GEOMETRIC AND COMBINATORIAL STRUCTURES ON GRAPHS

vorgelegt von
Dipl.-Math. Florian Zickfeld

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

Doktor der Naturwissenschaften
Dr. rer. nat.

genehmigte Dissertation

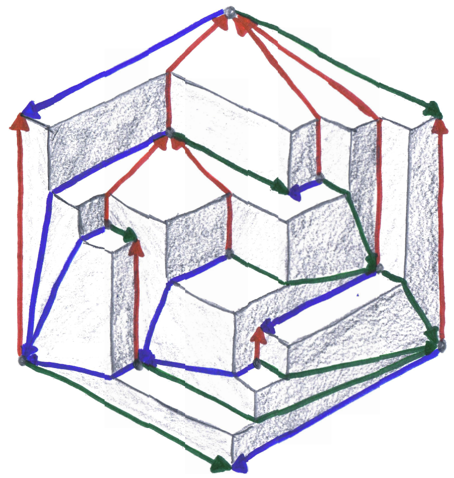

Promotionsausschuss

Vorsitzender:	Prof. Dr. Michael Scheutzow
Berichter:	Prof. Dr. Stefan Felsner
	Prof. Dr. Graham Brightwell

Tag der wissenschaftlichen Aussprache:

20. Dezember 2007

Berlin 2007
D 83

Acknowledgments

First of all I want to thank my advisor, Stefan Felsner. In lectures and many discussions I learned a lot from him not only about graph theory and combinatorics. He was also willing to share the large and little tricks and insights that make (scientific) working so much easier. During lunches and many coffee breaks I also learned a lot from Stefan about life, the universe, and everything [3] and I appreciate these insights as much as the mathematical ones.

Günter M. Ziegler has worked with me on the chapter about integer realizations of stacked polytopes and I have enjoyed learning math from him. I am also very grateful for the financial support, ranging from paying for dessert in Victoria, Canada to offering me a Günter M. Ziegler DFG Leibniz grant for more than a year.

I thank Paul Bonsma for introducing me to the subject of leafy trees and for inspiring joint work on the subject. I am grateful to Bruno Benedetti, Nicolas Bonichon, Páidí Creed, Eric Fusy, Bernd Gärtner, Michael Hoffmann, Mark Jerrum, Christian Krattenthaler, and Raman Sanyal for valuable hints and discussions related to this thesis.

I particularly thank Graham Brightwell for agreeing to referee this thesis. I am deeply indebted to Paul Bonsma, Sarah Kappes, Kolja Knauer, Christian Liebchen, Mareike Massow, Guido Schäfer, Björn Stenzel, Axel Werner, and Gregor Wünsch for proof-reading, comments, and technical support. Not only the aforementioned but the whole discrete math group at TU Berlin (including both coffee machines) has provided an outstanding working environment during the last two and a half years. I also enjoyed the two months that I spent with Emo Welzl's group at ETH Zurich very much.

Without the financial support of the Studienstiftung des deutschen Volkes and the research training groups "Combinatorics, Geometry, and Computation" and "Methods for Discrete Structures" I would not have been able to conduct the research presented in this thesis. I thank the organizations and the people behind them for making this possible, in particular Gabriele Klink and Elke Pose.

Finally I want to express my gratitude to those people who have not been involved in the actual process of writing this thesis. Their main points of contact with this process were my bad moods every once in a while, probably the least pleasant of all ways of being involved in this whole thing. Among these people are of course my folks, in particular my parents, who not only provided superb all inclusive holidays in Saarbrücken. One of the things I appreciate most is how good they were at knowing when "How is your thesis going?" is not the right question to ask. Greta has been equally good at avoiding this question and also at spending relaxing weekends with me that made the final months of working on the thesis much better than I had hoped for.

Well... that's it. So long, and thanks for all the coffee![1]

[1] See [4].

Contents

Introduction

The title “Geometric and Combinatorial Structures on Graphs” of this thesis is quite general. The reason is that we treat four main topics that are hard to summarize by a single keyword. The four topics are treated in Chapters 2-5, and although they are rather different a number of connections exist. We first sketch each of these four topics as well as the content of Chapter 1. We will then be able to explain how the different chapters are connected. The following chapter outlines are meant to give a succinct overview of the thesis. Each chapter starts with an introduction that includes more context and references.

Chapter 1: Introduction to Schnyder Woods. In this chapter we define Schnyder woods and summarize facts about them that we use in the rest of the thesis. We also present some results about the number of edge splits and edge merges that can be applied to a Schnyder wood. Furthermore, we characterize all planar maps with a unique Schnyder wood.

Chapter 2: Schnyder Woods and Orthogonal Surfaces. In this chapter we present joint work with Stefan Felsner. Parts of this chapter can be found in [45, 48].

Schnyder woods and orthogonal surfaces are closely related, and their connections yield fruitful insights about both objects. We exploit these connections for a new proof of the Brightwell-Trotter Theorem about the dimension of planar graphs. Our proof follows a very intuitive approach for the construction of a rigid orthogonal surface and can be translated into an efficient algorithm for the construction of a Brightwell-Trotter realizer. We also propose two types of efficient representations for orthogonal surfaces. Both representations use a Schnyder wood to encode the combinatorics of the surface and a small set of real numbers to encode the geometry of the surface. The first representation is restricted to coplanar orthogonal surfaces and generalizes the well known face-counting method by assigning a weight to each face. The second representation can also be used for non-coplanar surfaces and therefore needs a value, the so-called height, for every face and every vertex of the Schnyder wood. In this chapter we also show that the Schnyder wood shown on the title page cannot be embedded on a coplanar and simultaneously rigid orthogonal surface.

Chapter 3: The Number of Planar Orientations with Prescribed Out-Degrees. In this chapter we present joint work with Stefan Felsner. Parts of this chapter can be found in [46, 47].

The concept of orientations with prescribed out-degrees can be used to describe many well-known structures on planar maps. The use of this unifying description enables us to determine bounds for the maximum number of orientations for several out-degree functions and classes of planar maps. Besides proving bounds that are valid for every map and out-degree function we consider the numbers of Eulerian orientations, Schnyder woods, 2-orientations, and bipolar orientations in more detail. For each structure we present an infinite family of graphs to obtain a lower bound. These families are all close relatives of the grid graph. We obtain the upper bounds for Eulerian orientations, Schnyder woods, and 2-orientations as specializations of a rather general technique. This technique makes use of spanning trees whose leaves contain a big independent set. We conclude the chapter with a few results about the complexity of counting planar orientations with prescribed out-degrees.

Chapter 4: Spanning Trees with Many Leaves. In this chapter we present joint work with Paul Bonsma. Parts of this chapter can be found in [17].

We are concerned with the maximization of the number of leaves of spanning trees. This problem is known to be $\mathcal{NP}$-hard. We prove for different graph classes that a certain fraction of the number of vertices of the graph can be guaranteed to be leaves. Let $n_{\neq 2}(G)$ denote the number of vertices of a graph G that do not have degree 2 and $n_{\geq 3}(G)$ the number of vertices that have degree at least 3. The main result of the chapter is that every graph without certain subgraphs called necklaces and blossoms has a spanning tree with $n_{\geq 3}(G)/3 + 4/3$ leaves. We also discuss a few corollaries of this main result and some other bounds. For example we prove that every graph G has a spanning tree with at least $n_{\neq 2}(G)/4 + 3/2$ leaves. And if G contains no triangles, then it has a spanning tree with at least $n_{\neq 2}(G)/3 + 2/3$ leaves. Our results strengthen and generalize several known ones by extending an established proof method. This method iteratively extends a partial tree until it becomes spanning and guarantees that the number of leaves of every intermediate tree satisfies an inequality closely related to the bound that is to be proved.

Chapter 5: Small Integer Realizations of Stacked Polytopes. In this chapter we present joint work with Günter M. Ziegler.

We discuss realizations of stacked polytopes with integral vertex coordinates. We give polynomial bounds for the absolute value of the vertex coordinates for three subclasses of stacked polytopes. The first subclass are linear stacked polytopes and we construct realizations by explicitly defining a lifting function for tailored drawings of the skeleton graph. The second subclass are balanced stacked polytopes. We construct special Tutte embeddings for the skeleton graph and then use a known lifting framework for Tutte embeddings to construct the realizations. We then show how linear and balanced stacked polytopes can be glued together to form so-called brooms with small integer coordinates.

We use the remainder of this introduction to explain the connections between the individual chapters. One link between the chapters is the use of Schnyder woods. While Chapter 1 gives an introduction to these objects, all of Chapter 2 is concerned with Schnyder woods and their connections with orthogonal surfaces. The interest in the number of Schnyder woods was the starting point for the research that we present in Chapter 3. In Section 3.2 we study the maximum number of Schnyder woods that a planar map can have. While we studied bounds for the number of Schnyder woods it turned out that their encodings as 3-orientations are well suited for this problem. Furthermore, the methods that can be applied for 3-orientations are also useful for other interesting orientations with prescribed out-degrees that we consider in Chapter 3.

As mentioned above, spanning trees with a large independent set among their leaves play an important role in Chapter 3. In Section 3.3 we are concerned with 2-orientations of quadrangulations. In these bipartite graphs a spanning tree with k leaves automatically has an independent set of size $k/2$ among its leaves. We can therefore use the results from Chapter 4 which imply that every quadrangulation has a spanning tree with at least $n/3$ leaves to obtain a spanning tree with an independent set of leaves of size at least $n/6$.

Stacked triangulations also appear in several contexts throughout the thesis. Chapter 5 discusses small integer realization of polytopes whose skeleton is a stacked triangulation. Furthermore, stacked triangulations are exactly the triangulations with a unique Schnyder wood. We use this fact for example in Section 1.4, and in Section 1.5 we generalize it by giving a characterization of all planar maps with a unique Schnyder wood. Moreover, the unique Schnyder wood of a stacked triangulation T with n vertices can be used to show that T has a spanning tree with at least $2n/3 + 1/3$ leaves, see Section 4.5. We also show that the number of bipolar orientations of stacked triangulations can be determined exactly in Section 3.4. Bipolar orientations in turn also appear in Section 2.3 where we discuss the height representations of orthogonal surfaces.

We conclude with a few remarks about the interdependence of the chapters. The understanding of Chapter 2 requires that the reader is familiar with all aspects of the theory of Schnyder woods that we present in Sections 1.1–1.4 of Chapter 1. Section 3.2 relies on the encoding of Schnyder woods as orientations with prescribed out-degrees which is introduced in Section 1.3. The rest of Chapter 3 is self-contained and Chapter 4 can be read independently except for Section 4.5 which uses Schnyder woods. Although graph drawings using Schnyder woods are mentioned in Chapter 5 this chapter can be read independently of the rest of the thesis.

Chapter 1

Introduction to Schnyder Woods

In two fundamental papers [82, 83] Walter Schnyder developed a theory of Schnyder woods and Schnyder labelings for planar triangulations. In [82], he presented a characterization of planar graphs in terms of order dimension which has stimulated subsequent research, see e.g. [19, 20, 38]. Section 2.1 of this thesis is concerned with this aspect of the theory of Schnyder woods. In [83] Schnyder deals with grid drawings of planar graphs and gives the first of numerous applications of Schnyder woods in the area of graph drawing. The results in [66, 7, 13] are other examples of such applications and the topic of Section 2.2 is related to this aspect of Schnyder woods. More references related to Schnyder woods can be found in [37].

We make extensive use of known results about Schnyder woods in Chapter 2 and in Section 3.2. Therefore we introduce the necessary background in this chapter and complement it with a few new results in Sections 1.4 and 1.5. The chapter is organized as follows. We start with the definition and the essential properties of Schnyder woods in Section 1.1. In Section 1.2 we explain their connections with orthogonal surfaces. The encoding of Schnyder woods as graph orientations with prescribed out-degrees that we introduce in Section 1.3 will be used in Sections 2.3 and 3.2. In Section 1.4, we introduce the so-called edge split and edge merge operations. We also consider the minimum and maximum number of such operations that can be applied to a Schnyder wood. Finally, in Section 1.5 we present a constructive characterization of all graphs with a unique Schnyder wood.

In [40] Felsner gives a comprehensive introduction to Schnyder woods which also contains many of the proofs that we omit here. The omitted proofs of results that are not in [40] can be found in one of [38, 39, 41].

1.1 Basics on Schnyder Woods

A *planar map* M is a simple planar graph together with a fixed crossing-free embedding in the plane. In particular, M has a designated outer (unbounded) face. We denote the sets of vertices, edges and faces of a given planar map by $V(M)$, $E(M)$, and $\mathcal{F}(M)$, and their respective cardinalities by $n(M)$, $m(M)$, and $f(M)$. The degree of a vertex v will be denoted by $d(v)$. If it is clear from the context which map we refer to, we simply write V instead of $V(M)$, and similarly for the other parameters.

planar map

$V(M)$, $E(M)$, $\mathcal{F}(M)$

$n(M)$, $m(M)$, $f(M)$

$d(v)$

special/suspension vertex
suspended map M^σ

Let a_1, a_2, a_3 be three vertices occurring in clockwise order on the outer face of M. We call a_i a *special vertex* or a *suspension vertex* of M. A *suspended map M^σ* is obtained by attaching a half-edge that reaches into the outer face to each of the special vertices.

Schnyder wood

Let M^σ be a suspended 3-connected planar map. A *Schnyder wood* rooted at a_1, a_2, a_3 is an orientation and coloring of the edges of M^σ with the colors $1, 2, 3$ (alternatively: red, green, blue) satisfying the following rules. We assume a cyclic structure on the labels so that $i + 1$ and $i - 1$ are always defined.

(W1) Every edge e is oriented in one direction or in two opposite directions. If e is bidirected, then the two directions have different colors.

(W2) The half-edge at a_i is directed outwards and colored i.

(W3) Every vertex v has out-degree 1 in each color. The edges e_1, e_2, e_3 leaving v in colors $1, 2, 3$ occur in clockwise order. Each incoming edge of v in color i enters v in the clockwise sector between e_{i+1} and e_{i-1}, see Figure 1.1 (a).

(W4) The boundary of an interior face is not a monochromatic directed cycle.

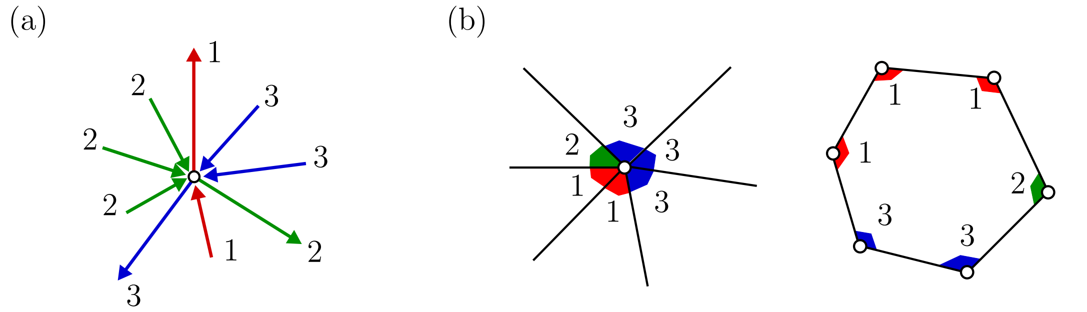

Figure 1.1. Part (a): Rule (W3). The numbers indicate the edge colors. Part (b): Rules (A2) and (A3). The numbers indicate the angle colors.

An existence proof for Schnyder woods on 3-connected planar maps can be found for example in [40]. We now also introduce Schnyder labelings, since they are useful for proving some facts about Schnyder woods. Let M^σ be a suspended

Schnyder labeling

3-connected planar map. A *Schnyder labeling* with respect to a_1, a_2, a_3 is a labeling of the angles of M^σ with the labels $1, 2, 3$ satisfying three rules.

(A1) The two angles at the half-edge of the special vertex a_i have labels $i + 1$ and $i - 1$ in clockwise order.

(A2) The labels of the angles at each vertex form, in clockwise order, nonempty intervals of 1's, 2's, and 3's, see Figure 1.1 (b).

(A3) The labels of the angles at each face form, in clockwise order, nonempty intervals of 1's, 2's, and 3's, see Figure 1.1 (b).

We want to point out a subtlety related to (A3). When M^σ endowed with a Schnyder labeling is embedded in the plane $\mathbb{R}^2$, then the labels of the outer face form non-empty intervals of 1's, 2's, 3's in *clockwise order*. When M^σ is embedded on the sphere, then the labels of the outer face form non-empty intervals of 1's, 2's, 3's in *counterclockwise order*.

The next theorem shows that Schnyder labelings and Schnyder woods are, essentially, the same. A proof can be found in [40].

Theorem 1.1. *Let M^σ be a suspended 3-connected planar map. The correspondence indicated in Figure 1.2 is a bijection between the Schnyder labelings and Schnyder woods of M^σ.*

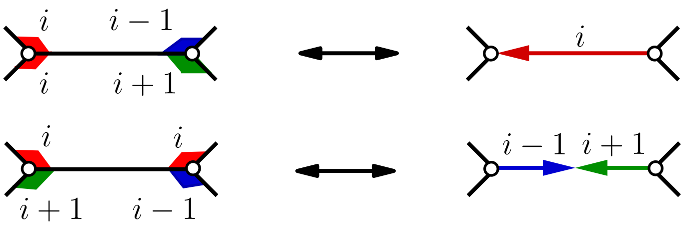

Figure 1.2. The correspondence between angle labels at an edge and the colored orientation of the edge.

Henceforth, when working with a Schnyder wood or a Schnyder labeling we may be sloppy and refer to properties of the corresponding other structure. We will also refer to the Schnyder wood of a planar map without choosing the special vertices explicitly.

Let M^σ be a planar map with a Schnyder wood. Let T_i denote the digraph induced by the directed edges of color i. Every inner vertex has out-degree 1 in T_i. Therefore, every vertex v is the starting vertex of a unique i-path $P_i(v)$ in T_i. The next lemma implies that each of the digraphs T_i is acyclic, and hence the $P_i(v)$ are simple paths. A proof can be found in [38] or [40].

Lemma 1.2. *Let M be a planar map with a Schnyder wood (T_1, T_2, T_3). Let T_i^{-1} be obtained by reversing all edges from T_i. Then the digraph*

$$D_i = T_i \cup T_{i-1}^{-1} \cup T_{i+1}^{-1}$$

is acyclic for $i = 1, 2, 3$.

By Rule (W3), every vertex has out-degree 1 in T_i. Disregarding the half-edge at a_i, this makes a_i the unique sink of T_i. Since T_i is acyclic and has $n-1$ edges we obtain the following statement.

Corollary 1.3. *T_i is a directed spanning tree rooted at a_i, for $i = 1, 2, 3$.*

The i-path $P_i(v)$ of a vertex v is the unique path in T_i from v to the root a_i. Lemma 1.2 implies that for $i \neq j$, the paths $P_i(v)$ and $P_j(v)$ have v as the only common vertex. Therefore, $P_1(v), P_2(v), P_3(v)$ divide the bounded faces of M^σ into $R_i(v)$ three regions $R_1(v)$, $R_2(v)$, and $R_3(v)$, where $R_i(v)$ denotes the region bounded by and including the two paths $P_{i-1}(v)$ and $P_{i+1}(v)$, see Figure 1.3.

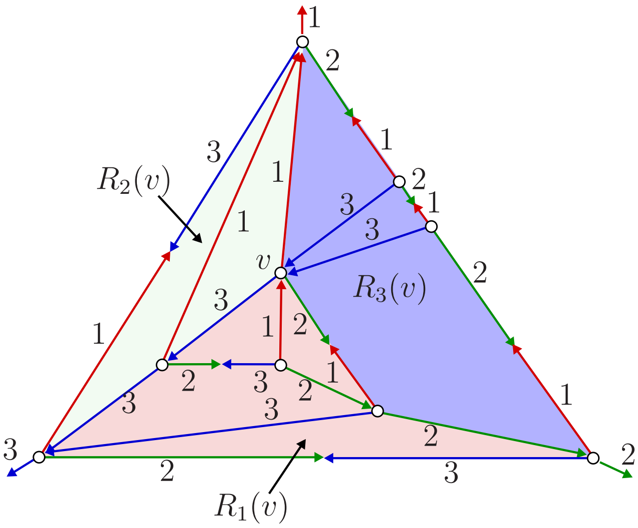

Figure 1.3. A Schnyder wood and the regions of the vertex v. The numbers indicate the edge colors.

