that is the common intersection of $\mathcal{R}(u')$, $\mathcal{R}(v')$ and $\mathcal{R}(w')$ lies within the region for Δ. It is easy to see that exactly one of the three rectangles has its bottom-left or top-right corner at $p_{u'v'w'}$. Hence,

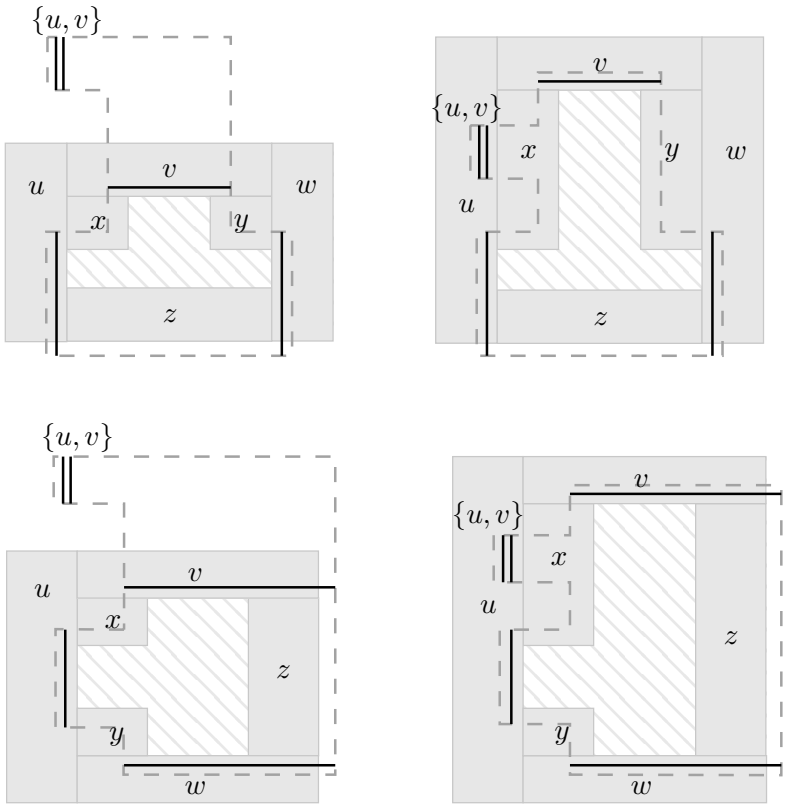

Figure 4.24: The placement of the rectangle-representation $\Gamma_{\mathcal{R}}$ of G'_Δ on top of the region for $\Delta = \{u, v, w\}$ depending on the region's type. **Top row:** Invariant for Δ is of type A. **Bottom row:** Invariant for Δ is of type B.

the 4-bend representation Γ_{EPG} locally looks like in one of the cases in Figure 4.26. In the figure, for each case a region for $\{u', v', w'\}$ is highlighted, which shows that this facial triangle is good.

It remains to identify a region for the three facial triangles $\{u, v, x\}$, $\{v, w, y\}$, and $\{u, w, z\}$ if Δ is of type A, as well as for the three facial triangles $\{u, v, x\}$, $\{u, w, y\}$, and $\{v, w, z\}$ if Δ is of type B. We slightly modify the paths $P(x)$ and $P(y)$ as follows. We move the top-end of the middle segment of $P(x)$ (the one that runs along the marking segment of $\mathcal{R}(x)$) vertically onto a horizontal grid line that crosses the displayed part for $\{u, v\}$. Then $P(x)$ bends there, goes up to this displayed segment, bends another time and ends shortly after this. See Figure 4.27 for an illustration. Note that in the left and the right of the figure the top-end was raised and lowered, respectively.

The path $P(y)$ is altered as follows. If the invariant for Δ is of type A, the right end of $P(y)$ is shortened, such that it lies completely inside the region for Δ. See the

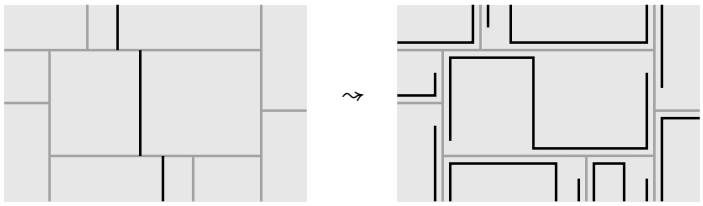

Figure 4.25: How to represent every vertex in $G'_\Delta \setminus \Delta$ by a 4-bend snake-like path.

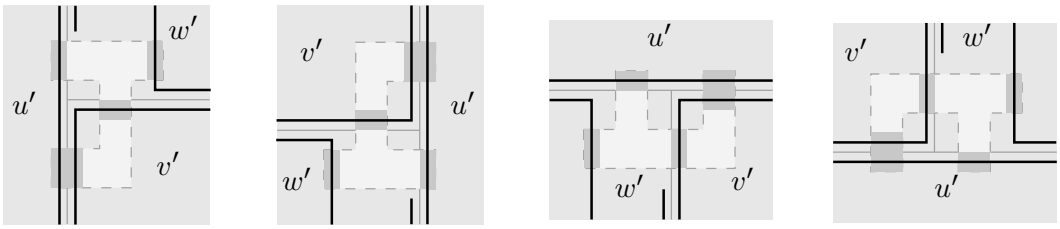

Figure 4.26: The four possibilities for an inner face $\{u', v', w'\}$ of G_Δ with at most one vertex from Δ. The highlighted region shows that such a face is good.

left of Figure 4.27 for an example. Similarly, if the invariant for Δ is of type B, the *left* end of $P(y)$ is shortened, such that it lies completely inside the region for Δ. An example is on the right of Figure 4.27.

Note that the modification of $P(x)$ and $P(y)$ results again in a 4-bend representation of the graph G'_Δ, which we again denote by Γ_{EPG}. Moreover, we can now find regions for the last three facial triangles in G_Δ, proving that each of these is good. Those regions are highlighted in Figure 4.27. Let us remark, that Figure 4.27 shows the situation only for one version of a type-A region and one version of a type-B region. However, it is not difficult to see, that the missing versions are a mixture of these two. Moreover, in case some two or more of the vertices x, y, and z coincide, the corresponding regions do still exist and are disjoint. One example is illustrated in Figure 4.28.

We have found a 4-bend path for every inner vertex of G_Δ, and a suitable region for every facial triangle of G_Δ, and hence have proven the claim. △

We can now extend the 4-bend representation Γ_i of $G[i]$ by the 4-bend representations Γ_{EPG} of the graphs G_Δ with Δ at depth $i+1$ in $\mathcal{T}_G$ one after another as provided by the above claim. This results in a 4-bend representation Γ_{i+1} of $G[i+1]$, such that every inner face of $G[i+1]$ is good. Thus, repeating this procedure until $G[i] = G$ we obtain a 4-bend representation of G. □

Putting Theorem 4.2.6 and Lemma 4.2.3 together we have shown the following.

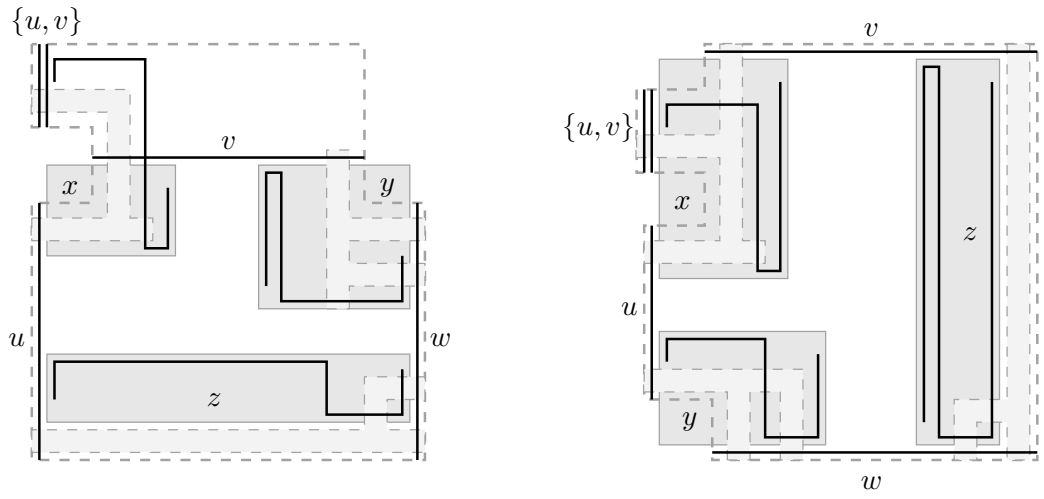

Δ of type A Δ of type B

Figure 4.27: The rectangle-representation Γ and the 4-bend representation Γ′ is positioned into the region of the good triangle $\Delta = \{u, v, w\}$. The regions of the three facial triangles that share an edge with Δ are highlighted in dark grey.

Theorem 4.2.7. *In an EPG representation of a planar graph* 4*-bend paths are always sufficient and* 3*-bend paths are sometimes necessary.*

Although we improved the previously known lower and upper bounds, it remains open to determine the maximum bend-number of planar graphs. We suspect that 4 is the right answer.

Conjecture 4.2.8. There is a planar graph G, such that every EPG representation of G contains at least one path with four bends.

4.3 The Bend-Number of Graphs with Fixed Degeneracy, Tree-Width, or Maximum Degree

In this section, we consider the class of graphs with fixed degeneracy (Subsection 4.3.1), fixed tree-width (Subsection 4.3.2) and fixed maximum degree (Subsection 4.3.3). We present lower and upper bounds on the maximum bend-number within these graph classes. The bounds are matching in case of the degeneracy, differ by an additive 1 in case of tree-width, and differ by a factor of 2 in case of maximum degree.