Lemma 1.4. *If u and v are vertices with $u \in R_i(v)$, then $R_i(u) \subseteq R_i(v)$. The inclusion is proper if $u \in R_i(v) \setminus (P_{i-1}(v) \cup P_{i+1}(v))$.*

Lemma 1.5. *If the directed edge $e = (u, v)$ is colored i, then $R_i(u) \subset R_i(v)$, $R_{i-1}(u) \supseteq R_{i-1}(v)$ and $R_{i+1}(u) \supseteq R_{i+1}(v)$. At least one of the latter two inclusions is proper.*

We remark that the equalities $R_{i-1}(u) = R_{i-1}(v)$ and $R_{i+1}(u) = R_{i+1}(v)$ hold if and only if e is bidirected in colors $i, i+1$ respectively $i, i-1$. The above lemmas are crucial for the applications of the face-count vector (v_1, v_2, v_3) of a vertex v of *face-count vector* a Schnyder wood S. The *face-count vector* is defined as

$$v_i = \text{ the number of faces of } M^\sigma \text{ contained in region } R_i(v) \text{ with respect to } S.$$

The classic application of the face-count vector is in graph drawing. Let three non-collinear points α_1, α_2, and α_3 in the plane be given. These points and the region vectors can be used to define an embedding of M in the plane. A vertex v is mapped to the point

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

An edge $\{u, v\}$ is mapped by μ to the line segment connecting $\mu(u)$ and $\mu(v)$. A *grid drawing* of a planar graph is a crossing-free straight line embedding with integral vertex coordinates and convex faces. *grid drawing*

Theorem 1.6. *For every 3-connected planar map M the drawing $\mu(M)$ is a grid drawing.*

The face-count vectors cannot only be used to obtain 2-dimensional grid drawings of graphs, but also 3-dimensional orthogonal surfaces. We explain in the next section what orthogonal surfaces are and how they are related to Schnyder woods.

1.2 Schnyder Woods and Orthogonal Surfaces

Consider $\mathbb{R}^3$ equipped with the *dominance order*. In the dominance order we have that $u \leq v$ if and only if $u_i \leq v_i$ holds for each component i. We write $u \vee v$ to denote the *join*, i.e. the component-wise maximum, of $u, v \in \mathbb{R}^3$. Let $\mathcal{V} \subset \mathbb{R}^3$ be an antichain, that is, a set of pairwise incomparable elements. The *filter* generated by $\mathcal{V}$ in $\mathbb{R}^3$ is the set *dominance order* $u \vee v$ *join*

$$\langle \mathcal{V} \rangle = \{\alpha \in \mathbb{R}^3 \mid \alpha \geq v \text{ for some } v \in \mathcal{V}\}.$$

The boundary $\mathfrak{S}_\mathcal{V}$ of $\langle \mathcal{V} \rangle$ is the *orthogonal surface* generated by $\mathcal{V}$, Figure 1.4 shows an example. *orthogonal surface* $\mathfrak{S}_\mathcal{V}$

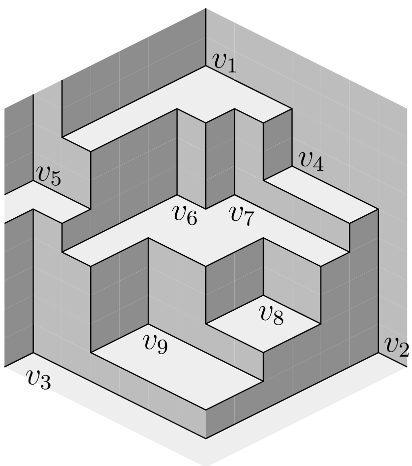

Figure 1.4. The orthogonal surface generated by $v_1 = (7, 0, 0)$, $v_2 = (0, 6, 0)$, $v_3 = (0, 0, 6)$, $v_4 = (5, 3, 0)$, $v_5 = (5, -1, 5)$, $v_6 = (4, 1, 2)$, $v_7 = (4, 2, 1)$, $v_8 = (2, 4, 2)$ and $v_9 = (1, 2, 4)$.

flat The *flats* of the surface are basically the connected regions of constant gray-value in our drawings of orthogonal surfaces. To make this precise, let H be the plane $x_i = \kappa$ and $\tilde{F}_1, \ldots, \tilde{F}_\ell$, the connected components of the interior of $H \cap \mathfrak{S}_\mathcal{V}$. The topological closures $F_1, \ldots, F_\ell$ of these components are the i-flats of $\mathfrak{S}_\mathcal{V}$ at $x_i = \kappa$, see Figure 1.5. The *i-flat* of $v \in \mathcal{V}$ is denoted by $F_i(v)$. In [44, 60] another definition of flats is given which captures interesting phenomena that appear in dimension 4 and higher.

i-flat $F_i(v)$

characteristic points We define the *characteristic points* of an orthogonal surface $\mathfrak{S}_\mathcal{V}$ as those points of $\mathfrak{S}_\mathcal{V}$ that are adjacent to an i-flat for every $i = 1, 2, 3$. We distinguish three different kinds of characteristic points. The first type are the local minima that generate the surface, see the point v in Figure 1.5. The second type are the local maxima of $\mathfrak{S}_\mathcal{V}$, see the point w in Figure 1.5. All other characteristic points are of the third type and we call them the *edge-points* of $\mathfrak{S}_\mathcal{V}$, see the point v_e in Figure 1.5. The name edge-point will be justified, since Theorem 1.7 shows that they are in bijection with the edges of a geodesically embedded map $M \hookrightarrow \mathfrak{S}_\mathcal{V}$. See also property (G2) of a geodesic embedding and Figure 1.7 (b). One can also think of the edge-points as those characteristic points, that can be obtained as the join $u \vee v$ of two minima $u, v \in \mathcal{V}$ of $\mathfrak{S}_\mathcal{V}$.

edge-point

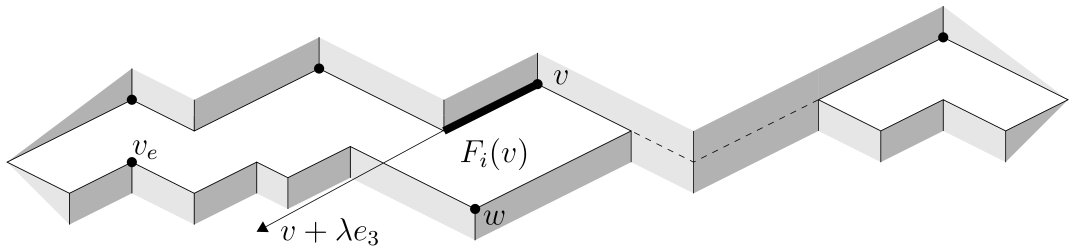

Figure 1.5. Two i-flats with the same i-coordinate.

If $u, v \in \mathcal{V} \subset \mathfrak{S}_\mathcal{V}$ and $u \vee v \in \mathfrak{S}_\mathcal{V}$, then $\mathfrak{S}_\mathcal{V}$ contains the union of the two line segments joining u and v to $u \vee v$. We refer to such arcs as *elbow geodesics* of $\mathfrak{S}_\mathcal{V}$. The *orthogonal arc* of $v \in \mathcal{V}$ in direction of the standard basis vector e_i is the part of the ray $v + \lambda e_i$, $\lambda \geq 0$ that lies on at least two flats of $\mathfrak{S}_\mathcal{V}$. In Figure 1.5 the part of the ray $v + \lambda e_3$ that forms the orthogonal arc is indicated by a bold line. Clearly every point $v \in \mathcal{V}$ has exactly three orthogonal arcs, one parallel to each coordinate axis. Some orthogonal arcs are unbounded while others are bounded, see Figure 1.4. Observe that $u \vee v$ shares two coordinates with at least one and possibly both of u and v, so every elbow geodesic contains at least one bounded orthogonal arc.

elbow geodesic

orthogonal arc

geodesic embedding Let M be a planar map. A drawing $M \hookrightarrow \mathfrak{S}_\mathcal{V}$ is a *geodesic embedding* of M into $\mathfrak{S}_\mathcal{V}$, if the following axioms are satisfied.

(G1) There is a bijection between $V(M)$ and $\mathcal{V}$.

(G2) There is a bijection between $E(M)$ and the edge-points of $\mathfrak{S}_\mathcal{V}$ and every edge is drawn as an elbow geodesic of $\mathfrak{S}_\mathcal{V}$.

(G3) There are no crossing edges in the embedding of M on $\mathfrak{S}_\mathcal{V}$.

An orthogonal surface $\mathfrak{S}_\mathcal{V} \subset \mathbb{R}^3$ is called *axial* if it contains exactly three unbounded orthogonal arcs. The example from Figure 1.4 is not axial. However, removing the point v_5 from the set $\mathcal{V}$ leads to an axial surface, see Figure 1.6 (a). These definitions have been proposed by Miller [72] who, essentially, also observed the following theorem. We give a proof sketch, since the connections between Schnyder woods and orthogonal surfaces are crucial for the understanding of Chapter 2. The complete proof can be found in [40]. *axial*

Theorem 1.7. *Let $\mathfrak{S}_\mathcal{V}$ be axial and $M \hookrightarrow \mathfrak{S}_\mathcal{V}$ be a geodesic embedding. Then the embedding induces a Schnyder wood of M^σ that is suspended at the unbounded orthogonal rays. Conversely, every Schnyder wood of a suspended map M^σ induces an axial geodesic embedding of M^σ.*

Proof sketch. Let $M \hookrightarrow \mathfrak{S}_\mathcal{V}$ be an axial geodesic embedding. The edges of M are colored with the direction of the orthogonal arc contained in the edge, that is arcs parallel to the x_i-axis are colored i. The orientation of an edge is chosen in accordance with the axis used to color the edge, Figure 1.6 shows an example. It can be verified that this rule for coloring and orienting edges yields a Schnyder wood on M^σ.

Conversely, given a Schnyder wood of M^σ, we embed every vertex v at its face-count vector $(v_1, v_2, v_3) \in \mathbb{N}^3 \subset \mathbb{R}^3$, that is

$$\mathcal{V} = \{(v_1, v_2, v_3) \mid v \text{ is a vertex of } M\}.$$

It can be verified that the canonical map $M \hookrightarrow \mathfrak{S}_\mathcal{V}$ is a geodesic embedding. The orthogonal surface in Figure 1.6 (a) can be constructed by this rule from the Schnyder wood in Figure 1.6 (b). □

The orthogonal surfaces constructed in the proof sketch for Theorem 1.7 have the additional property of being coplanar. A *coplanar orthogonal surface* is generated by vertices that all lie on a plane of the form $x_1+x_2+x_3 = c$, for some $c \in \mathbb{R}$. *coplanar orthogonal surface*

With an axial geodesic embedding $M^\sigma \hookrightarrow \mathfrak{S}_\mathcal{V}$ we can also associate a *Schnyder labeling*. Since every orthogonal arc leaving a vertex is occupied by an edge, every angle is completely contained in a flat. In the Schnyder labeling the angle φ at a vertex v is labeled i, if it is contained in $F_i(v)$. It is easy to verify (A1), (A2), and (A3) for this angle labeling. *Schnyder labeling*

Given Theorem 1.7 it is natural to ask whether there exists a geodesically embedded map on every orthogonal surface and whether this map is unique.

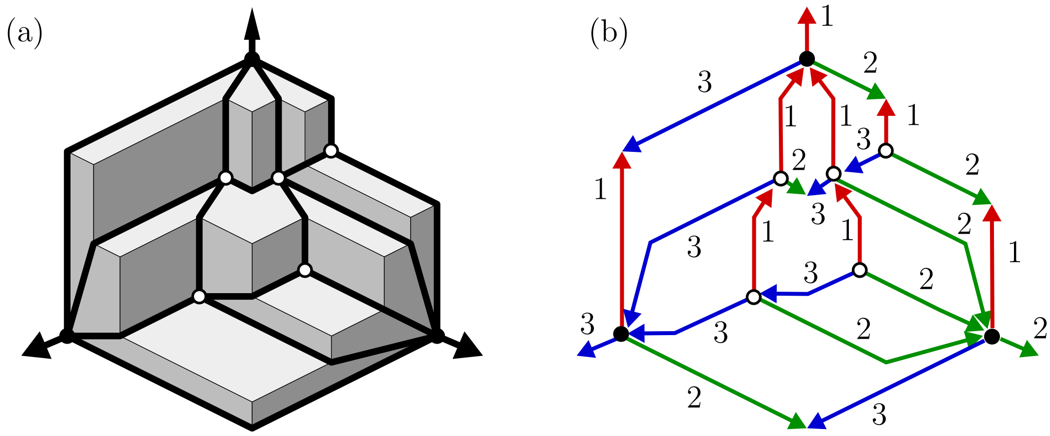

Figure 1.6. A geodesic embedding and the induced Schnyder wood. The numbers indicate edge colors.

As for the existence the answer is negative since a surface with three orthogonal arcs meeting in a single point does not support a Schnyder wood, see Figure 1.7 (a). *degenerate* We call a surface *degenerate* if such a pattern occurs. We omit the proof that every non-degenerate axial orthogonal surface $\mathfrak{S}_\mathcal{V}$ supports a Schnyder wood.

non-degenerate, axial From now on this thesis only deals with *non-degenerate* and *axial* orthogonal surfaces. For the sake of brevity we will usually omit these predicates.

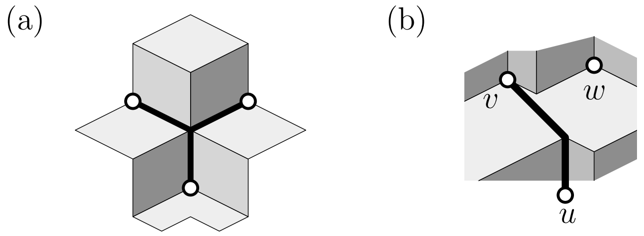

Figure 1.7. (a) A degenerate pattern, and (b) a non-rigid edge (u, v), where the edge-point $u \vee v$ dominates w.

As for the question about the uniqueness of geodesic embeddings, the answer is negative as well. For example in the situation shown in Figure 1.7 (b), the edge (u, v) can be replaced by the edge (u, w). Hence the surface supports two different graphs and also two different Schnyder woods. The reason for this ambiguity is a non-rigidity in the sense of the following definition. An elbow geodesic connecting *rigid elbow geodesic* vertices u and v is a *rigid elbow geodesic*, if u and v are the only vertices in $\mathcal{V}$ dominated by $u \vee v$. We will also call the edge of a geodesic embedding on $\mathfrak{S}_\mathcal{V}$ a *rigid edge* *rigid edge* if the corresponding elbow geodesic is rigid. An orthogonal surface $\mathfrak{S}_\mathcal{V}$ *rigid orthogonal surface* is a *rigid orthogonal surface* if all its elbow geodesics are rigid.

1.3 Schnyder Woods and Orientations with Prescribed Out-Degrees

The purpose of this section is to show how *orientations with prescribed out-degrees* can be used to encode Schnyder woods. Given a planar map M, and a function $\alpha : V \to \mathbb{N}$, an edge orientation X of M is called an *α-orientation* if for all $v \in V$ exactly $\alpha(v)$ edges are directed away from v in X. We call α an *out-degree function*. For the sake of brevity we will also use the term α-orientation to refer to orientations with prescribed out-degrees in general.

orientation with prescribed out-degrees
α-orientation
out-degree function

In a Schnyder wood on a triangulation only the three outer edges are bidirected. The reason is that the three spanning trees have to cover all $3n-6$ edges of the triangulation and the edges of the outer triangle must be bidirected because of Rule (W3). Theorem 1.8 says that not all edge colors are needed to encode a Schnyder wood. The edge orientations together with the colors of the half-edges at the special vertices are sufficient, the other edge colors can be deduced. For a proof see [31].

Theorem 1.8. *Let T be a plane triangulation, with vertices a_1, a_2, a_3 occuring in clockwise order on the outer face. Let $\alpha_T(v) := 3$ if v is an inner vertex and $\alpha_T(a_i) := 0$ for $i = 1, 2, 3$. Then there is a bijection between the Schnyder woods of T and the α_T-orientations of the inner edges of T.*

In the sequel we refer to an α_T-orientation of a triangulation T simply as a *3-orientation*. Schnyder woods on 3-connected planar maps are in general not uniquely determined by the edge orientations, see Figure 1.8.

3-orientation

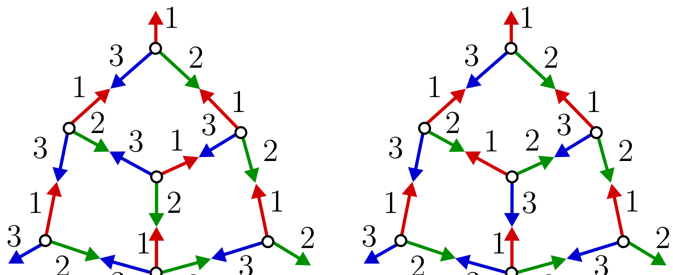

Figure 1.8. Two different Schnyder woods with the same underlying orientation. The numbers indicate edge colors.

Nevertheless Felsner describes a bijection between the Schnyder woods of a 3-connected planar map M and certain α-orientations on a related map $\widetilde{M}$, in [41]. In order to explain this bijection precisely, we first define the *suspension dual M^{σ^*}* of M^σ which is obtained from the dual M^* of M as follows, see Figure 1.9. Replace the vertex v^*_∞ that represents the unbounded face of M in M^* by a triangle on

suspension dual M^{σ^*}

three new vertices b_1, b_2, b_3. Let P_i be the path from a_{i-1} to a_{i+1} on the outer face of M that avoids a_i. In M^{σ^*} the edges dual to those on P_i are incident to b_i instead of v_∞^*. Adding a ray to each of the b_i yields M^{σ^*}. It is also nicely visible in Figure 1.9 that the maxima of an orthogonal surface are in bijection with the bounded faces of a geodesically embedded map, see also Proposition 2.4. The next proposition explains how the Schnyder woods on a graph and its suspension dual are related.

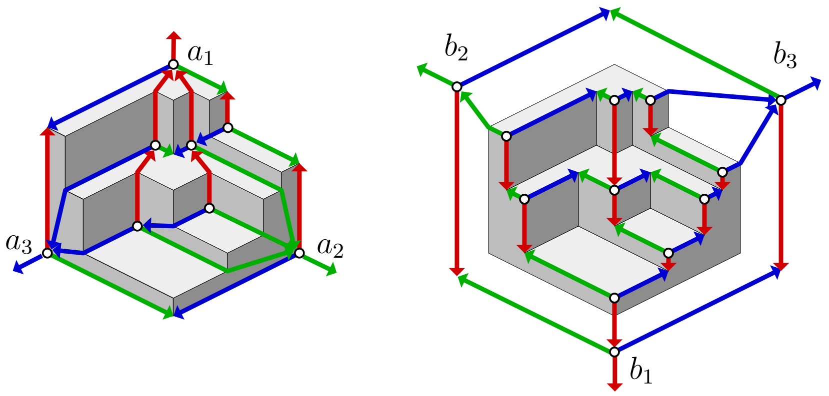

Figure 1.9. The Schnyder wood on the suspension dual of the map in Figure 1.6.

Proposition 1.9. *Let M^σ be a suspended planar map. There is a bijection between the Schnyder woods of M^σ and the Schnyder woods of the suspension dual M^{σ^*}.*

Schnyder labeling

Proof sketch. We work with *Schnyder labelings* instead of Schnyder woods. The inner angles of M^{σ^*} are in bijection with the angles of M^σ, see Figure 1.11. The inner angles of M^{σ^*} receive the same color as their counterparts in M^σ and for the outer angles the colors are prescribed by (A1). Using that (A2) holds for M^σ, it is easy to see that (A3) holds for the labeling of M^{σ^*}. Similarly (A3) for M^σ implies (A2) for M^{σ^*}. □

Figure 1.10 illustrates how the coloring and orientation of a pair of a primal and a dual edge are related.

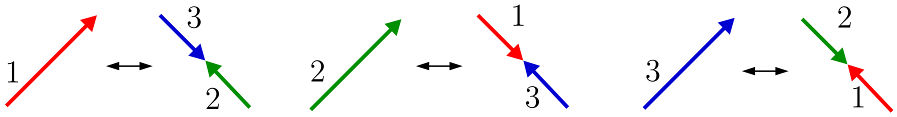

Figure 1.10. The three possible oriented colorings of a pair of a primal and a dual edge. The numbers indicate edge colors.

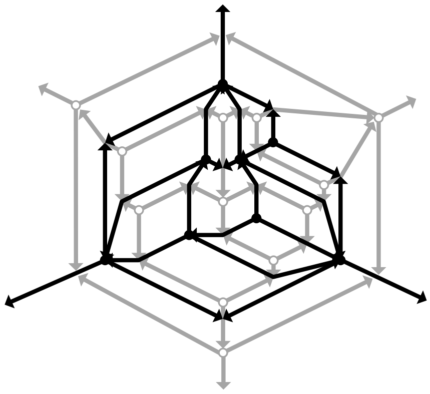

Figure 1.11. The primal-dual completion of the Schnyder woods shown in Figures 1.6 and 1.9.

The completion $\widetilde{M}'$ of M^σ and M^{σ^*} is obtained by superimposing the two graphs such that exactly the primal dual pairs of edges cross, see Figure 1.11.

In the *primal dual completion map* $\widetilde{M}$ the common subdivision of each crossing pair of edges of $\widetilde{M}'$ is replaced by a new *edge-vertex*. The rays emanating from the three special vertices of M^σ cross the three edges of the triangle induced by b_1, b_2, b_3 and thus produce edge-vertices. Note that all edge-vertices but these three correspond to edge-points of the orthogonal surface $\mathfrak{S}$ with $M^\sigma \hookrightarrow \mathfrak{S}$. The six rays emanating into the unbounded face of $\widetilde{M}'$ end at a new vertex v_∞ of $\widetilde{M}$ that is placed in this unbounded face. A pair of corresponding Schnyder woods on M^σ and M^{σ^*} induces an orientation of $\widetilde{M}$. We call this orientation an α_S*-orientation* and α_S can be defined as follows.

primal dual completion map $\widetilde{M}$

edge-vertex

α_S*-orientation*

$$\alpha_S(v) = \begin{cases} 3 & \text{for primal and dual vertices} \\ 1 & \text{for edge-vertices} \\ 0 & \text{for } v_\infty \end{cases}$$

For the sake of simpler notation we write α_S although the out-degree function depends on $\widetilde{M}$. Note that a pair of a primal and a dual edge always consists of a unidirected and a bidirected edge, as shown in Figure 1.10. This explains why $\alpha_S(v_e) = 1$ is the right choice for an edge vertex v_e. Theorem 1.10 says that the choice of the special vertices and the edge orientations of $\widetilde{M}$ are sufficient to encode a Schnyder wood of M^σ. For a proof see [41].

Theorem 1.10. *The Schnyder woods of a suspended planar map M^σ are in bijection with the α_S-orientations of $\widetilde{M}$.*

1.4 Edge Splits and Edge Merges on Schnyder Woods

In Section 2.1 we will need another tool from the theory of Schnyder woods, the edge split. In this section we introduce the edge split and the reverse operation, the edge merge. We start with a lemma from [13] about the generic appearance of a face in a Schnyder wood. This lemma will be used frequently in Chapter 2 and Section 3.2.

Lemma 1.11. *Given a Schnyder wood S, let F be an interior face. The edges on the boundary of F can be partitioned into six sets occurring in clockwise order around F. As illustrated in Figure 1.12, the sets are defined as follows (in case of bidirected edges the clockwise color is noted first).*

- *One edge from the set* {red-cw, blue-ccw, red-blue}
- *Any number (possibly 0) of edges* green-blue
- *One edge from the set* {green-cw, red-ccw, green-red}
- *Any number of edges* blue-red
- *One edge from the set* {blue-cw, green-ccw, blue-green}
- *Any number of edges* red-green

special edge Each of three edges from the first, third, and fifth set is called a *special edge* of the face F.

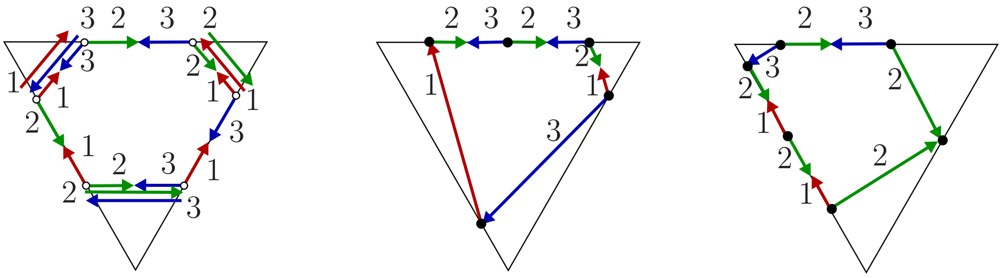

Figure 1.12. The generic appearance of a face as described by Lemma 1.11 and two concrete instances. The numbers indicate edge colors.

Proof sketch for Lemma 1.11. Recall part (A3) of the definition of *Schnyder*
Schnyder labeling *labelings*. Applying the rule depicted in Figure 1.2 for converting a Schnyder labeling into a Schnyder wood yields the claim of the lemma. □

edge split/merge We now introduce the operations *edge split* and *edge merge*. Given a Schnyder wood S, let e be a bidirected edge such that one of its directions is colored j,

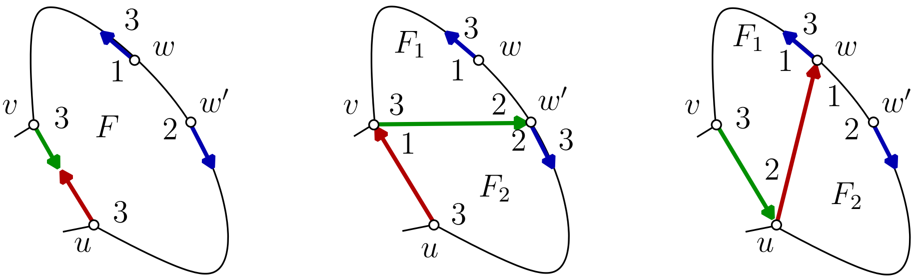

Figure 1.13. The two possible types of splits of a non-special bidirected red-green edge in F. The numbers indicate angle colors.

and let F be the incident face for that e is not special. Choose a vertex w of F such that the angle of w in F is labeled j. To *split e towards w* is to divide the bidirected edge e into two uni-directed copies and to move the head of the j-colored copy to connect to w. Figure 1.13 illustrates the operation. Note that $\{w, v\}$ and $\{w', u\}$ may be special edges of F. Furthermore, we observe that each direction of a bidirected edge e can be split into the face for which e is not special in at least one way, by part (A3) of the definition of Schnyder labelings. *split e towards w*

The reverse operation to an edge split is an edge merge. Given a Schnyder wood S, let e_1 and e_2 be two unidirected edges. Then, e_1 and e_2 form a *knee at v* if they form an angle at v and $e_1 = (u, v)$ is incoming at v while $e_2 = (v, w)$ is outgoing at v, see Figure 1.14. If the angle of e_1 and e_2 at v lies in the face F, then we also speak of the *knee at F*. To *merge edges* e_1, e_2 means to delete the edge e_2 from S and to make $\{u, v\}$ a bidirected edge. The direction (u, v) has the same color as e_1 and the direction (v, u) inherits the color of e_2. *knee at v, knee at F* *merge edges*

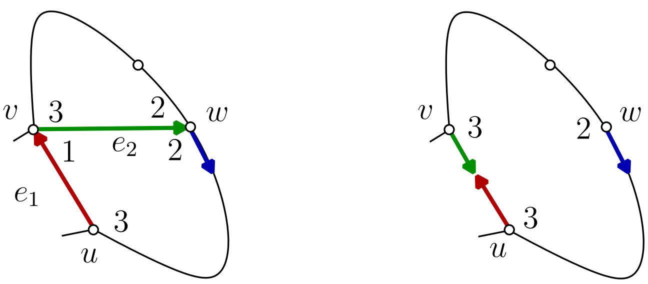

Figure 1.14. Merging the edges e_1 and e_2. The numbers indicate angle colors.

Lemma 1.12. *Let S be a Schnyder wood and S' obtained from S by an edge split or edge merge. Then S' is a Schnyder wood as well.*

Proof. Figures 1.13 and 1.14 show the *Schnyder labelings*. It is obvious that the labels at the angles of u, v, w, w', F_1, and F_2 obey (A2) respectively (A3). ☐ *Schnyder labeling*

In the rest of this section we will not distinguish between a Schnyder wood and the graph on which it is defined, but regard a Schnyder wood as a graph endowed with a colored orientation. Let $\mathcal{S}(n)$ denote the set of all Schnyder woods on n vertices with a triangular outer face. We define the *split-merge transition graph* on $\mathcal{S}(n)$ by saying that two Schnyder woods are adjacent if they can be obtained from each other by a single split or merge operation. We denote the transition graph by $\mathcal{S}(n)$ as well. That the transition graph is connected can be seen as follows. In [14] Bonichon et al. introduce colored diagonal flips of edges for Schnyder woods on triangulations. In the spirit of Wagner's Theorem [97] they show that the transition graph of these colored flips on the set of all Schnyder woods of triangulations with n vertices is connected. Since a colored diagonal flip is an edge merge followed by an edge split, this implies that the triangulations with n vertices are all in the same connected component of $\mathcal{S}(n)$. Every Schnyder wood that is not a triangulation has a bidirected inner edge and such an edge is splittable. Thus, every graph is connected to some triangulation in $S(n)$ and therefore $\mathcal{S}(n)$ is connected.

$\mathcal{S}(n)$

split-merge transition graph

$D(S)$

Let $S \in \mathcal{S}(n)$ be a Schnyder Wood. By $D(S)$ we denote the degree of S in the transition graph $\mathcal{S}(n)$. We now present a few results about the minimum and maximum degree of $\mathcal{S}(n)$. These results will not be needed in the rest of the thesis. Nevertheless we include them, since we think that the transition graph is interesting in its own right. For example it would be useful to obtain a random sampler for Schnyder woods using $\mathcal{S}(n)$.

Proposition 1.13. *Let $S \in \mathcal{S}(n)$ be a Schnyder wood and S^* be the suspension dual of S. Then,*

$$D(S) + D(S^*) \geq 2(m(S) - 3) \geq 3n(S) - 6.$$

Proof. Every edge e that does not lie on the boundary of the outer face is bidirected and thus splittable in either S or S^*. Every splittable edge contributes at least two to the split-merge degree. Since S is 3-connected we have $2m(S) \geq 3n(S)$, and the result follows. $\square$

Proposition 1.14. *For $S \in \mathcal{S}(n)$ we have that $D(S) \geq f(S) - 4$.*

Proof. Let F_0 be a face of S that is not incident to an outer edge of S. We prove the statement by showing that $D(S)$ suffices to distribute a charge of 1 to every such face. We first treat the case that $|F_0| \geq 4$. Then there is at least one edge e_0 that can be split into F_0. We charge this face with 1.

Now, we treat the case $|F_0| = 3$. If all the edges on the boundary of F_0 are undirected, then at least one of the three angles is a knee. We charge 1 to this face for the possible merge operation. If there is a bidirected edge e_0, then it splits into the face $F_1 \neq F_0$ on the boundary of that it lies, since F_1 is not the outer face. Since at most 1 has been charged from e_0 to F_1, we can charge 1 to F_0 as well, because there are at least two possible splits of e_0 into F_1. Thus, all faces not incident with an outer edge can be charged with 1 and the total charge does not exceed $D(S)$. □

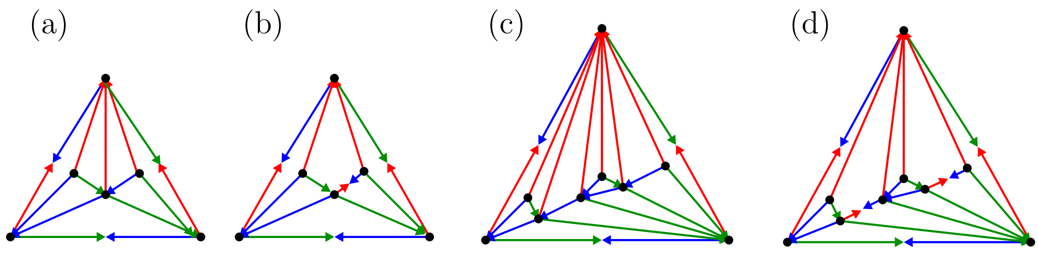

Figure 1.15. Schnyder woods with degree $f-4$ in $\mathcal{S}(n)$

Proposition 1.14 is tight, as the examples in Figure 1.15 show. The graphs in Figures 1.15 (a) and (c) are stacked triangulations. We now define this family of triangulations which we encounter several times throughout the thesis, see Chapter 5, Section 1.5, Proposition 3.26, and Section 4.5. Stacked triangulations can often be easily handled due to the inductive structure which is exhibited in the next definition.

- K_3 is a *stacked triangulation.* *stacked triangulation*
- Let $T=(V,E)$ be a stacked triangulation and $\{u,v,w\}$ a bounded face of T. Then, for a vertex $v' \notin V$,

$$T' = \Big(V \cup \{v'\}, E \cup \Big\{\{v'u\}, \{v'v\}, \{v'w\}\Big\}\Big)$$

  is a stacked triangulation.

The *height* of the outer vertices of a stacked triangulation is defined as -1. For an inner vertex v' stacked into a triangle $\{u,v,w\}$ we define its height as $h(v) = \max\{h(u), h(v), h(w)\}+1$. Similarly, the height of a face $F=\{u,v,w\}$ is $h(F) = \max\{h(u), h(v), h(w)\}+1$. *height*

The Schnyder woods depicted in Figures 1.15 (a) and (c) are stacked triangulations. Since every stacked triangulation has a unique Schnyder wood, see Section 1.5, this Schnyder wood has no cycles, see Theorem 1.8. Therefore every Schnyder wood of a stacked triangulation has only one knee per triangle, and the proposition is tight for all stacked triangulations.

The Schnyder woods from Figures 1.15 (b) and (d) are obtained from those in Figures 1.15 (a) respectively (c) by edge merges. These Schnyder woods allow for two edge splits into every quadrangular face. But the triangles for which the bidirected edges are special have no knee, and thus the proposition is tight for these graphs as well. The figure suggests how an infinite family of such non-triangular examples can be obtained.

Proposition 1.15. *For $S \in \mathcal{S}(n)$ we have that $D(S) \geq 4n(S)/3 - 6$.*

Proof. If $f(S) \geq 4n(S)/3 - 1$, the claim follows from Proposition 1.14. If $f(S) < 4n(S)/3 - 1$, then there are at least $2n(S)/3 - 3$ bidirected inner edges, and thus at least $4n(S)/3 - 6$ splits are possible. The lower bound for the number of bidirected edges can be obtained as follows. The inner vertices have in total exactly $3n - 9$ outgoing edges. Furthermore, $f(S) < 4n(S)/3 - 1$ implies that the number of inner edges is $m - 3 = n + f - 5 < 7n/3 - 6$. Together this implies that at least $2n/3 - 3$ edges are bidirected. □

We would like to point out that for triangulations Proposition 1.14 implies

$$D(S) \geq f(S) - 4 \geq 2(n(S) - 4).$$

In contrast to that the family of examples that we present now shows that the factor of 4/3 in Proposition 1.15 is best possible. This family was found by Stefan Felsner.

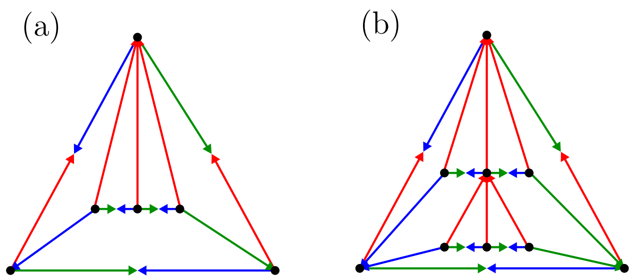

Figure 1.16. Schnyder woods with degree $4n/3 - 4$ in $\mathcal{S}(n)$.

We define an infinite family of Schnyder woods S_k. The first Schnyder wood of the family is S_1 as shown in Figure 1.16 (a), and S_2 is shown in Figure 1.16 (b). In general we denote by S_k the graph of this family with k levels of three vertices connected by green-blue edges. Since S_k admits no merges, two splits into every 4-face and four splits into the 5-face, we obtain

$$\lim_{k\to\infty} \frac{D(S_k)}{n(S_k)} = \frac{4}{3}.$$

Now we consider upper bounds for $D(S)$ for $S \in \mathcal{S}(n)$. We define an infinite family of Schnyder woods S'_k. The first Schnyder wood of the family is S'_1 as shown in Figure 1.17 (a), and S'_2 is shown in Figure 1.17 (b). In general we denote by S'_k the graph of this family in which the big central face has cardinality $3k+3$. Then, S'_k has $n(S'_k) = 3k+6$ vertices and $D(S'_k) \geq 6k^2$, since every non-special edge of the central face can be split in $2k$ different ways. Thus, $D(S)$ can be of order $\Omega(n^2)$ for $S \in \mathcal{S}(n)$.

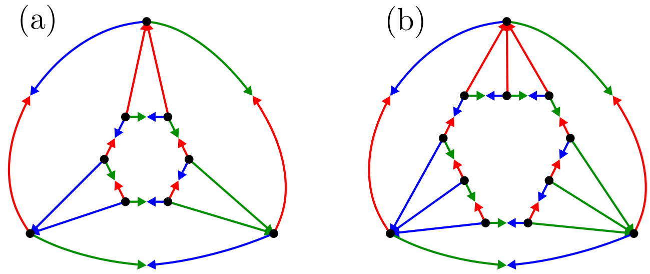

Figure 1.17. Schnyder woods with degree $6k^2 + 6$ in $\mathcal{S}(3k+6)$.

We now give a more restricted definition of an edge split which we call a *short split*. An edge e can now only split towards two vertices of the face F for that it is not special. A direction with color i of an edge e can split towards the first angle of color i that is encountered when walking around the face in this direction. It can be observed using Lemma 1.11 that the short splits suffice to generate all splits. We denote the transition graph produced by the merges and the short splits by $\mathcal{S}'(n)$ and the degree of $S \in \mathcal{S}'(n)$ by $D'(S)$. With this definition the degree of S'_k is $D(S'_k) = 6k+6$. Note that Propositions 1.13, 1.14, and 1.15 all hold for this more restricted notion of an edge split.

short split

$\mathcal{S}'(n)$, $D'(S)$

Proposition 1.16. *For $S \in \mathcal{S}'(n)$, we have that $D'(S) \leq 6n$.*

Proof. For $v \in V(S)$ let

$$D'(v) = (\text{number of knees at } v) + (\text{number of bidirected edges incident to } v).$$

Then we have

$$D'(S) = \sum_{v \in V(S)} D'(v),$$

and it is easy to see that $D'(v) \leq 6$. □

The family of *augmented triangular grids* $T^*_{k,\ell}$, see Figure 1.18, shows that Proposition 1.16 is essentially tight. A precise definition of this graph family is given in Section 3.1.2.

triangular grid

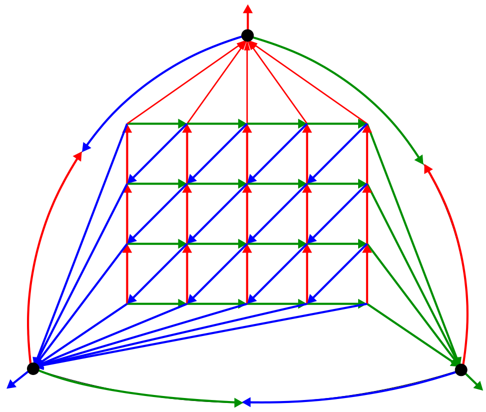

Figure 1.18. The triangular grid $T^*_{4,5}$ with a canonical Schnyder wood.

All vertices of $T^*_{k,\ell}$ that are not adjacent to an outer vertex have $D'(v) = 6$. Thus, the canonical Schnyder woods on $T^*_{k,k}$ satisfy

$$\lim_{k\to\infty} \frac{D'(T^*_{k,k})}{n(T^*_{k,k})} = 6.$$

1.5 Planar Maps with a Unique Schnyder Wood

The purpose of this section is to prove a constructive characterization of all 3-connected planar maps that have a unique Schnyder wood.

stacked triangulation

It is a well-known fact that the *stacked triangulations* are exactly the plane triangulations that have a unique Schnyder wood. We include the proof since it exemplifies our approach for proving Theorem 1.18.

Proposition 1.17. *A triangulation has a unique Schnyder wood if and only if it is a stacked triangulation.*

Proof. The proof uses the bijection between Schnyder woods and 3-orientations, see Theorem 1.8. It also uses that a triangulation has a unique 3-orientation if and only if it has an acyclic 3-orientation, see Theorem 3.1. Clearly, K_3 has a unique 3-orientations. Since every stacked triangulation T has a degree 3 vertex, it is easy to prove by induction that a 3-orientation of T must be acyclic, i.e. unique.

It remains to show that a 3-orientation of a non-stacked triangulation T is not acyclic. We remove degree 3 vertices from T until a triangulation T' of minimum degree 4 remains. Then, every vertex of T' has at least one incoming edge. This implies that there is an infinite sequence v_i, $i \in \mathbb{N}$ of inner vertices of T' such that (v_{i+1}, v_i) is a directed edge of T'. Since T' is finite, there must be a vertex repetition in this sequence, and this yields a directed cycle. $\square$

Theorem 1.18. *All 3-connected planar maps with a unique Schnyder wood can be constructed from the unique Schnyder wood on the triangle by the six operations shown in Figure 1.19 read from left to right.*

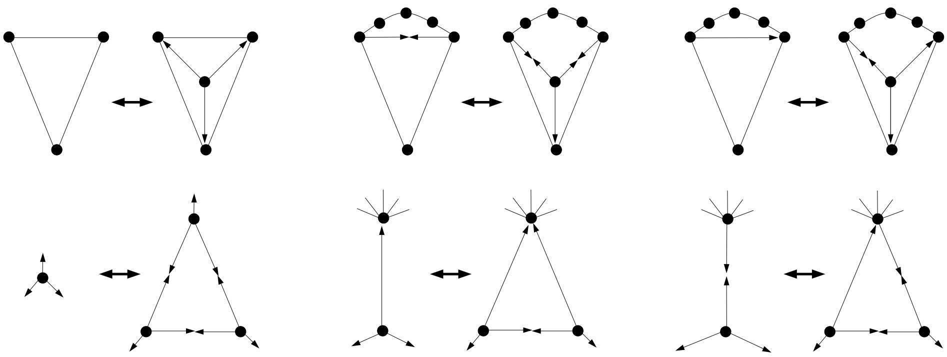

Figure 1.19. Every graph with a unique Schnyder wood can be constructed using the three primal operations in the first row and their duals in the second row.