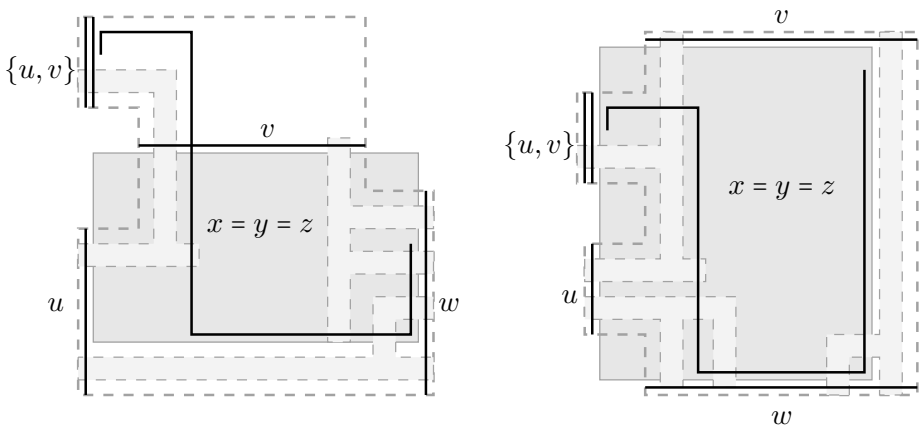

Δ of type A Δ of type B

Figure 4.28: The case $x = y = z$ for one version of a type-A region and one version of a type-B region.

4.3.1 The Bend-Number in Terms of the Degeneracy

Recall from Definition 1.1.1, that the degeneracy $d(G)$ is the smallest number k, such that there is a vertex ordering $(v_1, \ldots, v_n)$ of G with $\deg_{G_i}(v_i) \leq k$ for every $i = 1, \ldots, n$. First we provide an upper bound on the bend-number of graphs with a fixed degeneracy. The following result was already suspected to be true [BS09]. The representation we construct requires $d(G) \geq 2$. However, for in case $d(G) = 1$, i.e., G is a forest, it is already known [GLS07] that $b(G) \leq 1$.

Theorem 4.3.1. *For every graph G with $d(G) \leq k$ we have $b(G) \leq 2k - 1$.*

Proof. If $k = 1$, the result follows from Golumbic *et al.* [GLS07].

For $d(G) = k \geq 2$, take a vertex ordering of G such that every vertex has at most k neighbors with a smaller label. We consider the building sequence $G_0 \subset \cdots \subset G_n$ of G w.r.t. this ordering, i.e., where every new vertex v_i is connected to at most k vertices that are already represented by paths. We construct a $(2k-1)$-bend representation simultaneously to the building process of G. We maintain that for $i \geq 1$ Γ_i is a $(2k-1)$-bend representation of G_i, such that:

(I) For every vertex v in G_i there is a displayed vertical sub-segment of $P(v)$.
(II) All these vertical sub-segments pairwise see each other.
(III) For every vertex v in G_i there is a displayed horizontal sub-segment of $P(v)$.

Clearly, such a representation Γ_1 for the 1-vertex graph G_1 exists. For $i \geq 2$, the path for v_i is introduced into Γ_{i-1} as follows. All but one edge-intersection of $P(v_i)$

are placed vertically onto a small portion of the corresponding vertical displayed sub-segment. These short vertical segments of $P(v_i)$ are connected by horizontal segments without introducing further edge-intersections. The last edge-intersection is placed horizontally onto the corresponding horizontal displayed sub-segment. This short horizontal segment of $P(v_i)$ is connected to the rest by a horizontal and a vertical segment. The vertical edge-intersections can easily be chosen such that the vertical segment of $P(v_i)$ can see all the other vertical displayed sub-segments in Γ_{i-1}. See Figure 4.29 for an illustration.

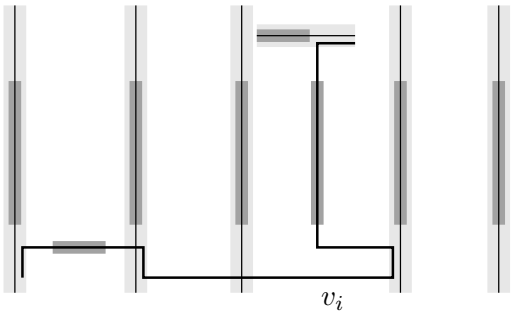

Figure 4.29: Building a $(2k-1)$-bend representation of G, the path for vertex v_i is inserted. Displayed sub-segments before and after the insertion are highlighted with light-grey and dark-grey, respectively.

It is easy to verify that $P(v_i)$ has exactly $2 \cdot \deg_{G_i}(v_i) - 1$ bends, which proves the theorem. (If $\deg_{G_i}(v_i) = 0$, we take a suitable 1-bend path as $P(v_i)$.) □

Theorem 4.3.1 is worst-case optimal, even for bipartite graphs.

Lemma 4.3.2. *For every k there is a bipartite graph G with $d(G) = k$ and $b(G) = 2k-1$.*

Proof. For $k = 1$ note that caterpillars are the only trees that are interval graphs and hence have bend-number 0. Thus there are trees with $d(G) = 1$ and $b(G) = 1$.

For $k \geq 2$, we construct a graph G, which is huge and arises from a $K_{k,n}$ with very large n. Let B be the independent set of vertices in $K_{k,n}$ of size n. We connect $k(2k-2)+1$ new vertices with every k-subset of B. The set of the $(k(2k-2)+1)\binom{n}{k}$ added vertices is denoted by C. Moreover, we connect $k(2k-2)+1$ new vertices with every k-subset of C. This set of vertices is denoted by D.

The above building sequence gives a vertex ordering with $\deg_{G_i}(v_i) \leq k$, for every vertex v_i. Hence we have $d(G) \leq k$. For the sake of contradiction assume that $b(G) \leq 2k-2$ and consider the induced $K_{k,n}$ the construction started with. By Lemma 4.1.6

all, but roughly $4k^4$ vertices in B, are represented by paths in which exactly every second segment is establishing one edge-intersection. By this –with increasing n– an arbitrarily large subset $B' \subset B$ looks like Figure 4.9 depicts, i.e., vertices in B' establish all further incidences only with vertical segments. Additionally, there are many pairs of k-subsets of B' such that every path in one subset lies to the left of every path in the other subset.

We fix k distinct k-subsets $B_1, \ldots, B_k$ of B', such that every path in B_1 lies to the left of every path in $B_2 \cup \ldots \cup B_k$. Hence every vertex in C that is connected to B_1 lies completely to the left of every vertex in C that is connected to one of $B_2, \ldots, B_k$.

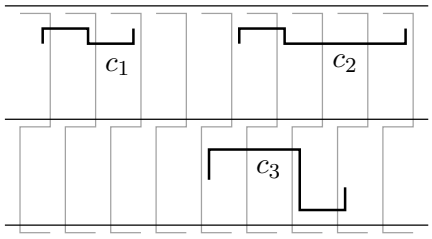

Figure 4.30: A part of a hypothetical $(2k-2)$-bend representation of G. The sets B' and $C' = \{c_1, c_2, c_3\}$ are depicted grey and thick, respectively.

We connected $k(2k-2)+1$ vertices with every B_i, that is, one more than there are bends in B_i. Hence for every $i \in \{1, \ldots, k\}$ there is at least one vertex $c_i \in C$ that is connected to B_i and whose vertical segments are contained in segments of B_i. See Figure 4.30 for an example. In G a set $D' \subset D$ of $k(2k-2)+1$ vertices is connected to $C' := \{c_1, \ldots, c_k\}$; again one more than there are bends in C'. Hence there is at least one vertex in D' for which all its horizontal segments are contained in segments of C'. But this is impossible since the path $c_1 \in C'$ lies left of all other paths in C'. □

Putting Theorem 4.3.1 and Lemma 4.3.2 together we have shown the following.

Theorem 4.3.3. *For every $k \geq 1$, $(2k-1)$-bend paths are always sufficient and sometimes necessary for an EPG representation of a graph with degeneracy k.*

4.3.2 The Bend-Number in Terms of the Tree-Width

Golumbic, Lipshteyn, and Stern [GLS07] showed that every tree, i.e., every graph with tree-width 1, has a single-bend representation. Since graphs with bend-number 0 are interval graphs, and not every tree is an interval graph, this implies that the maximum bend-number among all trees is 1. In this section, we consider the maximum bend-number of all graphs of tree-width $k \geq 2$. First we provide general lower bound, which immediately follows from Theorem 4.1.7 in Section 4.1.

Lemma 4.3.4. *For every $k \geq 2$ there is a graph G with $tw(G) = k$ and $b(G) = 2k-2$.*

Proof. Theorem 4.1.7 states, that the complete bipartite graph $K_{m,n}$ with $n > m^4 - 2m^3 + 5m^2 - 4m$ has bend-number $2m-2$. Clearly, this graph has tree-width m. □

The next lemma gives an upper bound for graphs with tree-width 2, that matches the lower bound from Lemma 4.3.4. Note that the upcoming statement seems stronger than Theorem 4.2.1 since every outer-planar graph has tree-width at most 2. However, Theorem 4.2.1 gives a *simple* 2-bend representation, which we do not obtain from the lemma below.

Lemma 4.3.5. *For every graph G with $tw(G) = 2$ we have $b(G) \leq 2$.*

Proof. Let $\tilde{G}$ be the 2-tree which contains G, and $(v_1, \ldots, v_n)$ be a vertex ordering of G, such that $\deg_{G_i}(v_i) \leq 2$ for every $i = 1, \ldots, n$. We construct a 2-bend representation of G along the building sequence $G_1 \subset \cdots \subset G_n = G$, where we add vertex v_i to G_i, such that the two neighbors of v_i in $\tilde{G}_i$ form a 2-clique in $\tilde{G}_i$ (and not necessarily in G_i). We maintain that Γ_i is a 2-bend representation of G_i, such that every 2-clique in $\tilde{G}_i$ satisfies one of the two invariants in Figure 4.31 **a)**, i.e., for $i = 2, \ldots, n$ and $(u, v) \in E(\tilde{G}_i)$ we have

(I) displayed sub-segments of $P(u)$ and $P(v)$ in Γ_i, which see each other as depicted in the top row of Figure 4.31 **a)**, or
(II) a displayed part in $P(u) \cup P(v)$ as depicted in the bottom row of Figure 4.31 **a)**.
(III) Moreover, all displayed parts are pairwise disjoint.

Now for $i \geq 2$, let Γ_i be a 2-bend representation of G_i that satisfies our invariant. It is easy to set up Γ_2 and we leave it to reader. Let (u, v) be the 2-clique that is the neighborhood of v_i in $\tilde{G}_i$. If v_i has an edge with both, u and v, we introduce the path for v_i as illustrated in Figure 4.31 **b)** depending on the type of the displayed part for (u, v). The displayed parts for the new 2-cliques (u, v_i) and (v, v_i) in $\tilde{G}_{i+1}$ are highlighted in dark grey, and the displayed part for (u, v) is highlighted in light grey.

If v_i has an edge with u and no edge with v, we introduce the path for v_i as illustrated in Figure 4.31 **c)** depending on the type of the displayed part for (u, v). Again the displayed parts for the new 2-cliques (u, v_i) and (v, v_i), and the old 2-clique (u, v) in $\tilde{G}_{i+1}$ are highlighted in dark grey, and light grey, respectively. If v_i has an edge with v and no edge with u, the roles of u and v are simply exchanged. In case, v_i has an edge neither with u nor with v, we introduce the path for v_i as illustrated in Figure 4.31 **d)** depending on the type of the displayed for (u, v). □

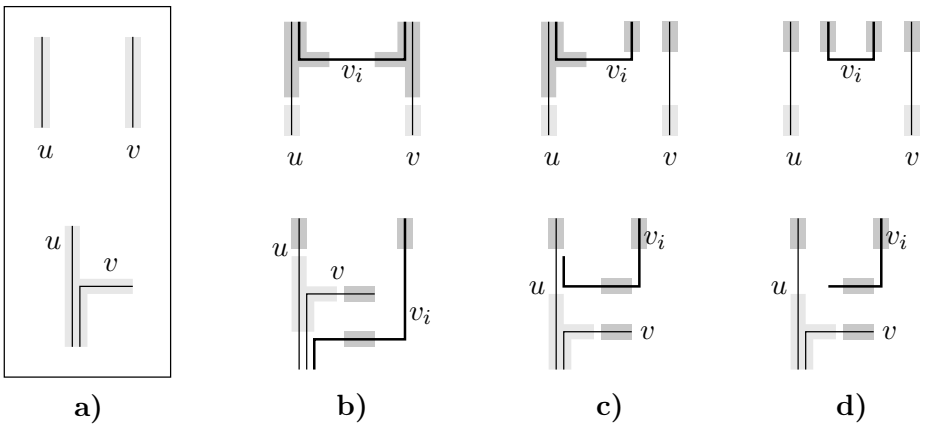

Figure 4.31: Invariants for a 2-clique and insertion rules for the path for the new vertex v_i (drawn bold). **a)** The two types of the invariant. **b)** Vertex v_i has an edge with both, u and v. **c)**: Vertex v_i has an edge with u and no edge with v. **d)**: Vertex v_i has no edge with u or v.

We could not generalize Lemma 4.3.5 to tree-width $k \geq 3$. However, since the tree-width of a graph is at most its degeneracy, we obtain an upper bound by Theorem 4.3.1 that is away from the lower bound only by an additive 1.

Corollary 4.3.6. *For every graph G with $tw(G) = k \geq 3$ we have $b(G) \leq 2k - 1$.*

It remains open whether the maximum bend-number among all graphs with tree-width $k \geq 3$ is $2k - 2$ or $2k - 1$. We suspect that the right answer is $2k - 1$ here; at least for some large k.

Conjecture 4.3.7. There exist a $k \geq 3$ and a graph G with $tw(G) = k$ and $b(G) = 2k - 1$.

4.3.3 The Bend-Number in Terms of the Maximum Degree

It was known [BS09] that every graph G with maximum degree Δ admits a $(2\lceil\frac{\Delta+1}{2}\rceil + 1)$-bend representation, and that in case G is bipartite, $2\lceil\frac{\Delta}{2}\rceil$ bends suffice. We improve this with a construction that is much simpler.

Theorem 4.3.8. *For every graph G with maximum degree Δ we have $b(G) \leq \Delta$. If G is bipartite, then $b(G) \leq \Delta - 1$.*

Proof. We actually prove that $b(G)$ is at most the edge-chromatic number $\chi'(G)$ of G minus one. Then Vizing's Theorem [Viz64], which states $\chi'(G) \leq \Delta + 1$, yields

$b(G) \leq \Delta$ for G with maximum degree Δ. Moreover, for bipartite graphs $\chi'(G) = \Delta(G)$.

Consider a proper edge-coloring of G with $k = \chi'(G)$ colors. We use an arbitrary, not self-intersecting $(k-1)$-bend path P as a template. (In Figure 4.32 we have chosen a snake-like path P.) Every vertex in G runs within a small distance of P. Along the i-th segment, two vertices run through the same grid line if and only if they are joint by an edge of color i. For example, in Figure 4.32 the edge (u, v) has color 1. □

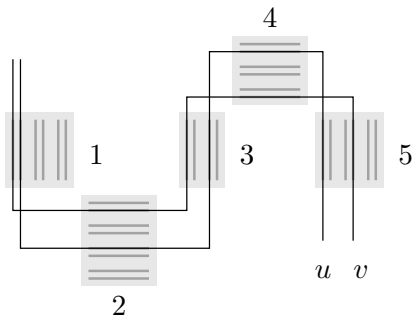

Figure 4.32: A $(k-1)$-bend representation based on a proper edge-coloring with k colors: The grey blocks correspond to the colors. Every edge of color i is assigned a grid line within the i-th block. Paths are inserted as demonstrated by u and v according to their incident edges.

The best known lower bound in terms of the maximum degree comes from the complete bipartite graph $K_{m,m}$. By Theorem 4.1.3, $b(K_{m,m}) = \lceil \frac{m}{2} \rceil$ and clearly $\Delta(K_{m,m}) = m$. Hence we conclude the following, which alternatively can be proven by applying the Lower-Bound-Lemma to *any* Δ-regular triangle-free graph.

Lemma 4.3.9. *For every $\Delta \geq 2$ there is a graph G with maximum degree Δ and $b(G) \geq \lceil \frac{\Delta}{2} \rceil$.*

Putting Theorem 4.3.8 and Lemma 4.3.9 together we have shown the following.

Theorem 4.3.10. *For every $\Delta \geq 2$, Δ-bend paths are always sufficient and $\lceil \frac{\Delta}{2} \rceil$-bend paths are sometimes necessary for an EPG representation of a graph with maximum degree Δ.*

4.4 Recognizing Single-Bend Graphs is NP-Complete

It has been asked several times [GLS07, BS09, Rie09] for the complexity of recognizing k-bend graphs. In general, the bend-number of a graph can be computed by solving a

mixed integer program (MIP). Unfortunately, the problem instance becomes so huge, that this approach is inapplicable even for graphs with only ten vertices. It is well-known that interval graphs, which are exactly 0-bend graphs, can be recognized in polynomial time [BL76]. In this section we prove that recognizing single-bend graphs (1-bend graphs) is NP-complete.

Definition 4.4.1 (SINGLE-BEND-RECOGNITION)**.**

INSTANCE:	Undirected graph G.
QUESTION:	Is there a single-bend representation of G, i.e., is $b(G) \leq 1$?

SINGLE-BEND-RECOGNITION is in NP, since a single-bend representation can be easily verified. For NP-hardness we set up a reduction from ONE-IN-THREE 3-SAT, i.e., we are given a formula $\mathcal{F} = (\mathcal{C}_1 \wedge \cdots \wedge \mathcal{C}_n)$ that is a conjunction of clauses $\mathcal{C}_1, \ldots, \mathcal{C}_n$. Each clause is the *exclusive disjunction* of exactly three literals $\mathcal{C}_i = (x_{i1} \veebar x_{i2} \veebar x_{i3})$, which are in turn either negated or non-negated boolean variables. Given such a formula $\mathcal{F}$, it is NP-complete [GJ79, Sch78] to decide, whether there is an assignment of the variables fulfilling $\mathcal{F}$, so that is in each clause there is *exactly one true literal*. Moreover ONE-IN-THREE 3-SAT remains NP-complete if each literal is a *non-negated variable* and each clause consists of three *distinct literals*. We use both additional assumptions on $\mathcal{F}$, even though the first is just for convenience. The distinctness assumption is crucial in the following reduction.

Given a ONE-IN-THREE 3-SAT formula $\mathcal{F}$ we define a graph $G_\mathcal{F}$, such that $b(G_\mathcal{F}) = 1$ if and only if $\mathcal{F}$ can be satisfied. The graph consists of an induced subgraph $G_\mathcal{C}$ for every clause $\mathcal{C}$ with 13 vertices, called the *clause gadget*, a vertex v_j for every variable x_j and 31 additional vertices.

4.4.1 Clause Gadgets

Constructing a clause gadget $G_\mathcal{C}$ starts with an induced *octahedral graph* O. Label the vertices by $\{a, A, b, B, c, C\}$ as in Figure 4.33. This way, $\{a, A\}$, $\{b, B\}$ and $\{c, C\}$ are the three *non*-edges and their complements $\{b, C, B, c\}$, $\{a, C, A, c\}$ and $\{a, B, A, b\}$ are the three induced 4-cycles in O.