Proof of Theorem 1.18. The proof uses the bijection with α_S-*orientations* from Theorem 1.10 and the *primal dual completion map* $\widetilde{M}$ of M^σ. One advantage of this approach is that we only have to consider the three cases shown in Figure 1.20. In this figure, the square vertices represent edge-vertices and the circular ones represent the vertices of the primal map M^σ respectively the dual map M^{σ^*}. *α_S-orientation* *primal dual completion map*

In each of the three parts of Figure 1.20 two operations on M^σ are indicated. One operation is obtained by choosing the black circular vertices as those of M^σ and the white circular vertices as those of M^{σ^*}. The other operation is obtained by choosing the white circular vertices as the primal vertices and the black ones as the dual vertices. Since Figure 1.20 shows how $\widetilde{M}$ can be reduced, not how M^σ can be constructed, left and right are switched with respect to the depiction of the operations in 1.19. Figure 1.21 shows how the two leftmost operations from Figure 1.19 are related to the topmost operation shown in Figure 1.20.

Observe that for each of the three operations from Figure 1.20, the six vertices whose incidences change are marked by a red circle. Every pair of such vertices is joined by a directed path in the graph on the left if and only if it is joined by such a path in the graph on the right. Thus, these operations cannot introduce a directed cycle in either direction. If $\widetilde{M}'$ is a planar map obtained from $\widetilde{M}$ by one of the operations from Figure 1.20, then $\widetilde{M}'$ has a directed cycle if and only if $\widetilde{M}$ does.

It remains to show that for a planar map M^σ with a unique Schnyder wood one of the operations from Figure 1.20 read from left to right can be applied to $\widetilde{M}$. We claim that $\widetilde{M}$ must have an edge-vertex v_e such that all three incoming edges

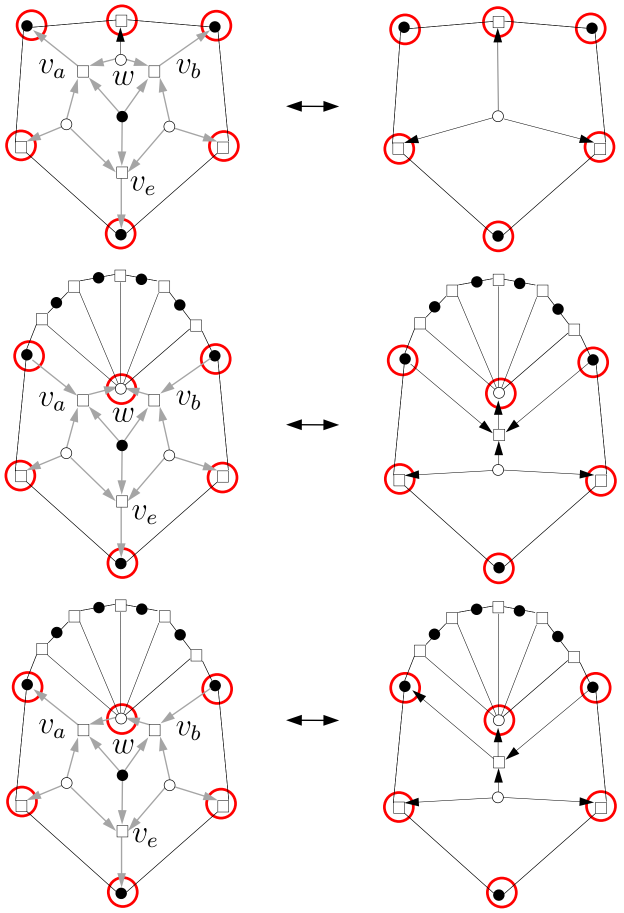

Figure 1.20. The three operations on $\widetilde{M}$ that can be used to reduce every planar map with a unique Schnyder wood to K_3.

at v_e start at a degree 3 vertex. Assume for the sake of contradiction that there is no such edge-vertex. Then we can construct an infinite sequence $(v_i)_{i \in \mathbb{N}}$ of vertices such that there is an edge from v_{i+1} to v_i for every i. Choose some edge-vertex to be v_1. By assumption, v_1 has one incoming edge that starts at a primal or dual vertex of degree at least 4 that we choose as v_2. As v_2 has degree at least 4 it has an incoming edge that starts at an edge-vertex. We choose this edge-vertex as v_3. This process can be continued infinitely, yielding the desired sequence. But as $\widetilde{M}$ is a finite graph, some vertex in $(v_i)_{i \in \mathbb{N}}$ has to be repeated and the first such repetition shows that there is a directed cycle in $\widetilde{M}$. Reversing the direction of all edges of this directed cycle yields another α_S-orientation. This orientation also corresponds to a Schnyder wood. This contradicts the assumption that M has a unique Schnyder wood. Thus, there must be an edge-vertex v_e such that all three incoming edges at v_e start at a degree 3 vertex.

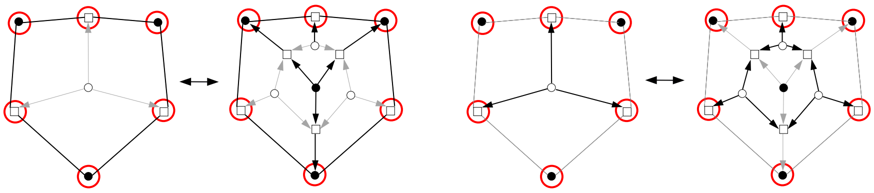

Figure 1.21. Translation of the operations from Figure 1.19 into the primal-dual completion map used in Figure 1.20.

We argue now that such an edge vertex v_e forces one of the three subgraphs highlighted by the gray edges in Figure 1.20. As indicated in Figure 1.20, let v_a and v_b be the two edge-vertices that are adjacent to two of the neighbors of v_e, and let w be the common neighbor of v_a and v_b. The three operations correspond to the cases where none, one, or both of the edges between v_a respectively v_b and w are directed towards w. Note that in the first case shown in Figure 1.20, the vertex w must have degree 3, otherwise there would be a directed 4-cycle. Hence, at least one of the three operations is applicable for every graph with a unique Schnyder wood that is not a triangle. This shows that every graph with a unique Schnyder wood can be reduced to a triangle with the six operations from Figure 1.19. □

Remark 1.19. The fact that the stacked triangulations are exactly the plane triangulations with a unique Schnyder wood is a special case of Theorem 1.18. The first operation shown in Figure 1.19 is exactly the operation used to construct stacked triangulations.

1.6 Conclusions

In this chapter we have introduced Schnyder woods. In Sections 1.1, 1.2, and 1.3 we have discussed the properties of Schnyder woods that we need in Chapter 2 and Section 3.2. In Section 1.5 we have given a constructive characterization of all planar maps with a unique Schnyder wood.

We have also introduced the operations edge split and edge merge in Section 1.4. We have studied the minimum and maximum degree of the split-merge transition graph $\mathcal{S}(n)$, and given examples that show that these bounds are essentially tight. We think that a better understanding of this transition graph and its structural properties could yield progress towards solving the following problem.

Problem 1.20. Can the transition graph $\mathcal{S}(n)$ be used to define a rapidly mixing Markov chain that yields a uniform random sampler for Schnyder woods?

Chapter 2

Schnyder Woods and Orthogonal Surfaces

In Chapter 1 we have introduced *Schnyder woods* and some of their properties. *Schnyder wood* In this chapter we use many of these properties, since we are concerned with the connections of Schnyder woods and *orthogonal surfaces*. *orthogonal surface*

This chapter is organized as follows. In Section 2.1 we are concerned with rigid orthogonal surfaces. We give a new proof of a theorem by Felsner [39] which says that every Schnyder wood can be geodesically embedded on a rigid orthogonal surface. Our proof uses the following very intuitive idea. The proof of Theorem 1.7 shows that every Schnyder wood can be geodesically embedded on some orthogonal surface. If this surface is not rigid, then we show that it can be transformed into a rigid surface by shifting some of its flats, as indicated in Figure 2.1. In Section 2.1 we also explain how the Brightwell-Trotter Theorem can be deduced from Felsner's result [39] and how our new proof can be used to design a simple algorithm that constructs a Brightwell-Trotter realizer for every 3-connected planar map.

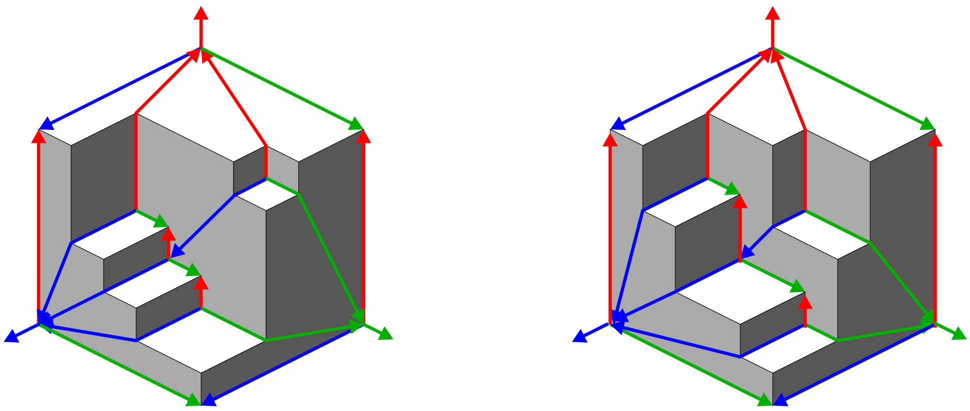

Figure 2.1. An orthogonal surface and an associated rigid surface

We recall from Section 1.2 that an orthogonal surface is coplanar if all its generating minima lie in the same plane. Section 2.2 is concerned with these coplanar surfaces. The interest in these surfaces originates from their close connection to planar straight line drawings that we have also pointed out after Theorem 1.7. Connecting the minima of a coplanar surface by straight line segments yields a plane and convex straight line drawing of the graph. We show in Section 2.2.1

that all coplanar surfaces supporting S can be obtained using Schnyder's original construction with appropriately weighted faces. An example of a Schnyder wood which has no supporting orthogonal surface that is simultaneously rigid and coplanar is the topic of Section 2.2.2.

In Section 2.3 we discuss height representations of orthogonal surfaces. These representations have a similar flavor as the face weight representations for coplanar surfaces, but they are not restricted to coplanar surfaces. For an orthogonal surface $\mathfrak{S}$, the height of a point $p \in \mathfrak{S}$ is the sum of its coordinates, that is $h(p) = p_1 + p_2 + p_3$ and the height-vector $h(\mathfrak{S})$ records the height of every minimum and maximum of $\mathfrak{S}$. We show in Section 2.3 that a Schnyder wood S supported by $\mathfrak{S}$ in conjunction with $h(\mathfrak{S})$ uniquely determines $\mathfrak{S}$.

2.1 Rigid Orthogonal Surfaces via Flat Shifting

In [82], Schnyder presented a characterization of planar graphs in terms of order dimension. We briefly introduce the terminology needed for the statement of this result. With a graph $G = (V, E)$ we associate an order P_G of height 2 on the set $V \cup E$. The order relation is defined by setting $x < e$ in P_G if $x \in V$, $e \in E$ and $x \in e$. The order P_G is called the *incidence order* of G.

incidence order

order dimension

The *order dimension* of an order P is the least k such that P admits an order preserving embedding in $\mathbb{R}^k$ equipped with the dominance order. We recall that in the *dominance order* we have that $u \leq v$ if and only if $u_i \leq v_i$ holds for each component i. For more on order dimension see [91, 92, 19, 38].

dominance order

Theorem 2.1 (Schnyder's Theorem)**.** *A graph is planar if and only if the dimension of its incidence order is at most 3.*

In the same paper Schnyder also shows that the incidence poset of vertices, edges, and faces of a planar triangulation has dimension 4, but the dimension drops to 3 upon removal of a face. Brightwell and Trotter [20] extended Schnyder's Theorem to the general case of embedded planar multigraphs. The main building block for the proof is the case of 3-connected planar graphs from [19].

Theorem 2.2 (Brightwell-Trotter Theorem)**.** *The incidence order of the vertices, edges, and faces of a 3-connected planar graph G has dimension 4. Moreover, if F is a face of G, then the incidence order of the vertices, edges, and all faces of G except F has dimension 3.*

Note that by Steinitz's Theorem, see Theorem 5.1 and [86, 87], the incidence poset of vertices, edges and faces of a 3-connected planar graph is the face lattice of a 3-polytope with **0** and **1** removed.

The original proof of Theorem 2.2 in [19] was long and technical and Felsner gave a simpler proof in [38]. Miller [72] observed that a *rigid orthogonal surface* induces a unique Schnyder wood and he proved the following two statements which also imply Theorem 2.2. Recall that an orthogonal surface is called rigid if it supports a unique graph, see Figure 1.7 (b). *rigid orthogonal surface*

Theorem 2.3. *Every suspended 3-connected planar map M^σ has a geodesic embedding $M^\sigma \hookrightarrow \mathfrak{S}$ on some rigid orthogonal surface $\mathfrak{S}$.*

Proposition 2.4. *Let $\mathfrak{S}_\mathcal{V}$ be a rigid orthogonal surface. Let $M^\sigma \hookrightarrow \mathfrak{S}_\mathcal{V}$ be a geodesic embedding and F a bounded face of M. If α_F is the join of the vertices of F, then $w \in F \Leftrightarrow w \leq \alpha_F$.* *join*

Note that α_F as defined above lies on $\mathfrak{S}_\mathcal{V}$ and is a maximum of the surface with respect to the *dominance order*. In fact, for any set $W \subseteq \mathcal{V}$ of vertices the join lies on $\mathfrak{S}_\mathcal{V}$ if and only if they all lie on a common face of M^σ. It is crucial here, that $\mathfrak{S}_\mathcal{V}$ is a rigid surface. If W contains a vertex from each of the three sides of the face F, as shown in Figure 1.12, then the join is a maximum of $\mathfrak{S}_\mathcal{V}$. *dominance order*

In [39] Felsner proved the following theorem thereby answering a question by Miller [72] and strengthening Theorem 2.3.

Theorem 2.5. *If S is a Schnyder wood of a map M^σ, then there is an axial rigid orthogonal surface $\mathfrak{S}$ and a geodesic embedding $M^\sigma \hookrightarrow \mathfrak{S}$ such that S is the unique Schnyder wood supported by $\mathfrak{S}$.*

In this section we present an intuitive proof of Theorem 2.5 which leads to a simple linear time algorithm for the computation of the rigid surface. The idea is to start with the orthogonal surface $\mathfrak{S}$ obtained from a Schnyder wood S by face-counting, see Theorem 1.7. If this surface is non-rigid it is possible to make some local adjustments at a non-rigid edge by moving some of the flats up or down in the direction of their normal vector, see Figure 2.1. The nontrivial point is to show that these adjustments can be combined in such a way that the whole surface becomes rigid.

Lemma 2.6 and Lemma 2.7 are part of our proof of Theorem 2.5. Let S be a Schnyder wood on a 3-connected planar map $M = (V, E)$ and let $\mathfrak{S}$ be the orthogonal surface obtained from S by face-counting, see Theorem 1.7. Let $\mathcal{F}_i$ be the set of *i-flats* of $\mathfrak{S}$. On the set $\mathcal{F}_i$ we define a relation Γ_i by three rules, Figure 2.2 shows an example. *i-flat*

(a) If (u, v) is an edge of color i, then $F_i(u) < F_i(v)$ in Γ_i.

(p) If (v, u) is unidirected in color $i-1$ or $i+1$, then $F_i(u) < F_i(v)$ in Γ_i.

(r) If (v, u) is unidirected in color $j \neq i$ and there is a $w \in V$ such that $F_j(w) = F_j(u)$ and $w_i > u_i$, then $F_i(v) < F_i(w)$ in Γ_i.

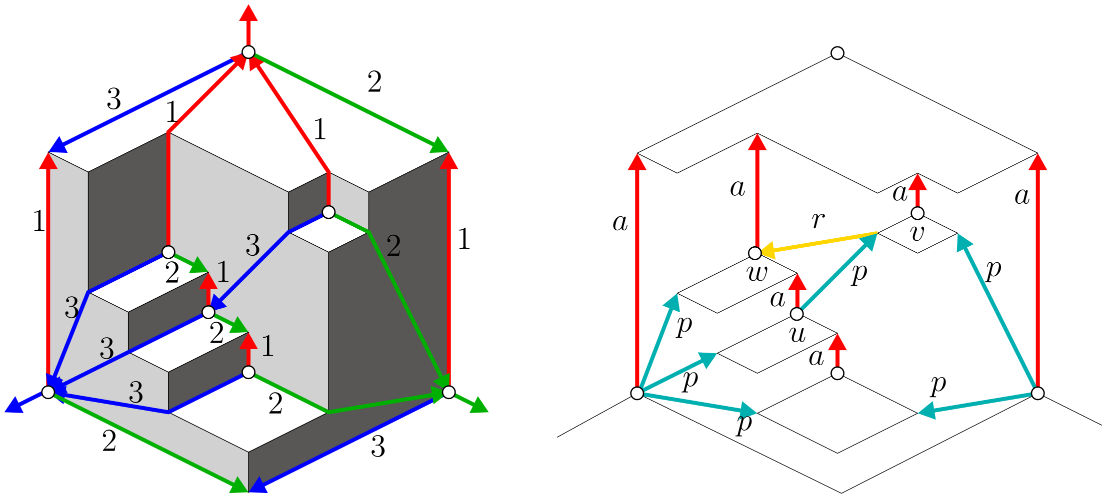

Figure 2.2. On the left a non-rigid surface with a Schnyder wood. On the right the corresponding relation Γ_1.

The pairs in Γ_i are classified as a-relations (*arc*), p-relations (*preserve*) and r-relations (*repel*). Lemma 2.6 is the heart of the proof of Theorem 2.5 as it justifies why the flat shifts (i.e. r-relations) can be combined to obtain a rigid surface.

Lemma 2.6. *The relation Γ_i defined on $\mathcal{F}_i$ is acyclic, for $i = 1, 2, 3$.*

Proof. By symmetry it is enough to prove the case $i = 1$. Recall that we identify the colors 1, 2, 3 with red, green, and blue in our figures. We identify the a- and p-relations with edges of the Schnyder wood S. We define a surjective map from the set of red edges in S to the set of a-relations by mapping an edge (u, v) to the relation $F(u) < F(v)$. Similarly, there is a surjective map from the blue and green unidirected edges in S to the p-relations (if (v, u) is such an edge, then $F(u) < F(v)$ is in Γ_1).

In order to deal with the r-relations we construct a Schnyder wood S' from S using *edge splits*, see Section 1.4. Let $e = (v, u) \in S$ be a unidirected blue edge and $F(u) < F(v)$ the corresponding p-relation. Let $F(u_k) > \ldots > F(u_1)$ be the set of flats that have an r-relation $F(v) < F(u_j)$ related to e. The order on this set comes from the red a-relations, since the edges $\{u, u_1\}$ and $\{u_{j-1}, u_j\}$ are bidirected in red and green in S. Construct S' by splitting the edges $\{u, u_1\}$, $\{u_1, u_2\}$, …, $\{u_{k-1}, u_k\}$ towards v. This is legal since the angle of v in the face in question has label 2 (green), see Lemma 1.12. We repeat this operation for other r-relations in Γ_1 that come from unidirected blue edges. A symmetric operation is used to introduce edges for all r-relations in Γ_1 that come from unidirected green edges in the Schnyder wood S.

edge split

In the Schnyder wood S' we associate an edge with every relation in Γ_1. With an a-relation $F(u) < F(v)$ we associate the red edge (u, v), and with a p-relation $F(u) < F(v)$ the blue or green edge (v, u). With an r-relation $F(u) < F(v)$ we associate the blue or green edge (v, u) that was introduced into S' by a split.

Now assume that C is a cycle in the relation Γ_1 on $\mathcal{F}_1$. The idea is to show that C induces a cycle C' in $T_1 \cup T_2^{-1} \cup T_3^{-1}$, where the T_i, $i \in \{1, 2, 3\}$, are the respective trees of S'. This yields a contradiction to Lemma 1.2.

The relations in C are associated with some edges in $T_1 \cup T_2^{-1} \cup T_3^{-1}$. However, consecutive relations $F(u) < F(v)$ and $F(u') < F(v')$ in C, i.e., $F(v) = F(u')$, may correspond to different vertices $v \neq u'$ from the flat $F(v)$. This yields gaps in the intended cycle C'. Note that the set of vertices lying on a common 1-flat is strongly connected in S' via bidirected green-blue edges. Therefore S' contains a path of bidirected green-blue edges connecting v and u', hence, the directed path required to close the gap in C' can be found in $T_1 \cup T_2^{-1} \cup T_3^{-1}$.

The contradiction shows that Γ_i is acyclic. $\square$

Let $\mathfrak{S}$ be the orthogonal surface supporting S that is generated by the face-count vectors, see Theorem 1.7. The transitive closure Γ_i^* of Γ_i is an order on $\mathcal{F}_i$ by Lemma 2.6. Let L_i be a linear extension of Γ_i^*. An i-flat F_i of $\mathfrak{S}$ is mapped to its position in L_i, more formally to

$$\alpha_{F_i} = |\{F_i' \in \mathcal{F}_i : F_i' < F_i \text{ in } L_i\}|.$$

With a vertex $v \in V$ we associate the point

$$v' = (\alpha_{F_1(v)}, \alpha_{F_2(v)}, \alpha_{F_3(v)}) \in \mathbb{R}^3.$$

To complete the proof of Theorem 2.5 it remains to verify that the orthogonal surface $\mathfrak{S}_{\mathcal{V}_\alpha}$ generated by $\mathcal{V}_\alpha = \{v' : v \in \mathcal{V}\}$ is rigid and supports the Schnyder wood S. The key for proving this is the following lemma.

Lemma 2.7. *If $R_i(u) = R_i(v)$, then $u_i' = v_i'$, and if $R_i(u) \subset R_i(v)$, then $u_i' < v_i'$.*

Proof. The first statement is immediate: From $R_i(u) = R_i(v)$ it follows that $F_i(u) = F_i(v)$ and hence $u_i' = v_i'$. For the proof of the second statement note that there exists an index $j \neq i$ such that $R_j(u) \supset R_j(v)$. Therefore, the j-path of v and the i-path of u have to cross in a vertex w. The edges of $P_i(u)$ imply that $F_i(w) > F_i(u)$ in Γ_i^* and hence in L_i. Let (x, y) be an edge of color j on $P_j(v)$. We distinguish the following three cases for the type of $e = \{x, y\}$.

(1) e is bidirected and the color of (y, x) is i.

(2) e is bidirected and the color of (y, x) is not i.

(3) e is unidirected.

In Case (1) we find the a-relation $F_i(y) < F_i(x)$ in Γ_i. In Case (2) vertices x and y are on the same i-flat, i.e. $F_i(y) = F_i(x)$. In Case (3) the relation $F_i(y) < F_i(x)$ is a p-relation in Γ_i. Transitivity yields $F_i(u) < F_i(w) < F_i(v)$ in Γ_i^* and hence in L_i. This implies $u'_i < v'_i$. □

Claim 1. $\mathcal{V}_\alpha$ is an antichain in $\mathbb{R}^3$.

Proof. This follows from Lemma 2.7 and the observation that for any two vertices $u, v \in S$, there are colors i and j with $R_i(u) \subset R_i(v)$ and $R_j(v) \subset R_j(u)$. △

Claim 2. $\mathfrak{S}_{\mathcal{V}_\alpha}$ is non-degenerate.

Proof. The linear extension L_i assigns different positions to different flats, therefore the situation from Figure 1.7 (a) cannot occur. △

Claim 3. $\mathfrak{S}_{\mathcal{V}_\alpha}$ supports the Schnyder wood S.

Proof. Let $e = \{u, v\}$ be an edge of S and $x \notin e$ a vertex. For some i the edge e is contained in region $R_i(x)$. This implies $R_i(u) \subseteq R_i(x)$ and $R_i(v) \subseteq R_i(x)$.

From Lemma 2.7 it follows that in the above setting $u'_i \leq x'_i$ and $v'_i \leq x'_i$. This shows that with $e = \{u, v\}$ the join $u' \vee v'$ and hence the elbow geodesic that connects u and v is on the surface $\mathfrak{S}_{\mathcal{V}_\alpha}$.

If the edge $e = (u, v)$ is directed in color i from u to v, then by Lemma 1.5 together with Lemma 2.7, we have $u'_i < v'_i$, $u'_{i+1} \geq v'_{i+1}$ and $u'_{i-1} \geq v'_{i-1}$. Therefore, the orthogonal arc of v in direction e_i is used by this edge. Since the orthogonal arcs of all vertices are already occupied by edges of S, there are no additional edges on $\mathfrak{S}_{\mathcal{V}_\alpha}$. △

Claim 4. $\mathfrak{S}_{\mathcal{V}_\alpha}$ is rigid.

Proof. Suppose not, then there is a unidirected edge (v, u) of color j and a vertex w such that, $w' \leq u' \vee v'$, and $F_j(u) = F_j(w)$. There is a bidirected path in colors i and k joining u and w. We may assume that $w \in P_i(u)$ and $u \in P_k(w)$. It follows that $R_i(w) \supset R_i(u)$, hence, $w_i > u_i$ and the relation $F_i(v) < F_i(w)$ is an r-relation in Γ_i. The unidirected edge (v, u) in color j induces the p-relation $F_i(u) < F_i(v)$ in Γ_i. Therefore, $u'_i < v'_i < w'_i$, in contradiction to $w' \leq u' \vee v'$. △

This completes the proof of Theorem 2.5. □

Next, we present a simple algorithm which, given a Schnyder wood S, computes a rigid orthogonal surface $\mathfrak{S}$ inducing S.

Corollary 2.8. *There is an $O(n)$ algorithm computing a rigid orthogonal surface for a given Schnyder wood S.*

Proof. We assume that S is given in the form of adjacency lists ordered clockwise around each vertex. With each edge in the adjacency list of a vertex v, the information about the coloring and orientation of that edge is given by its type relative to v. There are twelve such types, three outgoing types in each color (two of them for bidirected edges) and the unidirected incoming edges.

By symmetry it is sufficient to show how to obtain the first coordinate for all vertices of S in linear time. We produce a copy of the vertex set and build a digraph D_r on this copy. For every red edge there is an edge pointing in the same direction in D_r and for all blue and green unidirected edges there is an edge pointing in the opposite direction. We then check at each original vertex if its red outgoing edge is green in the reverse direction and if it has a unidirected blue incoming edge. If so, there is an edge from the start of the blue edge to the end of the red outgoing edge. This single repel edge is sufficient since other repel relations associated to the same unidirected blue edge are implied by transitivity. Analogously, check at each original vertex if its red outgoing edge is blue in the reverse direction and if it has a unidirected green incoming edge. If so, there is a repel edge from the start of the green edge to the end of the red outgoing edge in D_r. Finally, contract all blue-green edges from S in D_r, maintaining a pointer from the original vertices to their representatives in D_r.

Then, we compute a topological sorting of D_r and assign each vertex the topsort-number of its flat as first coordinate. All this can be done in $O(n)$ time. Three runs of this procedure, one for each coordinate are required. The correctness of the algorithm follows from Theorem 2.5. □

In the following theorem we assume that the *3-polytope* P is given as a *combinatorial 3-polytope*, that is as a planar map. For a *geometric 3-polytope* given by a set of points in $\mathbb{R}^3$, convex hull algorithms of the beneath-and-beyond type compute the combinatorial 3-polytope in $O(n \log n)$, see for example [76, 77].

3-polytope
combinatorial/geometric 3-polytope

Theorem 2.9. *Let P be a combinatorial 3-polytope with n vertices. Then, a Brightwell-Trotter realizer for P can be computed in $O(n)$ time.*

Proof. As shown by Fusy et al. in [52] a Schnyder wood S for P can be computed in $O(n)$ time. With little translational effort this also follows from algorithms for computing orderly spanning trees [23] or canonical orderings [25] which are based on Kant's algorithm [59, 24]. By Corollary 2.8 a rigid orthogonal surface that induces S can be computed in time $O(n)$. By Proposition 2.4 and Steinitz's Theorem, see Theorem 5.1 such an orthogonal surface yields a Brightwell-Trotter realizer of P. □

2.2 Coplanar Surfaces

coplanar surface Recall that an orthogonal surface is a *coplanar orthogonal surface* if there exists a constant $c \in \mathbb{R}$ such that every minimum v of the surface fulfills the equation $v_1+v_2+v_3 = c$. Schnyder's classic approach of drawing graphs using the *face-count* *face-count vector* *vectors* $\{(v_1, v_2, v_3) \mid v \in V\}$ yields a subclass of all coplanar surfaces, as described in the proof sketch of Theorem 1.7. Geodesic embeddings on coplanar surfaces have the pleasant property that the positions of the vertices in the plane yield a *grid drawing* *grid drawing* of the embedded graph.

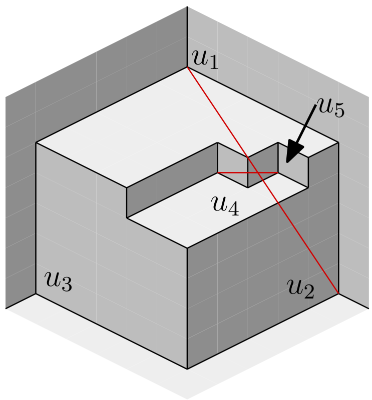

Figure 2.3. An orthogonal surface generated by $u_1 = (5,0,0)$, $u_2 = (0,5,0)$, $u_3 = (0,0,5)$, $u_4 = (4,3,2)$, and $u_5 = (4,4,1)$.

Similar approaches for non-coplanar surfaces, where the points are projected orthogonally to the plane $x + y + z = 1$, fail. This is because crossings between edges may be produced, see Figure 2.3 for an example. The coordinates of the orthogonally projected points are

$$\begin{aligned} u_1' &= (11,-4,-4)/3, & u_2' &= (-4,11,-4)/3, \\ u_4' &= (4,1,-2)/3, & u_5' &= (4,4,-5)/3. \end{aligned}$$

This implies that $8u_1'/15+7u_2'/15 = u_4'/3+2u_5'/3$, that is the straight line segments representing the edges u_1u_2 and u_4u_5 cross.

A representation of all coplanar surfaces using Schnyder woods and edge weights is given in Section 2.2.1. In Section 2.2.2 we present an example of a Schnyder wood that has no geodesic embedding on a surface that is rigid and simultaneously coplanar.

2.2.1 Coplanar Surfaces and Face Weights

We now generalize the classic approach of counting every bounded face with weight 1 (see Theorem 1.7) by allowing more general face weights. We then use coordinate vectors recording the sum of weights in the regions of a vertex. We show that this construction essentially yields all coplanar surfaces supporting a given Schnyder wood, and thus all non-degenerate coplanar surfaces can be obtained from some Schnyder wood in this way.

Theorem 2.10. *Let $\mathfrak{S}$ be a coplanar orthogonal surface generated by $\mathcal{V}$ supporting a Schnyder wood S on the vertex set $V(S) \equiv \mathcal{V}$. Then there is a unique weight function $w : F(S) \to \mathbb{R}$ on the set of bounded faces of S and a unique translation $t \in \mathbb{R}^3$ such that for all $v \in V(S)$ and $i \in \{1,2,3\}$ the coordinates are given by*

$$v_i = t_i + \sum_{F \in R_i(v)} w(F).$$

Remark 2.11. A Schnyder wood S and a weight function w define an orthogonal surface $\mathfrak{S}_{S,w}$. This surface, however, does not have to support the initial Schnyder wood. From the proof of Theorem 1.7 it follows that a necessary and sufficient condition for an embedding $S \hookrightarrow \mathfrak{S}_{S,w}$ is that

$$R_i(u) \subseteq R_i(v) \quad \Longrightarrow \quad \sum_{F \in R_i(u)} w(F) \leq \sum_{F \in R_i(v)} w(F)$$

with strict inequality whenever $R_i(u) \subset R_i(v)$.

Proof of Theorem 2.10. Let $\mathfrak{S}$ be a coplanar orthogonal surface and S a Schnyder wood induced by $\mathfrak{S}$. Note that $F_i(a_j) = F_i(a_k)$ for the suspension vertices a_1, a_2, a_3 of S, where $\{i,j,k\} = \{1,2,3\}$. Let t_i be the ith coordinate of a_j for $j \neq i$. Subtracting $t = (t_1, t_2, t_3)$ from all generating vectors $v \in \mathcal{V}$ yields a *normalized orthogonal surface* in the sense that the suspension vertices now have coordinates $(c,0,0), (0,c,0), (0,0,c)$ and $v_1 + v_2 + v_3 = c$ for all $v \in \mathcal{V}$. In the following we assume that $\mathfrak{S}$ is normalized in this sense. *normalized orthogonal surface*

Let f be the number of faces of S. With the region $R_i(v)$ of a vertex v we associate a row vector $r_i(v)$ of length $f-1$ with a component for each bounded face of F. The vector $r_i(v)$ is defined by

$$r_i(v)_F = \begin{cases} 1 & \text{if } F \in R_i(v) \\ 0 & \text{otherwise.} \end{cases}$$

The existence of a weight assignment to the faces realizing the normalized surface $\mathfrak{S}$ is equivalent to finding a vector $w \in \mathbb{R}^{f-1}$ such that

$$\forall v \in V, \forall i \in \{1,2,3\}: \quad r_i(v) \cdot w = v_i \tag{2.1}$$

Claim 1. The rank of the linear system (2.1) is at most $f-1$.

Proof. The vectors $r_i(v)$ have dimension $f-1$, and we show that the matrix with rows $(r_i(v), v_i)$ of length f has at most $f-1$ linearly independent rows. First suppose that S is the Schnyder wood of a triangulation. For the three special vertices, we only need the single equation $\mathbb{1} \cdot w = c$, where $\mathbb{1}$ denotes the vector whose components are all 1. This equation together with the three equations of an inner vertex v is a dependent system, since $\mathbb{1} = r_1(v) + r_2(v) + r_3(v)$ and $c = v_1 + v_2 + v_3$. Therefore, we have at most

$$1 + 2(n-3) = 2n - 5 = f - 1$$

linearly independent row vectors.

Let S be a Schnyder wood on a 3-connected planar map. If S has $k+3$ bidirected edges, then it has $f-1 = 2n-5-k$ bounded faces. If $e = \{v, w\}$ is a bidirected edge in colors $i-1$ and $i+1$, then $r_i(v) = r_i(w)$ and $v_i = w_i$. Therefore, among the $2n-5$ potentially independent vectors, there are at most $2n-5-k$ different ones. Hence, there are at most $f-1$ linearly independent row vectors. $\triangle$

We now show, that (2.1) has rank $f-1$ and therefore a unique solution. Let e_F be the $(f-1)$-dimensional row vector with a single one at the position corresponding to the bounded face F.

Claim 2. The vector e_F is in the linear span of the region-face incidence vectors $r_i(v)$, where $v \in V$ and $i \in \{1, 2, 3\}$.

Proof. Lemma 1.11 implies that it suffices to distinguish the following three cases.

Case 1. The boundary of F is a directed cycle C, where bidirected edges are allowed on C. From Lemma 1.11 or more directly from Rule (A3), see Section 1.1, it follows that the cycle C consists of three directed paths P_i in the three colors $i = 1, 2, 3$.

If C is clockwise, the order of the paths is P_1, P_2, P_3 and if C is counterclockwise the order of the paths is P_1, P_3, P_2, see Figure 2.4. Let v_i be the first vertex of path P_i. In the clockwise case consider the regions $R_2(v_1)$, $R_3(v_2)$ and $R_1(v_3)$, they are disjoint and cover the bounded area B except the face F. Hence

$$\mathbb{1} - (r_2(v_1) + r_3(v_2) + r_1(v_3)) = e_F.$$

In the counterclockwise case, the regions in question are $R_3(v_1)$, $R_1(v_2)$ and $R_2(v_3)$ and the equation is

$$\mathbb{1} - (r_3(v_1) + r_1(v_2) + r_2(v_3)) = e_F.$$

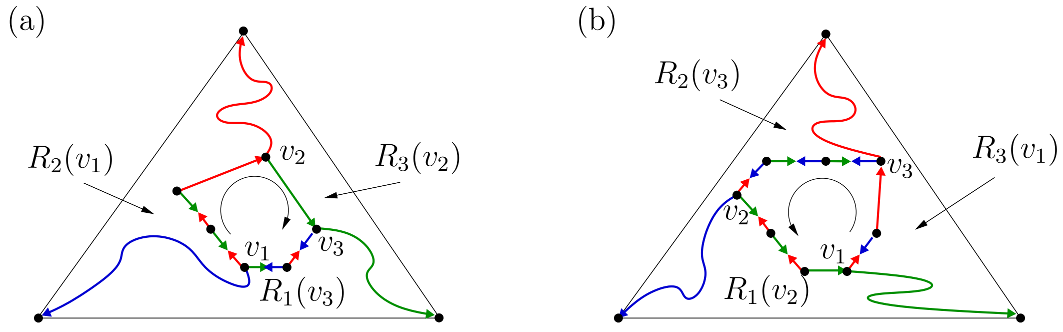

Figure 2.4. Faces with a directed cycle on the boundary.

Case 2. We assume that the boundary of F is not a directed cycle and that there are two unidirected special edges of the same color i.

We may assume that the three special edges e_1, e_2, e_3 have endvertices v_1, w_1, v_2, w_2, v_3, w_3 clockwise in this order on the boundary of F where $w_{j-1} = v_j$ is possible for every $j \in \{1, 2, 3\}$. By symmetry we may assume that $i = 1$ and $e_1 = (v_1, w_1)$, $e_2 = (w_2, v_2)$.

Subcase $\mathbf{w_1 = v_2}$. We first treat the case that e_3 is directed as (w_3, v_3). This includes the case that e_3 is bidirected. Figure 2.5 (a) shows the situation with $i = 1$.

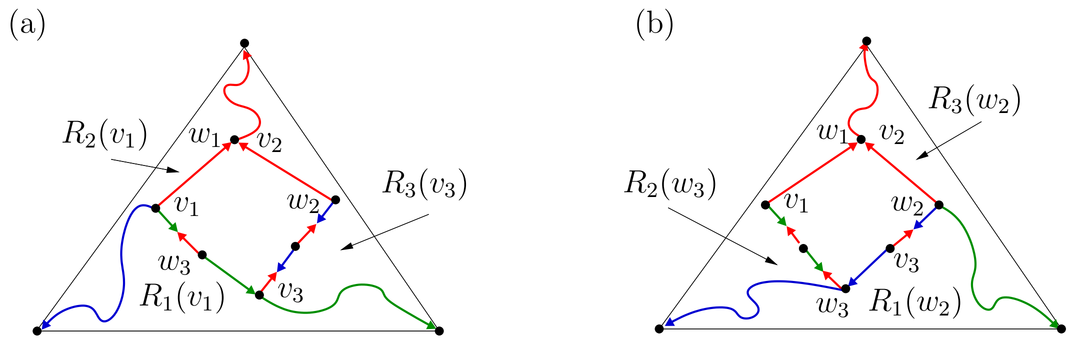

Figure 2.5. Faces without a directed cycle, and with unidirected edges of the same color.

As illustrated in the figure, $R_1(v_1)$, $R_2(v_1)$ and $R_3(v_3)$ partition $B \setminus F$, hence

$$\mathbb{1} - (r_1(v_1) + r_2(v_1) + r_3(v_3)) = e_F.$$

If e_3 is directed as (v_3, w_3), then, as shown in Figure 2.5 (b):

$$\mathbb{1} - (r_1(w_2) + r_3(w_2) + r_2(w_3)) = e_F.$$

Subcase $\mathbf{w_1 \neq v_2}$. In this case the boundary of F between w_1 and w_2 consists of edges bidirected in colors 2, 3. Let R be the region enclosed by this bidirected path, $P_1(w_1)$, and $P_1(v_2)$, and r the corresponding vector. As shown in Figure 2.6, $R_1(w_1)$, $R_2(w_1)$ and $R_3(v_2)$ partition $B \setminus R$, hence

$$\mathbb{1} - (r_1(w_1) + r_2(w_1) + r_3(v_2)) = r.$$

To represent the vector e_F we can now use the representations found in Subcase $w_1 = v_2$, we only have to add r on the right side.

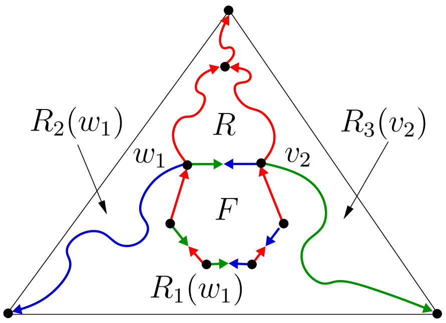

Figure 2.6. The region R in the case $i = 1$.

Case 3. We assume that the boundary of F is not a directed cycle and that there are no two unidirected special edges of the same color. Then, there are two unidirected special edges of different colors $i-1$, $i+1$, and the third special edge is bidirected in colors $i-1$, $i+1$. We assume that the two unidirected special edges are $e_1 = (v_1, w_1)$, $e_2 = (w_2, v_2)$ and that $i = 1$.

Subcase $\mathbf{w_1 = v_2}$. Figure 2.7 (a) shows the situation with $i = 1$.