Lemma 4.4.2. *We have $b(O) = 1$ and in every single-bend representation*

(i) there is a unique grid point, called the center*, that is contained in every path,*

(ii) every edge-intersection lies on a half ray starting at the center, called a center ray*,*

(iii) for every pair of center rays, there is a unique vertex in O supported by exactly these two center rays, and

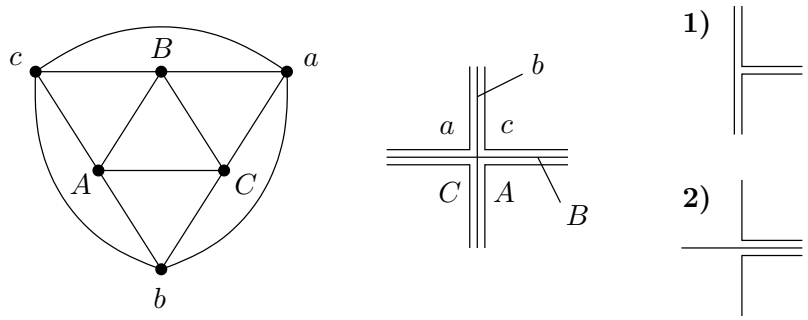

Figure 4.33: The labeled octahedral graph O, a single-bend representation of O, and the two possible ways a triangle of O is represented.

(iv) every triangle in O is represented in one of the two ways on the right of Figure 4.33.

Proof. Figure 4.33 shows $b(O) \leq 1$ and since O contains induced 4-cycles it is not an interval graph. Hence $b(O) = 1$.

It is known [GLS07] that every induced 4-cycle in a single-bend representation is either a *frame*, a *true pie*, or a *false pie*. These terms are illustrated in Figure 4.34. If an induced 4-cycle is represented by a frame, then the bends of its vertices are distinct. Thus in a single-bend representation no further path may overlap all of them. Since for each induced 4-cycle ($\{a, B, A, b\}$, $\{a, C, A, c\}$ and $\{b, C, B, c\}$) in O there is a vertex that is adjacent to all of its vertices, these 4-cycles are pies. Thus all pies share the middle point, the claimed center, and every vertex is supported by exactly two center rays. Since every edge in O is contained in an induced 4-cycle, no two vertices in O are supported by the same pair of center rays. This concludes (i)–(iii). Part (iv) is easily obtained from this. □

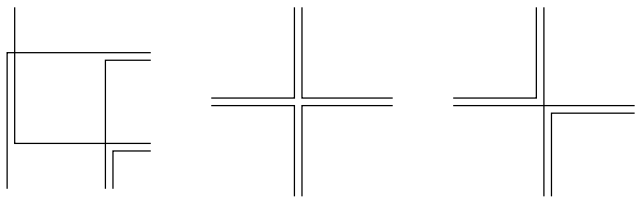

Figure 4.34: Single-bend representations of an induced 4-cycle: A frame, a true pie and a false pie.

To complete a clause gadget $G_{\mathcal{C}}$ seven vertices are added: W_{ABC} is connected to $\{A, B, C\}$, w_{abC}, w_{aBc} and w_{Abc} are connected to $\{a, b, C\}$, $\{a, B, c\}$ and $\{A, b, c\}$,

respectively, and s_{ab}, s_{ac} and s_{bc} are connected to $\{a, b\}$, $\{a, c\}$ and $\{b, c\}$, respectively. The resulting graph is depicted in Figure 4.35.

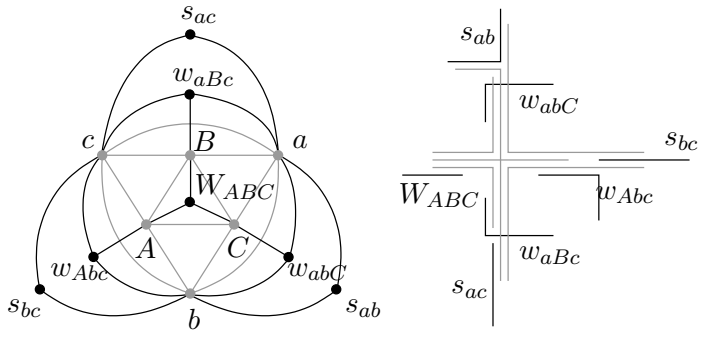

Figure 4.35: The clause gadget $G_{\mathcal{C}}$ with a single-bend representation.

Lemma 4.4.3. *We have $b(G_{\mathcal{C}}) = 1$ and in every single-bend representation*

(i) every center ray contains a segment of exactly one of W_{ABC}, w_{abC}, w_{aBc}, and w_{Abc} and

(ii) every such segment, except the one of W_{ABC}, is contained in a segment from $\{a, b, c\}$.

Proof. Let $w \in \{W_{ABC}, w_{abC}, w_{aBc}, w_{Abc}\}$. Then w is connected to the vertices of some triangle Δ in O. By Lemma 4.4.2–(iv), Δ is represented in one of the two ways that are illustrated on the right of Figure 4.33. In case **1)**, there are edge-intersections corresponding to w in two distinct center rays. But then, by Lemma 4.4.2–(iii), w edge-intersects vertices in O that are not adjacent to w. Hence Δ is represented as in case **2)** and one segment of w is contained in the center ray that carries all three vertices in Δ. This concludes part (i).

Now consider a pair (w, s) in $\{(w_{abC}, s_{ab}), (w_{aBc}, s_{ac}), (w_{Abc}, s_{bc})\}$. Both, w and s, are contained in at most one center ray. Moreover, it is the same center ray and it contains one of the capitalized vertices, that are adjacent to w but not to s. Hence the segment of s lies further away from the center than the segment of w. Thus the segment of w is contained in a segment of each neighbor of s. □

4.4.2 The Reduction

Given a formula $\mathcal{F} = (\mathcal{C}_1 \wedge \cdots \wedge \mathcal{C}_n)$ with clauses $\mathcal{C}_i = (x_{i1} \vee x_{i2} \vee x_{i3})$ for $i = 1, \ldots, n$ we are now ready to define the graph $G_{\mathcal{F}}$ as follows. See Figure 4.36 for an example.

1. For each clause $\mathcal{C}$ there is a clause gadget $G_{\mathcal{C}}$.

2. For each variable x_j there is a vertex v_j that is adjacent to w_{Abc}, w_{aBc}, or w_{abC}, whenever x_j is the first, second, or third variable in $\mathcal{C}$, respectively.
3. There is a vertex V adjacent to every W in the clause gadgets.
4. There is a $K_{2,4}$ with a specified vertex T of the larger part, called the *truth-vertex*. T is adjacent to every v_j and V.
5. There are two octahedral graphs O_1 and O_2. The vertex T is connected to the vertices of a triangle of each.
6. There are two more octahedral graphs O_3 and O_4. The vertex V is connected to the vertices of any triangle in O_3 and any triangle in O_4.

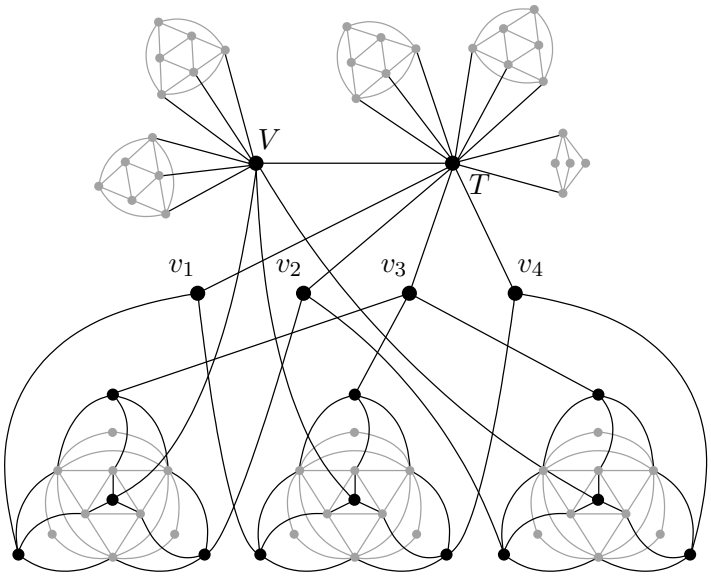

Figure 4.36: The graph $G_{\mathcal{F}}$ for $\mathcal{F} = (x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee x_3 \vee x_4) \wedge (x_2 \vee x_3 \vee x_4)$.

We prove that a ONE-IN-THREE 3-SAT-formula $\mathcal{F}$ can be satisfied if and only if $b(G_{\mathcal{F}}) = 1$.

Theorem 4.4.4. *SINGLE-BEND-RECOGNITION is NP-hard.*

Proof. First suppose $b(G_{\mathcal{F}}) = 1$ and consider a single-bend representation of $G_{\mathcal{F}}$. W.l.o.g. assume, that V and T edge-intersect with their horizontal segments. We set a variable x_j true if v_j edge-intersects the truth-vertex T with its horizontal segment and false if v_j edge-intersects T with its vertical segment.

It is known [AS09] that in every single-bend representation of a $K_{2,4}$, every vertex of the larger part, in particular T here, has its bend in a false pie. The truth-vertex T is attached to a triangle of O_1 and O_2. The proof of Lemma 4.4.3–(i) shows that a segment of T is contained in exactly one center ray of each, O_1 and O_2. As the

bend of T is in a false pie of $K_{2,4}$, the endpoints of T are contained in O_1 and O_2, respectively. Hence, every further segment that edge-intersects T is contained in a segment of T. Consequently, a vertex v_j edge-intersects the lower-case w-vertex in each clause gadget with its vertical segment if and only if x_j is true.

For the same reason V edge-intersects every neighbor other than T with its vertical segment. Since V is attached to a triangle of O_3 and O_4, the two endpoints of the vertical segment of V are contained in O_3 and O_4, respectively. Thus, the vertical segment of the upper-case W-vertex in each clause gadget is contained in the vertical segment of V. Thus by Lemma 4.4.3–(i) the horizontal segment of every such W is contained in a horizontal center ray. Hence, of the other three center rays, two are vertical and one is horizontal. Together with Lemma 4.4.3–(ii) this yields that in every clause gadget exactly two of $\{w_{abC}, w_{aBc}, w_{Abc}\}$ edge-intersect the corresponding v_j with their horizontal segment and exactly one with its vertical segment. In other words every clause contains exactly one true variable.

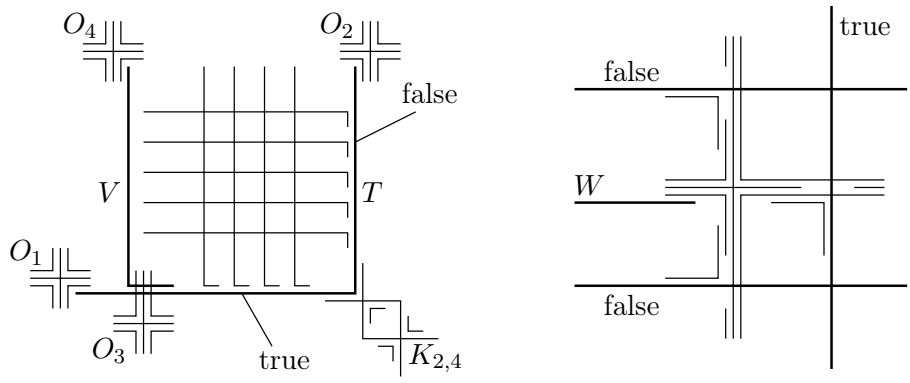

Figure 4.37: **On the left:** A single-bend representation of $G_{\mathcal{F}}$. The vertex V and the truth-vertex T are drawn bold. The clause gadgets are omitted. **On the right:** A single-bend representation of a clause gadget $G_{\mathcal{C}}$. The vertices v_j that correspond to the variables in the clause $\mathcal{C}$ and the vertex W of the clause gadget are drawn bold.

Now given a truth assignment fulfilling $\mathcal{F}$, we construct a single-bend representation of $G_{\mathcal{F}}$. First, represent all of $G_{\mathcal{F}}$ but the clause gadgets as on the left side in Figure 4.37. A vertex v_j is connected to the truth-vertex T horizontally if x_j is true and vertically if x_j is false.

To interlace a clause gadget $G_{\mathcal{C}}$, introduce a horizontal grid line l_h between the horizontal grid lines supporting the two false variables in $\mathcal{C}$. Then connect the W-vertex in $G_{\mathcal{C}}$ to V vertically with its bend on l_h. Furthermore introduce a vertical grid line l_v between the vertical grid lines supporting V and the true variable in $\mathcal{C}$. Where l_h and l_v cross, introduce the center of the clause gadget as illustrated on

the right side in Figure 4.37. Note that the clause gadget is symmetric in A, B and C and hence it can be represented with every center ray pointing into the desired direction. □

Given this NP-hardness result, a natural question follows:

Question 4.4.5. What is the complexity of recognizing k-bend graphs for $k \geq 2$?

4.5 Comparison with Interval-Number, Local Track-Number, and Track-Number

In this section, we briefly compare the bend-number with other graph parameters. Some ideas presented here are joint work with D. Heldt and K. Knauer [HKU10]. However, this is mostly ongoing research [HKU11], of which we want to discuss only selected parts here. We start with three concepts of intersection graphs, that are somehow similar to EPG graphs.

Interval graphs are intersection graphs of intervals on the real line. Every vertex is associated with an interval, in such a way that two intervals overlap with non-zero length if and only if the corresponding vertices are adjacent. This subject has been extended to intersection graphs of systems of intervals by Harary and Trotter [HT79] in 1979. In a *k-interval representation* of a graph G every vertex is associated with a set of at most k intervals on the real line, such that vertices are adjacent if and only if any of their intervals intersect with non-zero length. The *interval-number* $i(G)$ is then defined as the minimum k, such that G has a k-interval representation. Since its introduction, the maximum interval-number among all graphs of specific graph classes has been analyzed. Harary and Trotter [HT79] determine the maximum interval-number of trees and complete bipartite graphs, Scheinerman and West [SW83] for planar and outer-planar graphs, and Griggs and West [GW80] for graphs with fixed maximum degree Δ. (West gave a short proof later [Wes89].) For all these (and several more) graph classes matching upper and lower bounds are known. We provide an overview in Table 4.2 below.

In 1994 Kumar and Deo [DK94] proposed the study of the so-called *dimensionality* of a graph, a different but closely related relaxation of interval graphs. (Let us mention that Tibor Gallai suggested this subject already in 1968 and that combinatorial properties of k-interval graphs and k-track graphs have been studied already back in 1969 [GL69], however with different intention and questions in mind.) We stick here to the following notion of Gyárfás and West [GW95] from 1995. In a *k-track representation* of a simple graph G there are k parallel lines, called *tracks*. Every vertex is

Graph Class $\mathcal{G}$	$i(G)$	$t_\ell(G)$	$t(G)$	$b(G)+1$
Tree	2 [HT79]	2	2 [DK94]	2 [GLS07]
Outer-Planar	2 [SW83]	2	2 [KW99]	3 Thm. 4.2.1
Planar + Bipartite	3	3	4 [GO05, Gon07]	3 [BS09]
Planar + 4-Connected	3	$3 \dots 4$	4	$3 \dots 4$ Thm. 4.2.4
Planar 3-Tree	3	3	4	4 Thm. 4.2.2
Planar	3 [SW83]	$3 \dots 4$	4 [Gon07]	$4 \dots 5$ Thm. 4.2.7
$K_{m,n}$	$\lceil \frac{mn+1}{m+n} \rceil$ [HT79]	$\lceil \frac{mn}{m+n-1} \rceil$ [HKU11]	$\lceil \frac{mn}{m+n-1} \rceil$ [GW95]	$\lceil \frac{mn+\sqrt{mn}}{m+n} \rceil \dots 2m-1$ Lem. 4.1.1, [AS09]
Line Graph*	2	2	∞ [Mil11]	3 [BS09]
Tree-Width k	$k+1$	$k+1$	$k+1$ [DOSV98]	$\leq 2k$ Cor. 4.3.6
Degeneracy k	$k+1$	$k+1$ [HKU11]	$2k$ [HKU11]	$2k$ Thm. 4.3.3
Max Degree Δ	$\lceil \frac{\Delta+1}{2} \rceil$ [GW80]	$\leq \lceil \frac{\Delta}{2} \rceil + 1$ [HKU11]	$\leq \frac{\Delta}{2} + \mathcal{O}(\log \Delta)$ [AMR92]	$\lceil \frac{\Delta}{2} \rceil + 1 \dots \Delta + 1$ Thm. 4.3.10

Table 4.2: Summary of lower and upper bounds on the maximum interval-number $i(G)$, local track-number $t_\ell(G)$, track-number $t(G)$, and bend-number $b(G)$ among all graphs of some graph classes (* = without subgraphs)

associated with one interval from each track. Again vertex adjacency is equivalent to interval intersection and the *track-number* $t(G)$ is the minimum k, such that G has a k-track representation. It is easy to see that $i(G) \leq t(G)$, since a k-track representation can be transformed into a k-interval representation by putting the tracks in any order on a single real line. Equality does not hold in general as discussed below. The same set of extremal questions has been considered for the track-number as well. Kumar and Deo [DK94] determine the maximum track-number of trees, Kostochka and West [KW99] for outer-planar graphs, Gonçalves [GO05, Gon07] for planar graphs, and Gyárfás and West [GW95] for complete bipartite graphs. Again, for all these (and several more) graph classes matching upper and lower bounds are known and listed in Table 4.2.

We propose [HKU11] a third generalization as follows. In a *local k-track representation* of a simple graph G there are arbitrarily many parallel lines, called *tracks*. Every vertex is associated with at most k intervals, no two from the same track. Again vertex adjacency is equivalent to interval intersection and the *local track-number* $t_\ell(G)$ is the minimum k, such that G has a local k-track representation. It is easy to see that $i(G) \leq t_\ell(G) \leq t(G)$, since a k-track representation is a local k-track representation and a local k-track representation can be transferred into a k-interval representation by putting the tracks in any order on a single real line.

Now recall that in a k-bend representation every vertex is associated with a grid path with at most $k+1$ segments and edges correspond to intersections of segments of non-zero length. A k-bend representation can be transformed into a local $(k+1)$-track representation (and hence to a $(k+1)$-interval representation) by considering the grid lines supporting some segment as the set of tracks. However, this is a local track representation only if no two segments of the same vertex path are supported by the same grid line, which is the case in every EPG representation in the literature that we know of, including this thesis. Hence, all lower bounds for the interval-number and local track-number in Table 4.2 carry over to the bend-number plus 1. On the other hand, all upper bounds on the bend-number plus 1 in Table 4.2 carry over to the local track-number and hence to the interval-number.

In general, the bend-number of a graph G is at most $4(i(G) - 1)$, since a set of k intervals can be traced by a $4(i(G) - 1)$-bend path without creating unwanted edge-intersections. Similarly, one can show that $b(G) \leq 4(t_\ell(G) - 1) \leq 4(t(G) - 1)$. In particular, the bend-number is always bounded in terms of the interval-number and local track-number. Perhaps surprisingly, this is not true for the track-number. As we conjectured [HKU10] and K. Milans proved [Mil11], the track-number of $L(K_n)$, i.e., the line graph of the complete graph on n vertices, goes to infinity with $n \to \infty$. On the other hand, every line graph admits a 2-bend representation [BS09], which even inherits a local 2-track representation.