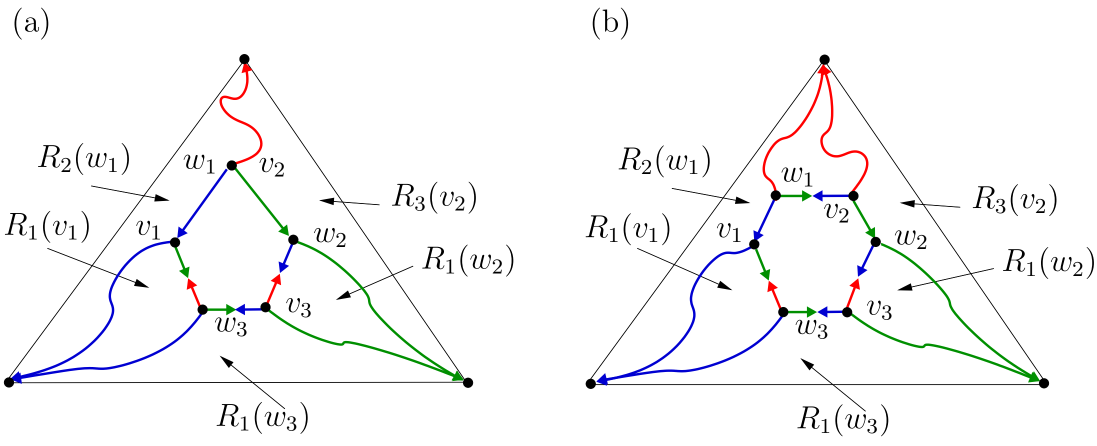

Figure 2.7. Faces without a directed cycle, and with unidirected edges of different colors.

As illustrated in the figure $R_3(v_2)$, $R_1(w_2)$, $R_1(v_1)$ and $R_2(w_1)$ cover $B \setminus F$, and exactly the faces in $R_1(w_3)$ are covered twice. Hence,

$$\mathbb{1} - (r_3(v_2) + r_1(w_2) + r_1(v_1) + r_2(w_1) - r_1(w_3)) = e_F.$$

Subcase $\mathbf{w_1 \neq v_2}$. This is analogous to Case 2, Subcase $w_1 \neq v_2$, Figure 2.7 (b) shows the situation. △

The dimension of the span of the face vectors e_F is $f-1$ and Claim 2 implies that the span of the face vectors has dimension $f-1$ as well. By Claim 1 the system (2.1) has at most $f-1$ linearly independent equations. Altogether, this implies that there exists a selection $r_1, \dots, r_{f-1}$ of linearly independent rows of (2.1). Note that in the proof of Claim 1 one such selection is explicitly given. Solving this subsystem with $f-1$ equations yields the unique solution of (2.1). □

Next we show how to obtain an efficient algorithm that computes the face weight representation from Theorem 2.10 for a given orthogonal surface $\mathfrak{S}$.

Theorem 2.12. *Let $\mathfrak{S}$ be a non-degenerate, axial, coplanar orthogonal surface generated by n minima. A Schnyder wood S for $\mathfrak{S}$ can be computed in $O(n \log n)$ time. Given S, the translation vector and the face weights can be computed in $O(n)$ time.*

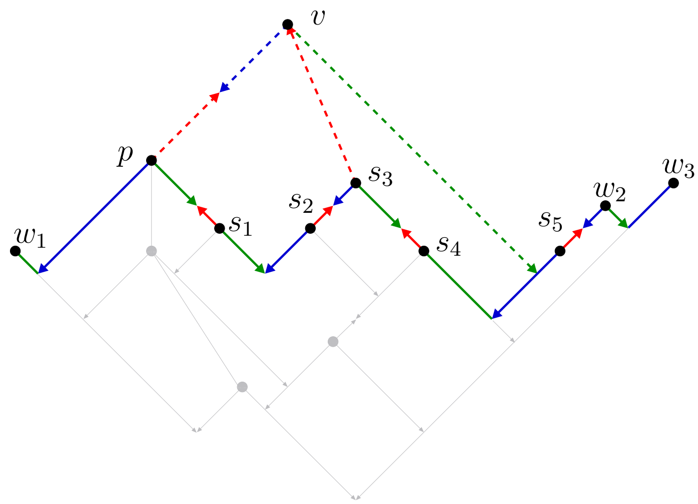

Figure 2.8. Projection of the explored part of $\mathfrak{S}$ onto the sweep plane. The dotted lines represent the new edges when v is added, the other colored lines the sweep front. The gray lines and vertices are the part of the surface that was already explored.

Proof. We will first describe how to extract the Schnyder wood S from $\mathfrak{S}$. The algorithm scans $\mathfrak{S}$ from bottom to top with a sweep plane orthogonal to the

x_1-axis. Figure 2.8 shows a snapshot of the intersection of P with $\mathfrak{S}$. For the sweep algorithm we need a data structure that maintains a finite ordered set of real numbers and allows us to insert and delete elements. Furthermore, we need access to the predecessor and successor of a given query value. Dynamic search trees perform all these operations in logarithmic time, for example the red-black trees presented in [26].

The algorithm builds a Schnyder wood S in the form of clockwise adjacency lists for the vertices, where we also store the information about the type of each edge relative to this vertex, see also the proof of Corollary 2.8. The correctness of the algorithm will follow from the invariant $(*)$.

$(*)$ Having seen a subset $W \subset V$ of the generators of $\mathfrak{S}$ the algorithm has constructed all colored and directed edges of S that are induced by W.

We give a description of the algorithm. A priority queue Q ordered lexicographically with respect to (x_1, x_2) and a dynamic tree P (the sweep front) ordered lexicographically with respect to $(x_2, -x_3)$ are the data structures used. Initialize S as a path of green-blue bidirected edges between the vertices with minimum x_1-coordinate which are ordered by increasing x_2-coordinate. Then, P is also initialized with these vertices, Q with all other vertices.

A step of the algorithm takes the first element v of Q, adds it to P and creates a representative for v in S. The blue outgoing edge of v to its predecessor p in P is added. If $p_1 = v_1$ the edge is green-blue bidirected, if $p_2 = v_2$ it is red-blue bidirected, otherwise it is unidirected. Let $s_1, \ldots, s_\ell$ be the successors of v in P, where s_ℓ is the first one with smaller or equal x_3-coordinate than v. Remove $s_1, \ldots, s_{\ell-1}$ from P adding a red unidirected edge from s_i to v in S for those s_i which do not yet have a red outgoing edge. Finally, check if u, the vertex to be added next, lies on the same x_1-flat as v. In this case u and v will be joined by a green-blue edge when u is considered. If not, add the green outgoing edge of v which ends at s_ℓ. If $(s_\ell)_3 = v_3$ this edge is green-red bidirected, otherwise it is green unidirected. This is done for all vertices in Q and the invariant $(*)$ guarantees that the result is a Schnyder wood S induced by $\mathfrak{S}$.

So we turn to proving that the invariant $(*)$ indeed holds. It is easy to see that it holds after the initialization. So we assume by induction, that only edges incident to the new vertex v have to be checked. There can be no incoming unidirected green or blue edges at v in S at this time, because their starting point has bigger x_1-coordinate than v. The red outgoing edge of v cannot be in S either. It is easy to check that the blue outgoing edge of v and its red incoming edges are geodesic arcs on $\mathfrak{S}$. If the green edge is added, it also corresponds to a geodesic arc. If the vertex u that is to be added next is the endvertex of the green outgoing edge of v, this edge is not induced by W yet. We have thus shown, that all edges added to S

belong to a Schnyder wood induced by $\mathfrak{S}$. Also, the induced orthogonal arcs are all used by an edge. In the case where the green outgoing edge of v is not added, this orthogonal arc is not induced by the explored part of the surface yet. This proves that the invariant $(*)$ holds.

We now show the $O(n \log n)$ complexity bound for the above algorithm. We access the predecessor of a vertex only when it is inserted and its successor only when it is inserted or deleted. As we insert and delete every vertex at most once, this proves the time bound of $O(n \log n)$. Edges can be added in constant time maintaining the clockwise ordering of the adjacency lists. This is possible since new adjacencies always can be added in front of or behind the so far newest vertex in an existing partial list.

The second part of the algorithm is the computation of the face weights. The translation (t_1, t_2, t_3) can be read off the coordinates of the three special vertices. We normalize all vertex coordinates by subtracting the translation. The faces are now considered one by one. When considering a face F, we first determine of which of the possible twenty types F is. As indicated in the proof of Theorem 2.10 there are two cases where the boundary of F is a clockwise or counterclockwise directed cycle. The other eighteen cases correspond to the four subcases of Case 2 and the two subcases of Case 3 in the proof, multiplied with the number of colors. These six cases are:

- F has two unidirected edges of the same color, $w_1 = v_2$, e_3 is directed (w_3, v_3)
- F has two unidirected edges of the same color, $w_1 = v_2$, e_3 is directed (v_3, w_3)
- F has two unidirected edges of the same color, $w_1 \neq v_2$, e_3 is directed (w_3, v_3)
- F has two unidirected edges of the same color, $w_1 \neq v_2$, e_3 is directed (v_3, w_3)
- F has two unidirected edges of different color, $w_1 = v_2$
- F has two unidirected edges of different color, $w_1 \neq v_2$

We determine the vertices v_1, v_2, v_3 respectively, $v_1, w_1, v_2, w_2, v_3, w_3$. As all coordinates are normalized, the coordinates of the vertices correspond to the respective regions' weights and the weight of F can be calculated as in the proof of Theorem 2.10. For example, in the case shown in Figure 2.5 on the right, the weight of F is $c - (w_2)_1 - (w_3)_2 - (w_2)_3$ where $x + y + z = c$ is the plane on which the minima lie after the translation.

The runtime of the procedure for one face F cannot be bounded by a constant, because the boundary of F has to be scanned. But every edge has to be considered only a constant number of times when calculating the weight of F, and every edge lies on at most two inner faces. Since the number of edges is linear in the number of vertices for planar graphs, this yields a linear runtime. $\square$

2.2.2 Rigidity and Coplanarity

coplanar surface
rigid surface

The face-counting produces a *coplanar orthogonal surface* for a given Schnyder wood and in Section 2.1 we have discussed how a *rigid orthogonal surface* for a given Schnyder wood can be constructed. Coplanarity and rigidity are both useful concepts in the realm of orthogonal surfaces. It is therefore natural to ask whether for every Schnyder wood there is an orthogonal surface that is rigid and coplanar. In this section we present an example of a Schnyder wood for which a geodesic embedding can be either rigid or coplanar, but not both.

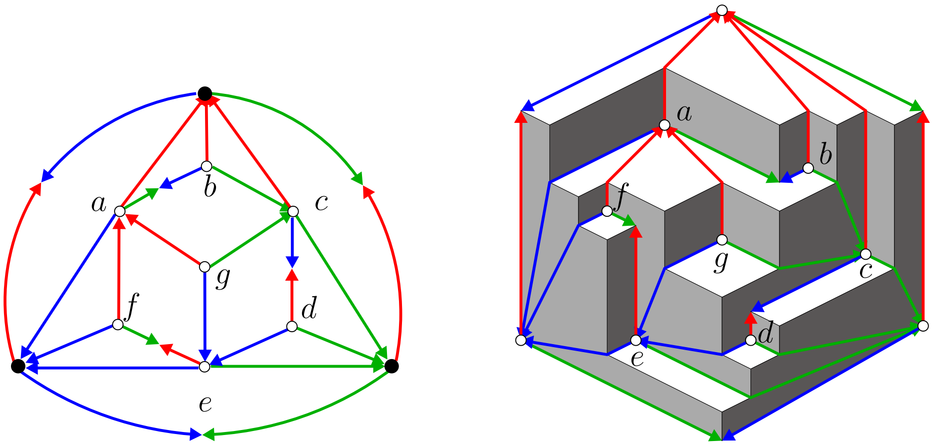

Figure 2.9. A Schnyder wood on a rigid, but not coplanar surface

Proposition 2.13. *The Schnyder wood shown in Figure 2.9 cannot be embedded on a rigid and simultaneously coplanar surface.*

Proof. Assume for the sake of contradiction that there is such an embedding. Coplanarity means that $v_1 + v_2 + v_3 = c = w_1 + w_2 + w_3$ for all $v, w \in V$, hence, $v_i = w_i$ implies $v_{i-1} - w_{i-1} = w_{i+1} - v_{i+1}$. In the Schnyder wood from Figure 2.9 rigidity requires $f_1 > g_1$, $b_2 > g_2$ and $d_3 > g_3$. We use the symbol $\prec$ to highlight the use of rigidity in the following calculation.

$$\begin{aligned}
c_3 < b_3 < a_3 < g_3 \prec d_3 \quad &\Rightarrow \quad a_3 - b_3 < d_3 - c_3 \\
& \qquad\; d_3 - c_3 = c_1 - d_1 \\
e_1 < d_1 < c_1 < g_1 \prec f_1 \quad &\Rightarrow \quad c_1 - d_1 < f_1 - e_1 \\
& \qquad\; f_1 - e_1 = e_2 - f_2 \\
a_2 < f_2 < e_2 < g_2 \prec b_2 \quad &\Rightarrow \quad e_2 - f_2 < b_2 - a_2 \\
& \qquad\; b_2 - a_2 = a_3 - b_3
\end{aligned}$$

Concatenating the inequalities from the right column of the calculation we obtain the contradiction $a_3 - b_3 < a_3 - b_3$. □

2.3 Height Representations of Orthogonal Surfaces

In Theorem 2.10 we have shown that a coplanar orthogonal surface $\mathfrak{S}$ can be represented by a Schnyder wood S and a vector $(w_F)_{F\in\mathcal{F}}$ of weights for the bounded faces of S. This section is concerned with *height representations* of orthogonal surfaces. Throughout this section $\mathfrak{S}$ is assumed to be a *normalized orthogonal surface*, that is the suspension vertices are

normalized surface

$$a_1 = (\alpha_1, 0, 0),\ a_2 = (0, \alpha_2, 0),\ a_3 = (0, 0, \alpha_3).$$

Let $\mathfrak{S}$ be generated by the set of minima $\mathcal{V}$. Let $\mathcal{W}$ denote the set of maxima of $\mathfrak{S}$ and $\mathcal{E}$ the set of edge-points, see Section 1.2 for the definitions of the *characteristic points*. Furthermore $\mathcal{F}$ denotes the set of bounded flats of $\mathfrak{S}$. We fix a geodesically embedded map M^σ on $\mathfrak{S}$. By Theorem 1.7 $\mathfrak{S}$ induces a Schnyder wood on M^σ, and we denote this Schnyder wood by S. The dual map of M^σ is denoted by $M^{\sigma*}$, the dual Schnyder wood by S^*, and the *primal dual completion map* by $\widetilde{M}$.

$\mathcal{V}, \mathcal{W}, \mathcal{E}$

characteristic points

primal dual completion map

Let $n = |\mathcal{V}|$, $m = |\mathcal{E}|$, $f = |\mathcal{W}| + 1$, and note that n, m, f are the number of vertices, edges, and faces of M^σ, respectively. The components of the *height-vector* $h(\mathfrak{S})$ which has dimension $n + f - 1 = m + 1$, are indexed by the elements of $\mathcal{V} \cup \mathcal{W}$, and defined as $h(\mathfrak{S})_x = x_1 + x_2 + x_3$. The purpose of this section is to prove the following theorem.

height-vector $h(\mathfrak{S})$

Theorem 2.14. *Let $\mathfrak{S}$ be a normalized orthogonal surface with Schnyder wood S and height-vector $h(\mathfrak{S})$. Then, S and $h(\mathfrak{S})$ uniquely determine $\mathfrak{S}$.*

It is well known that S uniquely determines the combinatorial type of $\mathfrak{S}$, that is the incidences between the characteristic points. What we prove here is that, under the assumption that the combinatorics of $\mathfrak{S}$ are known, $h(\mathfrak{S})$ uniquely determines the geometry of $\mathfrak{S}$. In other words $h(\mathfrak{S})$ uniquely determines the coordinates of all flats of $\mathfrak{S}$.

The heart of the proof of Theorem 2.14 is Lemma 2.19. The idea for the proof of Lemma 2.19 is due to Stefan Felsner, who used it in the context of triangle contact representations, see [6] for more on this topic. Before we come to this part of the proof we need some preparation.

Lemma 2.18 will show that it suffices to prove Theorem 2.14 for surfaces where only the three suspension vertices of M^σ are incident to the unbounded flats. Therefore, we do not give the following definitions in full generality but restrict ourselves to such surfaces. We call the three bounded flats on which the suspension vertices lie the *outer flats* of $\mathfrak{S}$, and the other bounded flats are the *inner flats* of $\mathfrak{S}$.

outer/inner flat

We now introduce the *skeleton* $R(\mathfrak{S})$ of an orthogonal surface, which is an important tool for the proof of Theorem 2.14. Recall from Section 1.3 that the

skeleton $R(\mathfrak{S})$

α_S-orientation Schnyder wood S corresponds to an α_S*-orientation* of the primal-dual completion map $\widetilde{M}$ in which the primal and dual vertices have out-degree 3 and the edge-vertices have out-degree 1. Let $\widetilde{M}'$ be obtained from $\widetilde{M}$ by deleting v_∞, its six neighbors, and all edges incident to these seven vertices. The skeleton is the graph obtained from $\widetilde{M}'$ by deleting all edges that are outgoing at an edge-vertex of $\widetilde{M}'$ in the orientation corresponding to S. It will be convenient in this section to refer

white/gray vertex to a primal vertex of $R(\mathfrak{S})$ as a *white vertex* and to a dual vertex as a *gray vertex*, see Figure 2.10 (b) for an example. Note that in particular the suspension vertices are colored white. We do not assign a color to the edge-vertices of $R(\mathfrak{S})$.

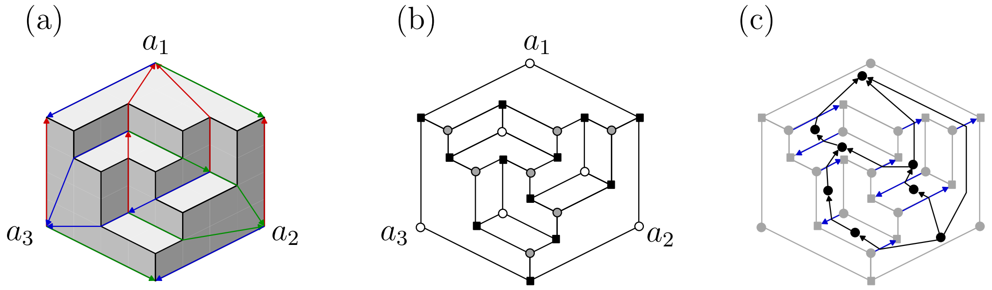

Figure 2.10. An orthogonal surface $\mathfrak{S}$ with a Schnyder wood, the skeleton $R(\mathfrak{S})$ and the graph $B(\mathfrak{S})$.

We regard $R(\mathfrak{S})$ as an embedded graph. More precisely we use the embedding that is obtained by projecting $\mathfrak{S}$ orthogonally to the plane $x_1 + x_2 + x_3 = 1$. Then, the characteristic points and orthogonal arcs of $\mathfrak{S}$ are the vertices respectively edges of $R(\mathfrak{S})$. Furthermore, the bounded faces of $R(\mathfrak{S})$ correspond to the bounded flats of $\mathfrak{S}$ and we refer to them as the flats of $R(\mathfrak{S})$. It can be easily checked, that $R(\mathfrak{S})$ has the following properties.

(P1) Every edge-vertex is adjacent to a white and a gray vertex.

(P2) Every flat has exactly two white-gray color changes on its boundary.

We call each of the two edge-vertices that are incident to a white as well as a

special edge-vertex gray vertex of a flat F a *special edge-vertex* of F. We will use the fact that, by property (P1), every edge-vertex v_e is special for exactly two of the three flats on which it lies. The three edge-vertices incident to the unbounded flats are special for their two bounded flats, not for the unbounded one. Hence, every bounded flat is special for two edge-vertices and every edge-vertex is special for two bounded flats. This implies that the number of flats of $R(\mathfrak{S})$ is $|\mathcal{F}| = |\mathcal{E}| = m$. We need the following lemma for the proof of Theorem 2.14.

Lemma 2.15. *Every recoloring of the gray and white non-suspension vertices of $R(\mathfrak{S})$ that satisfies (P1) and (P2) is the skeleton of some surface $\mathfrak{S}'$.*

Remark 2.16. Note that the recoloring itself is not the crucial part of the above statement. Essentially, every graph with the structural properties of a skeleton graph and a coloring that satisfies (P1) and (P2) is the skeleton graph of some orthogonal surface. But the chosen formulation is more compact and suffices for our purposes.

Proof of Lemma 2.15. On every bounded flat $F \in \mathcal{F}$ we choose a vertex $w(F)$ that is white and adjacent to one of the special edge-vertices. We also choose a gray vertex $g(F)$ that is adjacent to the other special edge-vertex of F. We now construct an oriented graph as follows. All edges of the skeleton are directed towards the incident edge-vertex. We add a directed edge from every edge-vertex which is adjacent to two gray vertices on some flat F to $w(F)$. Similarly directed edges are added from edge-vertices that are adjacent to two white vertices on F to $g(F)$. By construction none of these edges cross, see Figure 2.11.

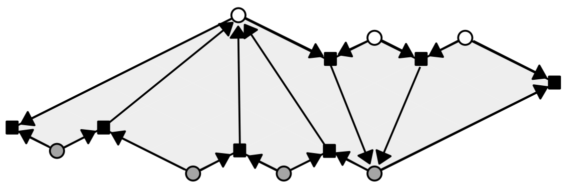

Figure 2.11. Constructing a primal dual completion map by adding directed edges.

The resulting graph can be augmented to a primal-dual completion map $\widetilde{N}$ by adding the seven vertices and oriented edges that have been deleted to obtain $\widetilde{M}'$ from $\widetilde{M}$. In this orientation of $\widetilde{N}$ the white and gray vertices have out-degree 3 and the edge-vertices have out-degree 1. Thus, this orientation yields a primal-dual pair of Schnyder woods. The surface that is obtained by face-counting from the primal Schnyder wood is $\mathfrak{S}'$, see the proof sketch of Theorem 1.7. $\square$

We need another fact about the skeleton graph. Let $R^*(\mathfrak{S})$ denote the dual of $R(\mathfrak{S})$. Let $B(\mathfrak{S})$ be the subgraph of $R^*(\mathfrak{S})$ induced by the vertices representing bounded 1-flats and 2-flats. An edge of $B(\mathfrak{S})$ is oriented from the 1-flat to the 2-flat if the dual edge in $R(\mathfrak{S})$ is incident to a minimum and from the 2-flat to the 1-flat otherwise, see Figure 2.10 (c). An edge orientation of a graph is a *bipolar orientation* with source s and sink t if it is a acyclic and every vertex but s and t has incoming as well as outgoing edges. More about properties of planar bipolar orientations can be found in Section 3.4, where we study the number of bipolar orientations of planar maps.

$R^*(\mathfrak{S})$

$B(\mathfrak{S})$

bipolar orientation

Lemma 2.17. *The orientation of $B(\mathfrak{S})$ is bipolar. The source is the outer 2-flat, the sink is the outer 1-flat.*

Proof. It is obvious that the vertices v_1 and v_2 representing the outer 1-flat and the outer 2-flat have only incoming respectively outgoing edges. Furthermore, every vertex of $B(\mathfrak{S})$ other than v_1, v_2 has incoming as well as outgoing edges. Thus, we only have to show that the orientation is acyclic.

orthogonal arc Note that all edges of $R(\mathfrak{S})$ dual to those of $B(\mathfrak{S})$ are *orthogonal arcs* that are part of a blue edge either in S or in S^*. If e is an edge from a 1-flat to a 2-flat, then the dual edge e^* is a blue edge of S that crosses e from right to left. If e goes from a 2-flat to a 1-flat, then the dual edge is a blue edge of S^* crossing it from left to right. Now suppose for the sake of contradiction that $B(\mathfrak{S})$ has a directed cycle C. The plane embedding of $R(\mathfrak{S})$ allows us to classify a directed *cw-cycle* simple cycle as clockwise (*cw-cycle*) if the interior, $\mathsf{Int}(C)$, is to the right of C or *ccw-cycle* as counterclockwise (*ccw-cycle*) if $\mathsf{Int}(C)$ is to the left of C. We may assume that C is a ccw-cycle. Then, the blue edges of S that are dual to the edges of C all point into $\mathsf{Int}(C)$. Since no blue unidirected edges of S cross 1-flats or 2-flats, no blue edge of S points from the interior to the exterior of C. Thus, the blue special vertex of S must be in the interior of C. But this is impossible since this vertex is not surrounded by bounded 1-flats and 2-flats. □

We need to introduce one more tool before we start with the proof of The- $\mathcal{L}(F)$ orem 2.14. We endow every flat F of $R(\mathfrak{S})$ with a linear order $\mathcal{L}(F)$ on the edge-vertices of F as follows. The outgoing edges of the edge-vertices of $\widetilde{M}$ partition F into quadrangles, and each of these quadrangles consists of a white vertex, a gray vertex and two edge-vertices. Choose one of the special edge-vertices as the first element v_1 of $\mathcal{L}(F)$, and the other edge-vertex on its quadrangle as v_2, see Figure 2.12. The third element v_3 of $\mathcal{L}(F)$ is the edge-vertex from the other quadrangle of v_2 and so on. Thus, the last element of $\mathcal{L}(F)$ is the other special edge-vertex. We refer to the inverse order as $\mathcal{L}^{-1}(F)$. Note that $\mathcal{L}^{-1}(F)$ is obtained by the same procedure starting with the other special edge-vertex.

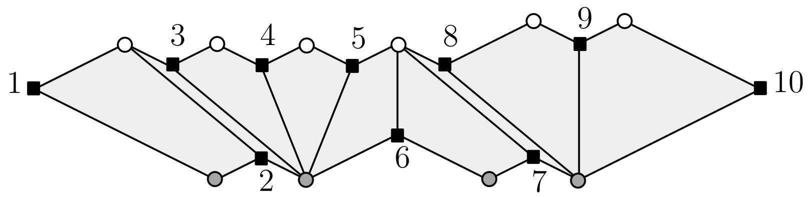

Figure 2.12. The order $\mathcal{L}(F)$ of the edge-vertices of a flat F.

$H(\mathfrak{S})$ We now sketch the proof of Theorem 2.14 for which we use the 0-1-matrix $H(\mathfrak{S})$. The rows of $H(\mathfrak{S})$ are indexed by $\mathcal{V} \cup \mathcal{W}$ and its columns by $\mathcal{F}$. Hence, $H(\mathfrak{S})$ has $m+1$ rows and m columns. An entry $(H(\mathfrak{S}))_{x,y}$ is 1 if x lies on the flat y and 0 otherwise. Table 2.1 shows the matrix $H(\mathfrak{S})$ and others matrices that will be introduced later for the surface $\mathfrak{S}_0$ depicted in Figure 2.13.

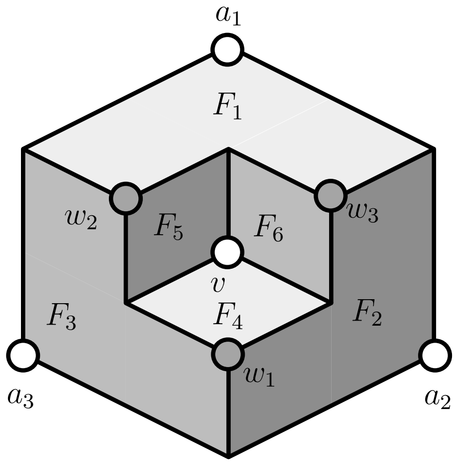

Figure 2.13. A small orthogonal surface $\mathfrak{S}_0$ for the illustration of the matrices $H(\mathfrak{S}_0)$, $H'(\mathfrak{S}_0)$, $C(\mathfrak{S}_0)$, and $C'(\mathfrak{S}_0)$, see Table 2.1.

$$H(\mathfrak{S}_0) = \left(\begin{array}{ccc|ccc} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 \end{array}\right) \begin{matrix} a_1 \\ a_2 \\ a_3 \\ v \\ w_1 \\ w_2 \\ w_3 \end{matrix}, \quad C(\mathfrak{S}_0) = \begin{pmatrix} -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix} \begin{matrix} F_4 \\ F_5 \\ F_6 \end{matrix},$$

$$H'(\mathfrak{S}_0) = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{matrix} v \\ w_1 \\ w_2 \\ w_3 \end{matrix}, \quad C'(\mathfrak{S}_0) = \begin{pmatrix} 0 & 0 & 1 & 1 \\ -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix} \begin{matrix} \vec{a} \\ F_4 \\ F_5 \\ F_6 \end{matrix}$$

Table 2.1. The matrices $H(\mathfrak{S}_0)$, $H'(\mathfrak{S}_0)$, $C(\mathfrak{S}_0)$, and $C'(\mathfrak{S}_0)$ for the orthogonal surface $\mathfrak{S}_0$ from Figure 2.13

We aim to prove that $H(\mathfrak{S})$ has full rank. This implies Theorem 2.14 since then the coordinates of the flats of $\mathfrak{S}$ are the unique solution of the linear equation system

$$H(\mathfrak{S}) \cdot x = h(\mathfrak{S}). \tag{2.2}$$

The first step in the proof is to show that we may work with orthogonal surfaces that have only three minima on the unbounded flats, see Lemma 2.18. We then work with the balance matrix $C(\mathfrak{S})$ to show that $H(\mathfrak{S})$ has full rank. We explain below how exactly $C(\mathfrak{S})$ is obtained from $H(\mathfrak{S})$, and content ourselves for the moment with the following explanation. A solution to the system

$$C(\mathfrak{S}) \cdot y = 0 \tag{2.3}$$

is an assignment of weights to the gray and white vertices of $R(\mathfrak{S})$ other than the three suspension vertices. This assignment has the property that the weights of the white vertices sum up to the same value as the weights of the gray vertices for every inner flat. The System (2.3) is a homogenous system and we will see later that it has one more variable than equations. Hence, (2.3) has a non-trivial solution y_0. The crucial step in the proof of Theorem 2.14 is showing the following statement, see Lemma 2.19. By exchanging the colors white and gray for all vertices of $R(\mathfrak{S})$ to which y_0 assigns negative weight we obtain another orthogonal surface $\mathfrak{S}'$ such that $C(\mathfrak{S}') \cdot |y_0| = 0$. Here $|y_0|$ denotes the componentwise absolute value of y_0 and we will verify this statement with the help of Lemma 2.15. This combinatorial interpretation subsequently helps us to argue that $C(\mathfrak{S})$ and $H(\mathfrak{S})$ have full rank.

This concludes the proof sketch and Lemma 2.18 is the first part of the proof of Theorem 2.14.

Lemma 2.18. *If Theorem 2.14 holds for orthogonal surfaces which have exactly three minima incident to the unbounded flats, then it holds for all orthogonal surfaces.*

Proof. Let $\mathfrak{S}$ be an orthogonal surface with more than three minima on the outer face. We construct a surface $\mathfrak{S}'$ with three minima on the outer face from $\mathfrak{S}$ and show that if Theorem 2.14 holds for $\mathfrak{S}'$ then it also holds for $\mathfrak{S}$. Since $\mathfrak{S}$ is normalized, the coordinates of its suspension vertices can be read off the height-vector $h(\mathfrak{S})$. We use these coordinates to define three new vertices.

$$a_1' = (0, \alpha_2 + 1, \alpha_3 + 1),\ a_2' = (\alpha_1 + 1, 0, \alpha_3 + 1),\ a_3' = (\alpha_1 + 1, \alpha_2 + 1, 0)$$

The surface $\mathfrak{S}'$ is the normalization of the surface $\mathfrak{S}'$ generated by the set $-\mathcal{W} \cup \{-a_1', -a_2', -a_3'\}$, see Figure 2.14.

The translation that normalizes the surface is $\vec{n} = (\alpha_1 + 1, \alpha_2 + 1, \alpha_3 + 1)$. The generating minima of $\mathfrak{S}'$ other than the three new suspension vertices $-a_1', -a_2', -a_3'$ are in bijection with the maxima of $\mathfrak{S}$. Hence, the maxima of $\mathfrak{S}'$ are the points in $-\mathcal{V} + \vec{n}$. Note that $h(\mathfrak{S}')$ can also be calculated from $h(\mathfrak{S})$ and $\vec{n}$.

The Schnyder wood associated with $\mathfrak{S}'$ is S^* and $-a_1', -a_2', -a_3'$ form the outer face of the underlying suspension map M^{σ^*}. The matrix $H(\mathfrak{S}')$ has three more rows and three more columns than $H(\mathfrak{S})$, corresponding to the three additional vertices respectively three new bounded flats. Furthermore, $H(\mathfrak{S})$ is a minor of $H(\mathfrak{S}')$ and the three new rows corresponding to the suspension vertices $-a_i' + \vec{n}$ have non-zero entries only in the three new columns. Thus, if the m columns of $H(\mathfrak{S})$ have a

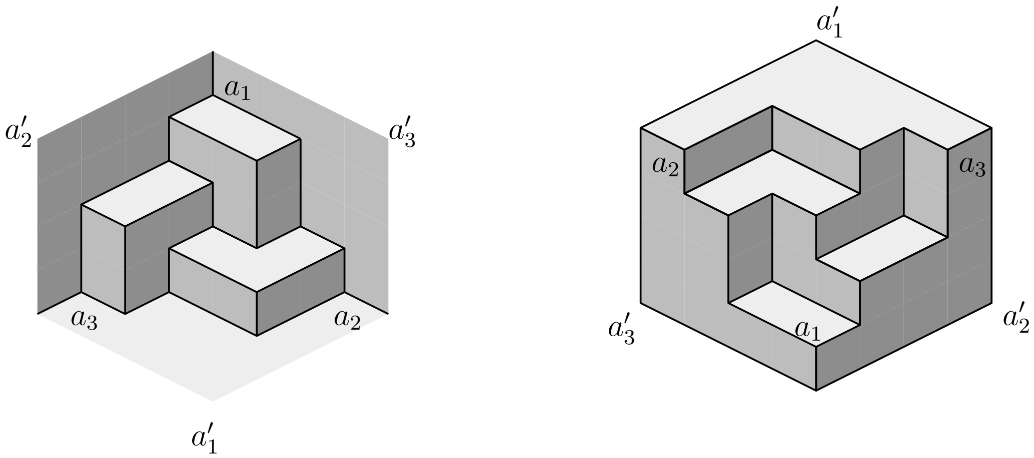

Figure 2.14. A surface $\mathfrak{S}$ and the associated surface $\mathfrak{S}'$ as described in Lemma 2.18.

non-trivial linear dependence, then the columns of $H(\mathfrak{S}')$ also have a non-trivial linear dependence. This implies that if $H(\mathfrak{S}')$ has full rank, then so does $H(\mathfrak{S})$. □

We now introduce the balance matrix that we need for the proof of Theorem 2.14, as we have mentioned in the proof sketch. For the definition of the balance matrix, we use the auxiliary matrix $H'(\mathfrak{S})$. This is a submatrix of $H(\mathfrak{S})$ of dimensions $(m-2)\times(m-3)$ which is obtained from $H(\mathfrak{S})$ by deleting the three rows corresponding to the outer vertices of S and the three columns corresponding to the outer flats, see Table 2.1. The three deleted column vectors each have an entry 1 in a row which is otherwise 0, since it belongs to an outer vertex. It follows that if $H'(\mathfrak{S})$ has full rank, then so does $H(\mathfrak{S})$. It remains to prove that $H'(\mathfrak{S})$ indeed has full rank and we define the balance matrix $C(\mathfrak{S})$ for this purpose.

The *balance matrix* $C(\mathfrak{S})$ is obtained from $H'(\mathfrak{S})^T$ by multiplying all columns corresponding to minima of $\mathfrak{S}$ by -1, see Table 2.1. A solution of (2.3) corresponds to an assignment of weights to the inner white and gray vertices of $R(\mathfrak{S})$ such that for every inner flat of the skeleton the sum of weights of its white vertices equals the sum of weights of its gray vertices. Since this is a homogeneous system with $m-3$ equations and $m-2$ variables it has a non-trivial solution y_0. We define $\Sigma(y_0)$ to be the diagonal matrix which has ith diagonal entry -1 if the ith component of y_0 is negative and 1 otherwise. Thus, $|y_0|$ is a solution of the system *balance matrix $C(\mathfrak{S})$*

$$C(\mathfrak{S}) \cdot \Sigma(y_0) \cdot y = 0. \tag{2.4}$$

We call a vertex of $R(\mathfrak{S})$ *positive/negative* if the corresponding component of y_0 is positive or negative. The next claim is the core of the proof of Theorem 2.14. *positive/negative vertex*

Lemma 2.19. *Let y_0 be a non-trivial solution of (2.3).*

(i) All gray vertices lying on an outer flat have the same sign with respect to y_0.

(ii) There exists a surface $\hat{\mathfrak{S}}$ such that

$$C(\hat{\mathfrak{S}}) = C(\mathfrak{S}) \cdot \Sigma(y_0) \text{ or } C(\hat{\mathfrak{S}}) = C(\mathfrak{S}) \cdot \Sigma(-y_0).$$

Figure 2.15 shows an orthogonal surface, and the surface $\hat{\mathfrak{S}}$, obtained by switching the color of the encircled vertex from gray to white. Since the proof of Lemma 2.19 is rather long we first show how it can be used to prove Theorem 2.14.

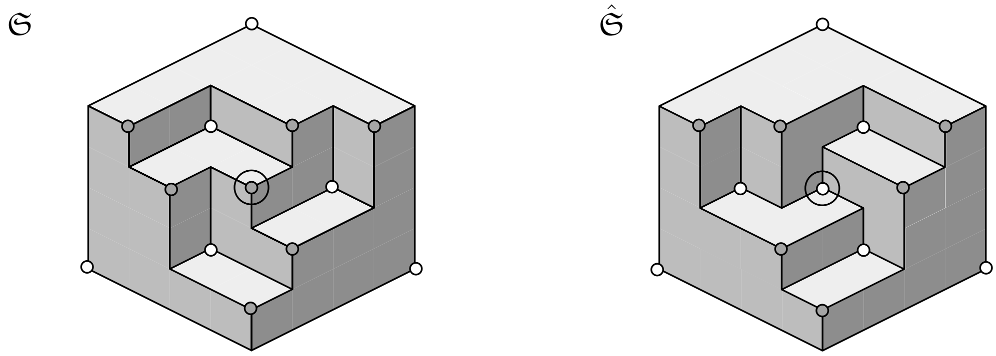

Figure 2.15. A surface $\mathfrak{S}$ and the corresponding surface $\hat{\mathfrak{S}}$, as described in Lemma 2.19.

augmented balance matrix $C'(\mathfrak{S})$

Proof of Theorem 2.14. Consider the indicator vector $\vec{a}$ of the maxima that lie on the outer 1-flat of $\mathfrak{S}$. Let the *augmented balance matrix* $C'(\mathfrak{S})$ denote the $(m-2) \times (m-2)$ matrix obtained from $C(\mathfrak{S})$ by adding $\vec{a}$ as the first row, see Table 2.1. We aim to show that $C'(\mathfrak{S})$ has full rank. This implies that $C(\mathfrak{S})$ also has full rank and thus suffices to prove the theorem. We first show that the system

$$C'(\mathfrak{S}) \cdot y = \vec{e_1} \tag{2.5}$$

has a solution, where $\vec{e_1}$ has first component 1 and all other components are 0. Let y_0 be a non-trivial solution of (2.3). By Lemma 2.19 there is a surface $\hat{\mathfrak{S}}$ with $C(\hat{\mathfrak{S}}) = C(\mathfrak{S}) \cdot \Sigma(y_0)$ or $C(\hat{\mathfrak{S}}) = C(\mathfrak{S}) \cdot \Sigma(-y_0)$. Since $-y_0$ also is a non-trivial solution of (2.3), we may assume that $C(\hat{\mathfrak{S}}) = C(\mathfrak{S}) \cdot \Sigma(y_0)$.

We show that scaling y_0 yields a solution of (2.5). The suspension vertices of $\hat{\mathfrak{S}}$ are white and cannot change their color and all gray vertices of the outer flats have the same sign in y_0. Since no flat of $\hat{\mathfrak{S}}$ can have only white vertices because of (P2) these gray vertices must all be non-negative with respect to y_0. Thus, we have that $\vec{a} \cdot \Sigma(y_0) = \vec{a}$ and we aim to show that $\langle \vec{a}, |y_0| \rangle > 0$.

Some component of $|y_0|$ has positive value α_0. We use Lemma 2.17 to show that $\langle \vec{a}, |y_0| \rangle \geq \alpha_0$. We may assume without loss of generality that some gray vertex v_0 of $R(\hat{\mathfrak{S}})$ has value α_0. Then the sum of the values of the white vertices incident to $F_1(v_0)$, the 1-flat of v_0, must be at least α_0, since $|y_0|$ is non-negative. Each of these white vertices is the minimum of some 2-flat that shares an orthogonal arc of $\hat{\mathfrak{S}}$ with $F_1(v_0)$. Thus, the sum of the maxima of these 2-flats must be at least α_0, and these maxima in turn lie on 1-flats. This illustrates that the value α_0 is passed on from flat to flat in accordance with the edge directions of the graph $B(\hat{\mathfrak{S}})$ from Lemma 2.17. Since $B(\hat{\mathfrak{S}})$ is bipolar it follows that the values of the gray vertices of the outer 1-flat sum up to at least α_0. This implies $\langle \vec{a}, |y_0| \rangle = \alpha \geq \alpha_0 > 0$. Thus,

$$C'(\mathfrak{S}) \cdot y_0 = C'(\mathfrak{S})\Sigma(y_0) \cdot |y_0| = C'(\hat{\mathfrak{S}}) \cdot |y_0| = \alpha \cdot \vec{e_1}$$

and scaling yields the desired solution of (2.5).

Now, let z_0 be some non-trivial solution of (2.5). Lemma 2.19 shows that there is a surface $\mathfrak{S}(z_0)$ with $C(\mathfrak{S}(z_0)) = C(\mathfrak{S}) \cdot \Sigma(z_0)$ or $C(\mathfrak{S}(z_0)) = C(\mathfrak{S}) \cdot \Sigma(-z_0)$. Thus, $|z_0|$ is a solution of $C'(\mathfrak{S}(z_0)) \cdot y = \vec{e_1}$. Using Lemma 2.17, this implies that the absolute value of every component of z_0 is bounded by 1.