Another interesting aspect concerns the Lower-Bound-Lemma (c.f. Lemma 4.1.1), which states that in a simple k-bend representation of a graph $G = (V, E)$ with $|L|$ supporting grid lines, we have $|E| + |L| \leq (k+1)|V|$. The same statement (and the same proof) holds for simple k-interval representations, simple local k-track representations, and simple k-track representations. In the first and last case we know the number $|L|$ of supporting grid lines exactly, i.e., $|L| = 1$ and $|L| = k$, respectively. For a local k-track representation, we have $|L| \geq k$. In any case, the Lower-Bound-Lemma proves tight lower bounds for the interval-number, track-number and local track-number of all complete bipartite graphs. Indeed, the Lower-Bound-Lemma has been discovered several times [DK94, HT79, HKU10]. In case of the bend-number of $K_{m,n}$, the Lower-Bound-Lemma is tight only for some values of m and n, e.g., when $n = m$ (Theorem 4.1.3) and when $n = (m+1)(m-2)$ for even m and $n = m(m-2)$ for odd m (Theorem 4.1.4).

To end this section, consider the recognition problem associated with each graph representation. It is known that interval graphs can be recognized in polynomial, even linear, time [BL76]. But recognizing graphs with a fixed interval number $k \geq 2$ is NP-complete [SW84]. Moreover, Gyárfás and West [GW95] proved NP-completeness of recognizing 2-track graphs, and recently Jiang [Jia10] proved this for k-track graphs with every fixed $k \geq 2$. In Section 4.4 we presented the first NP-completeness result for k-bend graphs, i.e., we show that recognizing 1-bend graphs is NP-complete. It remains open whether the recognition of k-bend graphs is NP-complete for every fixed $k \geq 2$ (c.f. Question 4.4.5).

Open Questions

Let us list some of the problems that remain unsolved in this thesis. Some of them evolved just recently, some were already posed by several people. The questions are ordered according to appearance of relevant material in this document.

Question 1. Is the α-orientation problem for planar graphs solvable in linear time?

Description: We are given a planar graph $G = (V, E)$ and a mapping $\alpha : V \to \mathbb{N}$. Can one decide in linear time whether there exists an orientation of the edges of G, such that the out-degree at v equals $\alpha(v)$ for every vertex $v \in V$?

Comments: If the answer is 'YES', one can compute the minimal α-orientation in linear time by Remark 1.2.6. Conversely, computing the minimal α-orientation in linear time, would answer Question 1 in the affirmative. Furthermore, a similar question asks, whether there is an orientation of G, such that out-deg$(v) \leq \alpha(v)$ for every $v \in V(G)$. As before, it is not know how to answer it in linear time.

Question 2. What is the relation between $\square$-representations and Schnyder woods with $L_1 \cap L_2 = \{v_{n-1}\}$?

Description: We are given a 4-connected maximally planar graph G with an embedding with outer triangle $\{v_1, v_2, v_n\}$ and the corresponding 4-connected near-triangulation $G \setminus (v_1, v_n)$ with fourth outer vertex v_{n-1}. When do different $\square$-representations of $G \setminus (v_1, v_n)$ define the same Schnyder wood of G? What is a good characterization of the Schnyder wood corresponding to the minimal $\square$-representation?

Comments: How a Schnyder wood of G is inherited from a $\square$-representation of $G \setminus (v_1, v_n)$ is given by Definition 2.1.2. We refer to Remark 2.1.5 for further information.

Question 3. Does every maximally planar n-vertex graph admit a hole-free rectilinear representation with at most $\lceil \frac{4n}{3} \rceil$ supporting grid lines?

Comments: By Theorem 2.1.8 $\lfloor \frac{5n-2}{3} \rfloor$ supporting grid lines are always enough. Liao, Lu and Yen [LLY03] prove the same upper bound, as well as, a lower bound of $\lceil \frac{4n}{3} \rceil$.

Note that the 2×2 grid is considered to consist of four grid lines. For additional considerations see Subsection 2.1.1.

Question 4. Can a nesting assignment be found in linear time?

Description: We are given a maximally planar graph. Can one decide in linear time whether it admits a nesting assignment? And if so, can it be found in linear time?

Comments: Nesting assignments are defined in Definition 2.2.2. One may assume or not, that the graph is given with a fixed embedding. Sun and Sarrafzadeh [SS93] present an algorithm with $\mathcal{O}(n^{3/2})$ runtime, where n is the number of vertices.

Question 5. When is a rectilinear layout rectilinear area-universal?

Comments: The definition of rectilinear area-universality is given in Section 3.1. By Theorem 3.1.3 a $\square$-layout is rectilinear area-universal if and only if it is one-sided. From Lemma 3.1.5 follows that a layout is area-universal if some refinement of it is a one-sided $\square$-layout. However, this is not necessary as illustrated in Figure 4.38.

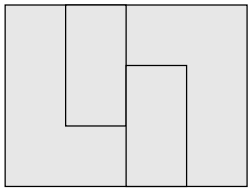

Figure 4.38: An area-universal rectilinear layout, which does not have a refinement that is a one-sided rectangle-representation.

Question 6. Can a cartogram for a one-sided $\square$-layout be computed in polynomial time?

Description: We are given a one-sided $\square$-layout Λ for a near-triangulation $G = (V, E)$ and a weight function $w : V \to \mathbb{R}^+$. Can one compute in polynomial time a $\square$-representation of G that realizes Λ and the given weights?

Comments: See Section 3.1 for one-sided layouts. There is an algorithm based in numerical iteration and some very practical implementations of this and other ideas. However, none of them ends in finite time with the exact solution. We briefly discuss this issue in Remark 3.1.8. Computing *any* $\square$-representation of G respecting the given weights, and not necessarily realizing Λ, in polynomial time would be equally good.

Question 7. Can a hole-free rectilinear cartogram of a maximally planar graph with complexity at most 8 be computed in linear time?

Description: We are given a maximally planar graph $G = (V, E)$ and a weight function $w : V \to \mathbb{R}^+$. Can one compute in linear time a hole-free rectilinear cartogram Γ of G w.r.t. w with complexity at most 8?

Comments: By Theorem 3.1.7 such a cartogram exists. By Lemma 3.1.6 a suitable rectilinear area-universal layout can be computed in linear time. The answer to Question 7 is 'YES' for Hamiltonian maximally planar graphs, given with a Hamiltonian cycle, (Theorem 3.2.2), as well as for planar 3-trees. For the latter case, one easily sets up a recursive construction similar to the one in Subsection 3.3.1.

Question 8. Does every 4-connected maximally planar graph admit a cartogram with solely convex polygons?

Description: We are given a 4-connected maximally planar graph $G = (V, E)$, i.e., G has no separating triangle, and a weight function $w : V \to \mathbb{R}^+$. Is there a (not necessarily rectilinear nor hole-free) cartogram of G w.r.t. w in which every vertex is represented by a convex polygon?

Comments: By Lemma 3.3.1 the answer is 'NO' for general maximally planar graphs; even for Hamiltonian 3-trees. We conjecture that the answer to Question 8 is 'YES' (c.f. Conjecture 3.4.1). Conjecture 3.4.4, which states that every two-sided $\square$-layout is area-universal, is even much stronger, since this would give a cartogram of G that additionally is hole-free and has polygonal complexity at most 4.

Question 9. Does every 4-connected maximally planar graph admit a Hamiltonian cycle that is a canonical order?

Description: We are given a 4-connected maximally planar graph G on n vertices, i.e., G has no separating triangle. Is there an embedding of G and a Hamiltonian cycle $(v_1, \ldots, v_n)$ in G, such that (v_1, v_n) is an outer edge and $(v_1, v_2, \ldots, v_n)$ is a canonical order of G w.r.t. (v_1, v_n).

Comments: Canonical order are defined in Subsection 1.1.3. In Subsection 3.2.1 such a Hamiltonian cycle is called one-sided Hamiltonian cycle. Moreover it is shown that equivalently G admits a Schnyder wood (T_1, T_2, T_3) in which every inner vertex is a leaf in T_1 or T_2. An affirmative answer to Question 9 would imply that every 4-connected maximally planar graph admits a rectilinear cartogram of polygonal complexity at most 6 (c.f. Conjecture 3.4.5).

Question 10. Is the worst-case complexity of hole-free cartograms for tri-connected planar graphs finite?

Description: We are given a pair (G, w) of a tri-connected planar graph $G = (V, E)$ and a weight function $w : V \to \mathbb{R}^+$. Is there a $k \in \mathbb{N}$ such that every (G, w) admits a hole-free cartogram of complexity at most k?

Comments: The cartogram does not have to be rectilinear. Indeed, inner faces of length 5 or more are not representable with a rectilinear cartogram. We know very little about the question, despite that tri-connectedness is necessary, since every hole-free representation of $K_{2,n}$ for large n has large complexity.

Question 11. What is the bend-number of $K_{3,11}$ and $K_{4,5}$?

Comments: The bend-number is defined in Definition 4.0.7. We know that either number is 2 or 3, and we suspect that 3 is the right answer for both. This question illustrates our lack of understanding about the bend-number of complete bipartite graphs. For instance, $b(K_{2,n})$ is known for every n, but $b(K_{3,n})$ is not. See Figure 4.11 for all we know about $b(K_{5,n})$. We refer to Section 4.1 and in particular to Question 4.1.10 for further information.

Question 12. Is there a planar graph with bend-number 4?

Comments: By Theorem 4.2.7 every planar graph has bend-number at most 4 and there is a planar graph with bend-number 3. We suspect the answer to Question 12 to be 'YES', c.f. Conjecture 4.2.8.

Question 13. What is the maximum bend-number in terms of the maximum degree?

Comments: By Theorem 4.3.10 the bend-number of any graph is at most its maximum degree, and there is a graph whose maximum degree is twice its bend-number. In other words, the maximum bend-number among all graphs of maximum degree Δ lies between $\lceil \frac{\Delta}{2} \rceil$ and Δ, both included. We strongly suspect that neither the upper nor the lower bound already gives the right value.

Question 14. What is the complexity of recognizing k-bend graphs?

Comments: We have shown in Theorem 4.4.4 that recognizing 1-bend graphs, i.e., single-bend graphs, is NP-complete. However, we do not know how to modify the reduction to larger k. See also Question 4.4.5.

Bibliography

[ABF+11a] M. J. Alam, T. Biedl, S. Felsner, A. Gerasch, M. Kaufmann, and S. G. Kobourov. Linear-time algorithms for proportional contact graph representations. In *Proceedings of the 22nd International Symposium on Algorithms and Computation (ISAAC 2011)*, 2011.

[ABF+11b] M. J. Alam, T. Biedl, S. Felsner, M. Kaufmann, and S. G. Kobourov. Proportional contact representations of planar graphs. In *Proceedings of the 19th International Symposium on Graph Drawing (GD 2011)*, 2011.

[ABF+11c] M. J. Alam, T. Biedl, S. Felsner, M. Kaufmann, S. G. Kobourov, and T. Ueckerdt. Computing cartograms with optimal complexity. submitted, 2011.

[AC84] N. Alon and Y. Caro. On the number of subgraphs of prescribed type of planar graphs with a given number of vertices. In *Annals of Discrete Mathematics (20) Convexity and Graph Theory Proceedings of the Conference on Convexity and Graph Theory*, volume 87, pages 25–36. 1984.

[ACP87] S. Arnborg, D. G. Corneil, and A. Proskurowski. Complexity of finding embeddings in a k-tree. *SIAM J. Algebraic Discrete Methods*, 8:277–284, 1987.

[AHAdS05] J. S. Andrade, H. J. Herrmann, R. F. S. Andrade, and L. R. da Silva. Apollonian networks: Simultaneously scale-free, small world, euclidean, space filling, and with matching graphs. *Phys. Rev. Lett.*, 94:018702, 2005.

[AKS84] T. Asano, S. Kikuchi, and N. Saito. A linear algorithm for finding hamiltonian cycles in 4-connected maximal planar graphs. *Discrete Applied Mathematics*, 7(1):1–15, 1984.

[AMR92] N. Alon, C. McDiarmid, and B. Reed. Star arboricity. *Combinatorica*, 12:375–380, 1992.

[APC90] S. Arnborg, A. Proskurowski, and D. G. Corneil. Forbidden minors characterization of partial 3-trees. *Discrete Math.*, 80(1):1–19, 1990.

[AS09] A. Asinowski and A. Suk. Edge intersection graphs of systems of paths on a grid with a bounded number of bends. *Discrete Appl. Math.*, 157(14):3174–3180, 2009.

[BF10a] O. Bernardi and É. Fusy. A bijection for triangulations, quadrangulations, pentagulations, etc. 2010.

[BF10b] O. Bernardi and É. Fusy. Schnyder decompositions for regular plane graphs and application to drawing. 2010.

[BFM07] N. Bonichon, S. Felsner, and M. Mosbah. Convex drawings of 3-connected plane graphs. *Algorithmica*, 47:399–420, 2007.

[BGPV08] A. L. Buchsbaum, E. R. Gansner, C. M. Procopiuc, and S. Venkatasubramanian. Rectangular layouts and contact graphs. *ACM Transactions on Algorithms*, 4(1), 2008.

[BL76] K. S. Booth and G. S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using *PQ*-tree algorithms. *J. Comput. System Sci.*, 13(3):335–379, 1976.

[Bre00] E. Brehm. 3-*orientations and Schnyder* 3-*tree-decompositions*. Diplomarbeit, Freie Universität, Berlin, 2000.

[BRV11] T. Biedl and L. E. Ruiz Velázquez. Orthogonal cartograms with few corners per face. In *Proceedings of the Data Structures and Algorithms Symposium (WADS 2011)*, pages 98–109, 2011.

[BS90] M. L. Brady and M. Sarrafzadeh. Stretching a knock-knee layout for multilayer wiring. *IEEE Trans. Comput.*, 39(1):148–151, 1990.

[BS09] T. Biedl and M. Stern. Edge-intersection graphs of k-bend paths in grids. In *Proceedings of the 15th Annual International Conference on Computing and Combinatorics (COCOON '09)*, pages 86–95, 2009.

[BW00] U. Brandes and D. Wagner. A linear time algorithm for the arc disjoint menger problem in planar directed graphs. *Algorithmica*, 28:16–36, 2000.

[CG09] J. Chalopin and D. Gonçalves. Every planar graph is the intersection graph of segments in the plane. In *Proceedings of the 41st annual ACM symposium on Theory of computing (STOC 2009)*, pages 631–638, 2009.

[CGH+98] R. Chuang, A. Garg, X. He, M.-Y. Kao, and H.-I Lu. Compact encodings of planar graphs via canonical orderings and multiple parentheses. In *Automata, Languages and Programming*, volume 1443, pages 118–129. 1998.

[Chv85] V. Chvátal. *Hamiltonian cycles*, chapter 11, pages 403–429. Wiley, New York, 1985.

[CLL05] Y.-T. Chiang, C.-C. Lin, and H.-I Lu. Orderly spanning trees with applications. *SIAM J. Comput.*, 34:924–945, 2005.

[Cou90] B. Courcelle. Graph rewriting : An algebraic and logic approach. *Handbook of Theoretical Computer Science*, pages 194–242, 1990.

[dBMS06] M. de Berg, E. Mumford, and B. Speckmann. On rectilinear duals for vertex-weighted plane graphs. In *Proceedings of the 14th International Symposium on Graph Drawing (GD 2006)*, pages 61–72, 2006.

[dFOdM01] H. de Fraysseix and P. Ossona de Mendez. On topological aspects of orientations. *Discrete Mathematics*, 229(1-3):57–72, 2001.

[dFOdMR94] H. de Fraysseix, P. Ossona de Mendez, and P. Rosenstiehl. On triangle contact graphs. *Combinatorics, Probability and Computing*, 3(02):233–246, 1994.

[dFPP90] H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.

[Die10] R. Diestel. *graph theory*. Graduate Texts in Mathematics. Springer, 4th edition, 2010.

[DK94] N. Deo and N. Kumar. Multidimensional interval graphs. In *Proceedings of the 25th Southeastern International Conference on Combinatorics, Graph Theory, and Computing*, volume 102, pages 45–56, 1994.

[DOSV98] G. Ding, B. Oporowski, D. P. Sanders, and D. Vertigan. Partitioning graphs of bounded tree-width. *Combinatorica*, 18(1):1–12, 1998.

[EFK01] M. Eiglsperger, S. Fekete, and G. Klau. Orthogonal graph drawing. In *Drawing Graphs*, volume 2025 of *Lecture Notes in Computer Science*, pages 121–171. 2001.

[EH66] P. Erdős and A. Hajnal. On chromatic number of graphs and set-systems. *Acta Math. Acad. Sci. Hungar*, 17:61–99, 1966.

[EMC90] E. S. El-Mallah and C. J. Colbourn. On two dual classes of planar graphs. *Discrete Math.*, 80(1):21–40, 1990.

[EMSV09] D. Eppstein, E. Mumford, B. Speckmann, and K. Verbeek. Area-universal rectangular layouts. In *Proceedings of the 25th annual symposium on Computational geometry (SOCG 2009)*, pages 267–276, 2009.

[Epp10] D. Eppstein. Regular labelings and geometric structures. In *Proceedings of 22nd Canadian Conference on Computational Geometry (CCCG 2010)*, 2010.

[Fel01] S. Felsner. Convex drawings of planar graphs and the order dimension of 3-polytopes. *Order*, 18:19–37, 2001.

[Fel03] S. Felsner. Geodesic embeddings and planar graphs. *Order*, 20:135–150, 2003.

[Fel04] S. Felsner. Lattice structures from planar graphs. *Electron. J. Combin.*, 11(1):Research Paper 15, 24 pp., 2004.

[Fel11] S. Felsner. Rectangle and square representations of planar graphs. submitted, 2011.

[FFNO11] S. Felsner, É. Fusy, M. Noy, and D. Orden. Bijections for Baxter families and related objects. *Journal of Combinatorial Theory, Series A*, 118(3):993–1020, 2011.

[FK09] S. Felsner and K. Knauer. ULD-lattices and Δ-bonds. *Combinatorics, Probability and Computing*, 18(5):707–724, 2009.

[FK11] S. Felsner and K. Knauer. Distributive lattices, polyhedra, and generalized flows. *European Journal of Combinatorics*, 32(1):45–59, 2011.

[Fus07] É. Fusy. *Combinatoire des cartes planaires et applications algorithmiques*. PhD Thesis, 2007.

[GHK10] E. R. Gansner, Y. Hu, and S. G. Kobourov. On touching triangle graphs. In *Proceedings of the 18th International Symposium on Graph Drawing (GD 2010)*, pages 250–261, 2010.

[GHKK10] E. R. Gansner, Y. Hu, M. Kaufmann, and S. G. Kobourov. Optimal polygonal representation of planar graphs. In *Theoretical Informatics (LATIN 2010)*, volume 6034 of *Lecture Notes in Computer Science*, pages 417–432. 2010.

[GJ79] M. R. Garey and D. S. Johnson. *Computers and intractability.* W. H. Freeman and Co., San Francisco, Calif., 1979. A guide to the theory of NP-completeness, A Series of Books in the Mathematical Sciences.

[GL69] A. Gyárfás and J. Lehel. A Helly type problem in trees. *Coll. Math. Soc. J'anos Bolyai*, (4):571–584, 1969.

[GL86] P. M. Gilmer and R. A. Litherland. The duality conjecture in formal knot theory. *Osaka J. Math.*, 23(1):229–247, 1986.

[GLS07] M. C. Golumbic, M. Lipshteyn, and M. Stern. Edge intersection graphs of single bend paths on a grid. In *Proceedings of the 6th Cologne Twente Workshop on Graphs and Combinatorial Optimization (CTW 2007)*, pages 53–55, 2007.

[GLS08] M. C. Golumbic, M. Lipshteyn, and M. Stern. Representing edge intersection graphs of paths on degree 4 trees. *Discrete Math.*, 308(8):1381–1387, 2008.

[GO05] D. Gonçalves and P. Ochem. On some arboricities in planar graphs. *Electronic Notes in Discrete Mathematics*, 22:427–432, 2005. 7th International Colloquium on Graph Theory.

[Gon07] D. Gonçalves. Caterpillar arboricity of planar graphs. *Discrete Math.*, 307(16):2112–2121, 2007.

[GW80] J. R. Griggs and D. B. West. Extremal values of the interval number of a graph. *SIAM J. Algebraic Discrete Methods*, 1(1):1–7, 1980.

[GW95] A. Gyárfás and D. B. West. Multitrack interval graphs. In *Proceedings of the 26th Southeastern International Conference on Combinatorics, Graph Theory, and Computing*, volume 109, pages 109–116, 1995.

[Hal76] R. Halin. S-functions for graphs. *Journal of Geometry*, 8:171–186, 1976.

[He93] X. He. On finding the rectangular duals of planar triangular graphs. *SIAM J. Comput*, 22:1218–1226, 1993.

[HKPS04] R. Heilmann, D. A. Keim, C. Panse, and M. Sips. Recmap: Rectangular map approximations. In *10th IEEE Symp. on Information Visualization (InfoVis 2004)*, pages 33–40, 2004.

[HKRS97] M. R. Henzinger, P. Klein, S. Rao, and S. Subramanian. Faster shortest-path algorithms for planar graphs. *Journal of Computer and System Sciences*, 55(1):3–23, 1997.

[HKU10] D. Heldt, K. Knauer, and T. Ueckerdt. Edge-intersection graphs of grid paths – the bend-number. submitted, 2010.

[HKU11] D. Heldt, K. Knauer, and T. Ueckerdt. Track-number and caterpillar arboricity of graphs of fixed treewidth, 2011. manuscript of work in progress.

[HP68] F. Harary and E. M. Palmer. On acyclic simplicial complexes. *Mathematika*, 15(01):115–122, 1968.

[HT79] F. Harary and W. T. Trotter. On double and multiple interval graphs. *J. Graph Theory*, 3(3):205–211, 1979.

[Jia10] M. Jiang. Recognizing d-interval graphs and d-track interval graphs. In *Proceedings of the 4th international conference on Frontiers in algorithmics (FAW 2010)*, pages 160–171, 2010.

[Kan92] G. Kant. Drawing planar graphs using the lmc-ordering. In *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science (FOCS 1992)*, pages 101–110, 1992.

[Kan96] G. Kant. Drawing planar graphs using the canonical ordering. *Algorithmica*, 16:4–32, 1996.

[KH97] G. Kant and X. He. Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. *Theoretical Computer Science*, 172(1-2):175–193, 1997.

[KK85] K. Koźmiński and E. Kinnen. Rectangular dual of planar graphs. *Networks*, 15(2):145–157, 1985.

[KMv+09] H. Kierstead, B. Mohar, S. Špacapan, D. Yang, and X. Zhu. The two-coloring number and degenerate colorings of planar graphs. *SIAM J. Discret. Math.*, 23:1548–1560, 2009.

[KMW10] P. N. Klein, S. Mozes, and O. Weimann. Shortest paths in directed planar graphs with negative lengths: A linear-space $\mathcal{O}(n \log^2 n)$-time algorithm. *ACM Trans. Algorithms*, 6:30:1–30:18, 2010.

[KN07] A. Kawaguchi and H. Nagamochi. Orthogonal drawings for plane graphs with specified face areas. In *Proceedings of the 4th international conference on Theory and applications of models of computation (TAMC 2007)*, pages 584–594, 2007.

[Kna10] K. Knauer. *Lattices and polyhedra from graphs.* PhD Thesis, Berlin, 2010.

[Koe36] P. Koebe. Kontaktprobleme der konformen Abbildung. *Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl.*, 88:141–164, 1936.

[KP88] P. G. Karger and B. T. Preas. Placement, assignment and floorplanning. *Physical Design Automation of VLSI Systems*, pages 87–155, 1988.

[KRRS94] P. N. Klein, S. Rao, M. Rauch, and S. Subramanian. Faster shortest-path algorithms for planar graphs. In *Proceedings of the 26th annual ACM symposium on Theory of computing (STOC '94)*, pages 27–37, 1994.

[KW99] A. V. Kostochka and D. B. West. Every outerplanar graph is the union of two interval graphs. In *Proceedings of the 30th Southeastern International Conference on Combinatorics, Graph Theory, and Computing*, volume 139, pages 5–8, 1999.

[LL84] Y.-T. Lai and S. M. Leinwand. An algorithm for building rectangular floor-plans. In *21st Design Automation Conference (DAC '84)*, pages 663–664, 1984.

[LLY03] C.-C. Liao, H.-I Lu, and H.-C. Yen. Compact floor-planning via orderly spanning trees. *Journal of Algorithms*, 48(2):441–451, 2003.

[LZ03] P. C. B. Lam and H. Zhang. A distributive lattice on the set of perfect matchings of a plane bipartite graph. *Order*, 20(1):13–29, 2003.

[Mil11] K. Milans. personal communication, 2011.

[MN95] G. L. Miller and J. Naor. Flow in planar graphs with multiple sources and sinks. 24(5):1002–1017, 1995.

[MWN10] S. Mozes and C. Wulff-Nilsen. Shortest paths in planar graphs with real lengths in $\mathcal{O}(n \log^2 n / \log \log n)$ time. In *Algorithms – ESA 2010*, volume 6347 of *Lecture Notes in Computer Science*, pages 206–217. 2010.

[OdM94] P. Ossona de Mendez. *Orientations bipolaires*. PhD Thesis, Paris, 1994.

[Ott88] R. H. J. M. Otten. Graphs in floor-plan design. *International Journal of Circuit Theory and Applications*, 16(4):391–410, 1988.

[Pro93] J. Propp. Lattice structure for orientations of graphs. 1993.

[PS03] D. Poulalhon and G. Schaeffer. Optimal coding and sampling of triangulations. In *Automata, Languages and Programming*, volume 2719, pages 191–191. 2003.

[Rai34] E. Raisz. The rectangular statistical cartogram. *Geographical Review*, 24(3):292–296, 1934.

[Rém04] E. Rémila. The lattice structure of the set of domino tilings of a polygon. *Theoret. Comput. Sci.*, 322(2):409–422, 2004.

[Rie09] B. Ries. Some properties of edge intersection graphs of single bend path on a grid. *Electronic Notes in Discrete Mathematics*, 34:29–33, 2009. European Conference on Combinatorics, Graph Theory and Applications (EuroComb 2009).

[Rin87] I. Rinsma. Nonexistence of a certain rectangular floorplan with specified areas and adjacency. *Environment and Planning B: Planning and Design*, 14:163–166, 1987.

[RS86] N. Robertson and P. D. Seymour. Graph minors. ii. algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986.

[RT86] P. Rosenstiehl and R. Tarjan. Rectilinear planar layouts and bipolar orientations of planar graphs. *Discrete and Computational Geometry*, 1:343–353, 1986.

[Sch78] T. J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th annual ACM symposium on Theory of computing (STOC '78)*, pages 216–226, 1978.

[Sch89] W. Schnyder. Planar graphs and poset dimension. *Order*, 5:323–343, 1989.

[Sch90] W. Schnyder. Embedding planar graphs on the grid. In *Proceedings of the 1st ACM-SIAM Symposium on Discrete Algorithms (SODA '90)*, pages 138–148, 1990.

[She74] G. Shephard. Subpolytopes of stack polytopes. *Israel Journal of Mathematics*, 19:292–296, 1974.

[SS93] Y. Sun and M. Sarrafzadeh. Floorplanning by graph dualization: L-shaped modules. *Algorithmica*, 10:429–456, 1993.

[SW83] E. R. Scheinerman and D. B. West. The interval number of a planar graph: three intervals suffice. *J. Combin. Theory Ser. B*, 35(3):224–239, 1983.

[SW84] D. B. Shmoys and D. B. West. Recognizing graphs with fixed interval number is NP-complete. *Discrete Appl. Math.*, 8(3):295–305, 1984.

[SY93] M. Sarrafzadeh and K.-H. Yeap. Floor-planning by graph dualization: 2-concave rectilinear modules. *SIAM Journal on Computing*, 22(3):500–526, 1993.

[Tam87] R. Tamassia. On embedding a graph in the grid with the minimum number of bends. 16(3):421–444, 1987.

[Tho86] C. Thomassen. Interval representations of planar graphs. *Journal of Combinatorial Theory, Series B*, 40(1):9–20, 1986.

[Thu90] W. P. Thurston. Conway's tiling groups. *Amer. Math. Monthly*, 97(8):757–773, 1990.

[Ung53] P. Ungar. On diagrams representing graphs. *Journal of the London Mathematical Society*, 28:336–342, 1953.

[Viz64] V. G. Vizing. On an estimate of the chromatic class of a p-graph. *Diskret. Analiz No.*, 3:25–30, 1964.

[vKS07] M. van Kreveld and B. Speckmann. On rectangular cartograms. *Computational Geometry*, 37(3):175–187, 2007.

[Wag37] K. Wagner. Über eine Eigenschaft der ebenen Komplexe. *Mathematische Annalen*, 114:570–590, 1937.

[Wes89] D. B. West. A short proof of the degree bound for interval number. *Discrete Math.*, 73(3):309–310, 1989.

[Whi31] H. Whitney. A theorem on graphs. *Annals of Mathematics*, 32:378–390, 1931.

[WKC88] S. Wimer, I. Koren, and I. Cederbaum. Floorplans, planar graphs, and layouts. *Circuits and Systems, IEEE Transactions on*, 35(3):267–278, 1988.

Index