We are ready to conclude that $C'(\mathfrak{S})$ has full rank. Otherwise the kernel of $C'(\mathfrak{S})$ would be a non-trivial vector space $\vec{V}$ that of course contains vectors of arbitrary length. Since all vectors in $z_0 + \vec{V}$ are solutions of $C'(\mathfrak{S}) \cdot y = \vec{e_1}$ this contradicts that all solutions have components of absolute value at most 1. This concludes the proof of Theorem 2.14 based on Lemma 2.19. □

Proof of Lemma 2.19. If all components of y_0 are non-positive or all components of y_0 are non-negative, then we can simply choose $\hat{\mathfrak{S}} = \mathfrak{S}$. So we may assume that y_0 has negative and positive components. We work with a signing of $R(\mathfrak{S})$ where negative vertices are signed $-$ and non-negative vertices are signed $+$. Then there is an edge-vertex v_0 that has a neighbor signed $-$ and a neighbor signed $+$. The first step in the proof is to construct a closed Jordan curve $J(v_0)$ through v_0 that intersects the embedding of $R(\mathfrak{S})$ only in edge-vertices that have negative and non-negative neighbors. Moreover, the negative neighbors of every edge-vertex on $J(v_0)$ are separated from the non-negative ones by $J(v_0)$. This construction will imply Part (i) of the lemma. We then proceed by showing that if there is an edge-vertex $v_0' \notin J(v_0)$ with negative and non-negative neighbors, then the closed Jordan curve $J(v_0')$ through v_0' cannot intersect $J(v_0)$. In conjunction with Lemma 2.15 this enables us to prove that switching the color of all negative vertices, with respect to y_0 or with respect to $-y_0$, yields the surface $\hat{\mathfrak{S}}$. The following observation is essential for the construction of the Jordan curve $J(v_0)$.

Observation 1. It is not possible that all gray vertices of a flat are negative, while all white ones are non-negative or vice versa, since this would imply that a negative number equals a non-negative one.

Claim 1. There is a simple, closed Jordan curve $J(v_0)$ through a sequence of edge-vertices v_i and flats F_i

$$v_0, F_0, v_1, F_1, v_2, \ldots, F_{k-1}, v_k = v_0.$$

This sequence has the following properties. The vertex v_i is a special vertex of the flat F_i and a non-special vertex of F_{i-1}. Furthermore, the clockwise predecessor of v_i on F_i, denoted u_i, is signed $-$ and the clockwise successor w_i is signed $+$, or vice versa. Thus, either all neighbors of the v_i in the interior of $J(v_0)$ are negative and the neighbors in the exterior are non-negative or vice versa.

As above, let v_0 be an edge-vertex that is adjacent to a non-negative and a negative vertex. This vertex v_0 cannot lie on an unbounded face, since these edge-vertices have only one signed neighbor, the suspension vertices are not signed. Since there are two sign changes around v_0 and it is special for two flats, one of them, call it F_0, must have a vertex u_0 signed $-$ and a vertex w_0 signed $+$, both adjacent to v_0. We may assume that u_0 is the predecessor of v_0 on a clockwise walk on the boundary of F_0 and w_0 the successor. We may furthermore assume that v_0 is the minimum of the linear order $\mathcal{L}(F_0)$.

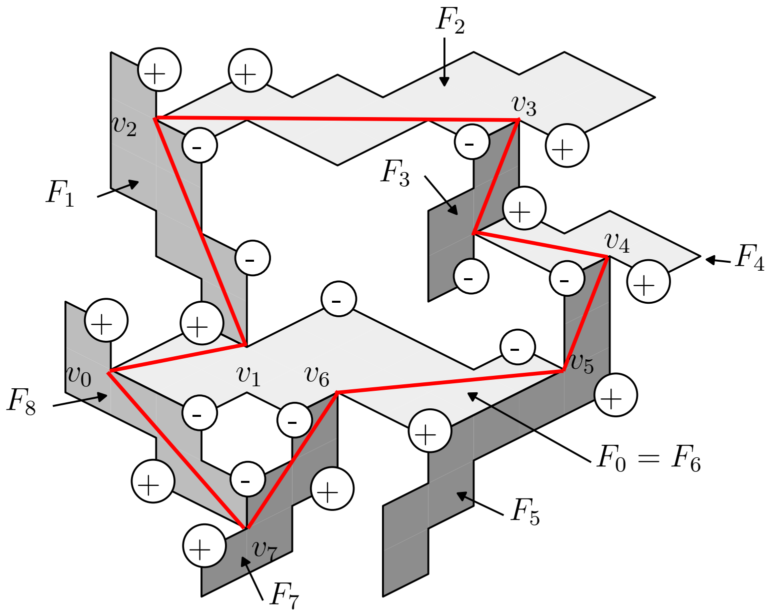

Figure 2.16. The curve $J(v_0)$ separating a positive and a negative region.

We define v_1 to be the smallest edge-vertex in $\mathcal{L}(F_0)$ which has differently signed neighbors on F_0. Then, v_1 cannot be the other special vertex of F_0, since this would contradict Observation 1. Let F_1 be the flat other than F_0 on which v_1 has differently signed neighbors.

By construction v_1 is special for F_1, and since we assumed that u_0 is signed $-$ we have that u_1 is signed $-$ and w_1 is signed $+$. Thus, we can repeat the same construction as for v_1 and F_1 to obtain a vertex v_2 on F_1 and a flat F_2 that have the desired properties. In this way we find a sequence of the form

$$v_0, F_0, v_1, F_1, v_2, \dots v_{k-1}, F_{k-1}.$$

Note that every flat can appear in this sequence at most twice if there are no vertex repetitions, since it is entered through one of its two special vertices. When we enter a flat $F_i = F_j$ for the second time from the vertex v_j we consider the edge-vertices in the inverse order $\mathcal{L}^{-1}(F_i)$. Since u_j, the clockwise predecessor of v_j on F_i, is negative a sign change must occur on both boundary paths between v_i and v_j. Since v_{i+1} is the largest edge-vertex on the boundary of F_i with respect to $\mathcal{L}^{-1}(F_i)$, at which a sign change occurs, we know that $v_{j+1} \neq v_{i+1}$. Since the skeleton is finite the sequence must finally end with a vertex repetition

$$v_0, F_0, v_1, F_1, v_2, \dots, F_{k-1}, v_k$$

where $v_k = v_i$ for some $i \in \{0, \dots, k-1\}$. Note that v_i is not special for F_{i-1}, that is we enter each vertex from a flat for which it is not special. The only vertex, whose non-special flat is not among $F_0, \dots, F_{k-1}$ is v_0. Thus, the sequence is of the form $v_0, F_0, v_1, \dots, F_{k-1}, v_0$.

The sequence $v_0, F_0, v_1, F_1, v_2, \dots, F_{k-1}, v_0$ is now used to obtain $J(v_0)$ as announced at the beginning of the proof. The closed Jordan curve $J(v_0)$, connects the edge-vertices v_i and v_{i+1} within F_i. From the above considerations it follows that $J(v_0)$ can be drawn without self-intersection also on flats which appear twice in the sequence, that is $J(v_0)$ is simple, see Figure 2.16. Note that either the u_i lie in the interior of $J(v_0)$ and the w_i in the exterior, or vice versa. △

We are now ready to prove Part (i) of the lemma. The special edge vertices of the outer flats are special for two outer flats and non-special for an unbounded flat. Thus, by the construction from Claim 1 the curve $J(v_0)$ cannot pass through an outer flat. Therefore every edge-vertex on an outer flat has neighbors with the same sign on this outer flat. Part (i) follows since any two outer flats share a gray vertex. We may therefore assume from now on that all gray vertices on the outer flats are non-negative with respect to y_0, otherwise we may work with $-y_0$ instead of y_0.

We now proceed to prove Part (ii) of the lemma. Suppose there is an edge-vertex v_0' that does not lie on $J(v_0)$. Then, by Claim 1 there exists a closed curve $J(v_0')$ through edge-vertices and flats

$$v_0', F_0', v_1', F_1', v_2', \dots, F_{\ell-1}', v_0'.$$

Claim 2. Two Jordan curves $J(v_0)$ and $J(v_0')$ with $v_0' \notin J(v_0)$ do not intersect.

We show below, that if $J(v_0)$ and $J(v_0')$ enter the same flat from different edge-vertices, then they continue on different flats. Since v_0' does not lie on $J(v_0)$, this implies, that if both curves enter a flat $F_j = F_i'$, then they do so from flats $F_{j-1} \neq F_{i-1}'$, and thus from the two different special vertices of F_j.

An argument similar to the one that $J(v_0)$ does not self-intersect when it enters the same flat F_j twice shows that $J(v_0')$ leaves $F_j = F_i'$ through an edge-vertex on the same side of $J(v_0)$ as v_i'. We may assume that the u_k are signed $-$. If the u_k' are also negative, the argument is the same as in Claim 1. The other case is that u_i', the clockwise predecessor of v_i' on F_i', is signed $+$. In this situation both v_j-v_i'-paths on F_j start and end with the same sign. Since there is an exit vertex v_{j+1} on one of these paths, there has to be another non-special edge-vertex on the same path with differently signed neighbors. Thus, $v_{i+1}' \neq v_{j+1}$ which implies that v_{i+1} lies on the same side of $J(v_0)$ as v_i', and thus $J(v_0)$ and $J(v_0')$ do not intersect.

$\triangle$

We now show Part (ii) of the lemma. We claim that switching the colors of all negative vertices does not violate properties (P1) and (P2). We observe that for every edge-vertex w_e on some $J(v_e)$ we enter w_e from a flat on which it has neighbors of the same color, and that these two neighbors are separated by $J(v_e)$. Thus, the two neighbors of w_e which lie on the same side of $J(v_e)$ have different colors and this implies (P1). Property (P2) still holds, since a color change only involves one or two consecutive segments of monochromatic vertices on the boundary of a flat F. Since both these segments start at a special vertex of F they do not introduce new changes of color on the boundary of F, but only shift the old ones. Thus, we have obtained a surface $\mathfrak{S}'$ with skeleton $R(\mathfrak{S}')$, such that $C(\mathfrak{S}') = C(\mathfrak{S})\Sigma(y_0)$ if the gray vertices of the outer flats are signed $+$ and $C(\mathfrak{S}') = C(\mathfrak{S})\Sigma(-y_0)$ otherwise. $\square$

2.4 Conclusions

In this chapter we have studied connections of orthogonal surfaces and Schnyder woods. In Section 2.1, we presented a new proof of the Brightwell-Trotter Theorem. This approach uses facts about Schnyder woods to prove that the intuitive method of shifting flats can be used to obtain a rigid surface in time $O(n)$.

In Section 2.2 we have presented a generalization of the face-counting approach for the generation of coplanar surfaces from Schnyder woods. We have shown that every coplanar surface can be obtained by generalized face-counting. We also gave an example of a Schnyder wood that cannot be geodesically embedded on a rigid and simultaneously coplanar surface.

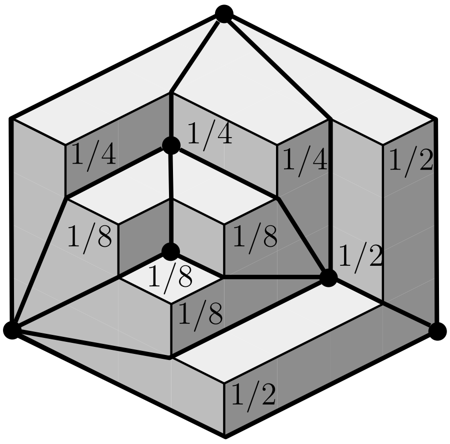

Figure 2.17. An orthogonal surface with geodesically embedded stacked triangulation.

In Section 2.3 we have discussed height representations of orthogonal surfaces. The proof that orthogonal surfaces are uniquely determined by a Schnyder wood plus the heights of all minima and maxima uses the augmented balance matrix $C'(\mathfrak{S})$. We show that $C'(\mathfrak{S})$ has full rank and thus (2.5) has a unique solution. In order to simplify the proof of Theorem 2.14, it would be useful to understand this solution from a combinatorial point of view.

In the case that M^σ is a stacked triangulation we are able to give an interpretation of this solution. Let the solution α of (2.5) be given for some stacked triangulation T. It is easy to check that the solution for $T+v'$ obtained from T by stacking the vertex v' into face F is obtained as follows. The value of v' is defined as $\alpha(F)/2$ and the values for the three new faces as $\alpha(F)/2$. All other values are kept except $\alpha(F)$. A solution for the K_3 simply assigns 1 to the unique bounded face. Then, the above method yields that every vertex v and bounded face F of height h, has value 2^{-r}, see Figure 2.17.

In general we do not have a combinatorial interpretation for the solution of (2.5), and therefore we need to reinterprete it in Lemma 2.19. Hence, solving the following problem could potentially simplify the proof of Theorem 2.14.

Problem 2.20. What is the combinatorial interpretation of the solution of (2.5)?

As we have mentioned above, Felsner proved Lemma 2.19 when working on triangle contact representations, see [6]. Progress on Problem 2.20 could conversely help to answer open questions related to triangle contact representations.

Chapter 3

The Number of Planar Orientations with Prescribed Out-Degrees

Many different combinatorial structures on *planar maps* have attracted the attention of researchers. *planar map* Among them are spanning trees, bipartite perfect matchings (or more generally bipartite f-factors), Eulerian orientations, Schnyder woods, bipolar orientations and 2-orientations of quadrangulations. *Orientations with prescribed out-degrees* are a quite general concept. *orientation with prescribed out-degrees* Remarkably, all the above structures can be encoded as orientations with prescribed out-degrees. Let a planar map M with vertex set V and a function $\alpha : V \to \mathbb{N}$ be given. An orientation X of the edges of M is an *α-orientation* if every vertex v has out-degree $\alpha(v)$. *α-orientation* For the sake of brevity, we refer to orientations with prescribed out-degrees simply as α-orientations in the rest of this chapter.

For some of the above mentioned structures it is not obvious how to encode them as α-orientations. For Schnyder woods on triangulations the encoding by 3-orientations goes back to de Fraysseix and de Mendez [31], see Theorem 1.8. We have seen in Section 1.3 that encoding Schnyder woods on 3-connected planar maps as α-orientations requires the use of an auxiliary map, the primal dual completion map $\widetilde{M}$. Similarly, the encoding for bipolar orientations proposed by Woods [98] and independently by Tamassia and Tollis [88] uses the angle graph $\widehat{M}$ as an auxiliary map. For bipartite f-factors and spanning trees Felsner [41] describes encodings as α-orientations.

Given the existence of a combinatorial structure on a class $\mathcal{M}_n$ of planar maps with n vertices, one of the questions of interest is how many instances of this structure there are for a given map $M \in \mathcal{M}_n$. Especially, one is interested in the minimum and maximum that this number attains on $\mathcal{M}_n$. This question has been treated quite successfully for spanning trees and bipartite perfect matchings. For spanning trees the Kirchhoff Matrix Tree Theorem allows to bound the maximum number of spanning trees of a planar graph with n vertices between 5.02^n and 5.34^n, see [81, 73]. Pfaffian orientations can be used to efficiently calculate the number of bipartite perfect matchings in the planar case, see for example [69]. Kasteleyn [61] has shown that the $k \times \ell$ square grid has asymptotically $e^{0.29 \cdot k\ell} \approx 1.34^{k\ell}$ perfect matchings. The number of Eulerian orientations is studied in statistical physics under the name of ice models, see [9] for an overview. In particular Lieb [65] has shown that the $k \times \ell$ square grid on the torus has asymptotically $(8\sqrt{3}/9)^{k\ell} \approx 1.53^{k\ell}$ Eulerian orientations and Baxter [8] has worked out the asymptotics for the triangular grid on the torus as $(3\sqrt{3}/2)^{k\ell} \approx 2.598^{k\ell}$.

In many cases it is relatively easy to see which maps in a class $\mathcal{M}_n$ carry a unique object of a certain type, while the question about the maximum number is rather intricate. Therefore, we focus on finding the asymptotics for the maximum number of α-orientations that a map from $\mathcal{M}_n$ can carry. Table 3.1 gives an overview of the results of this chapter for different instances of $\mathcal{M}_n$ and α. The entry c in the "Upper Bound" column is to be read as $O(c^n)$, in the "Lower Bound" column as $\Omega(c^n)$ and for the "$\approx c$" entries the asymptotics are known.

Graph class and orientation type	**Lower bound**	**Upper bound**
α-orientations on planar maps	2.59	3.73
Eulerian orientations	2.59	3.73
Schnyder woods on triangulations	2.37	3.56
Schnyder woods on the square grid	≈ 3.209	
Schnyder woods on planar maps	3.209	8
2-orientations on quadrangulations	1.53	1.91
bipolar or. on stacked triangulations	≈ 2	
bipolar orientations on outerplanar maps	≈ 1.618	
bipolar orientations on the square grid	2.18	2.62
bipolar orientations on planar maps	2.91	3.97

Table 3.1. An overview of the results presented in Chapter 3.

The chapter is organized as follows. In Section 3.1 we treat the most general case, where $\mathcal{M}_n$ is the class of all planar maps with n vertices and α can be any integer valued function. We prove an upper bound that applies to every map and every α and in Section 3.1.3 we prove a lower bound for the number of Eulerian orientations. In Section 3.2.1 we consider Schnyder woods on plane triangulations and in Section 3.2.2 the more general case of Schnyder woods on 3-connected planar maps is discussed. We split the treatment of Schnyder woods because the more direct encoding of Schnyder woods on triangulations as α-orientations yields stronger bounds. In Section 3.2.2 we also discuss the asymptotic number of Schnyder woods on the square grid. Section 3.3 is dedicated to 2-orientations of quadrangulations. In Section 3.4, we study bipolar orientations. The square grid is treated in Section 3.4.1 while stacked triangulations, outerplanar maps and planar maps are studied in Section 3.4.2. The upper bound for planar maps relies on a new encoding of bipolar orientations of inner triangulations by $+/-$ vectors. In Section 3.4.3 we characterize the $+/-$ vectors that induce a bipolar orientation. In Section 3.5.1 we discuss the complexity of counting α-orientations. In Section 3.5.2 we show how counting α-orientations can be reduced to counting (not necessarily planar) bipartite perfect matchings and the consequences of this connection are explained as well. We conclude with some open problems.

3.1 The Number of α-Orientations

For convenience, we remind the reader of the following definitions from Section 1.1. A planar map M is a simple planar graph G together with a fixed crossing-free embedding of G in the Euclidean plane. In particular, M has a designated outer (unbounded) face. Recall that we denote the sets of vertices, edges and faces of a given planar map M by $V(M)$, $E(M)$, and $\mathcal{F}(M)$, and their respective cardinalities by $n(M)$, $m(M)$ and $f(M)$. If ambiguities can be excluded we omit the parameter M. The degree of a vertex v will be denoted by $d(v)$.

Let M be a planar map and $\alpha : V \to \mathbb{N}$. An orientation X of the edges of M is an *α-orientation* if every $v \in V$ has out-degree $\alpha(v)$ in X. *α-orientation*

Let X be an α-orientation of M and let C be a directed cycle in X. Define X^C as the orientation obtained from X by reversing all edges of C. Since the reversal of a directed cycle does not affect out-degrees, the orientation X^C is also an α-orientation of M. In the sequel we refer to such a reorientation of a directed cycle as a *cycle flip*. The plane embedding of M allows us to classify a directed simple cycle as clockwise (*cw-cycle*) if the interior, $\mathsf{Int}(C)$, is to the right of C or as counterclockwise (*ccw-cycle*) if $\mathsf{Int}(C)$ is to the left of C. If C is a ccw-cycle of X then we say that X^C is *left of* X and X is *right of* X^C. Felsner proved the following theorem in [41]. *flip* *cw-cycle* *ccw-cycle*

Theorem 3.1. *Let M be a planar map and $\alpha : V \to \mathbb{N}$. The set of α-orientations of M endowed with the transitive closure of the 'left of' relation is a distributive lattice.*

Theorem 3.1 found applications in drawing algorithms in [50, 13], and for enumeration and random sampling of graphs in [52]. In [64] Knauer studies the structure of the set of α-orientations on non-planar graphs and other aspects of α-orientations.

The following observation is easy but useful. Let M and $\alpha : V \to \mathbb{N}$ be given, let $W \subset V$, and let E_W be the edges of M with one endpoint in W and the other endpoint in $V \setminus W$. Suppose all edges of E_W are directed away from W in some α-orientation X_0 of M. The demand of W for $\sum_{w \in W} \alpha(w)$ outgoing edges forces all edges in E_W to be directed away from W in every α-orientation of M. Such *rigid edge* an edge with the same direction in every α-orientation is a *rigid edge* of M. Note that the rigidity of an edge in this sense is not related to the notion of an rigid edge on an orthogonal surface, as defined at the end of Section 1.2.

$r_\alpha(M)$ We denote the number of α-orientations of M by $r_\alpha(M)$. Let $\mathcal{M}$ be a family of pairs (M, α) of a planar map and an out-degree function. Most of this chapter is concerned with lower and upper bounds for $\max_{(M,\alpha)\in\mathcal{M}} r_\alpha(M)$ for some family $\mathcal{M}$. In Sections 3.1.1 and 3.1.3, we deal with bounds that apply to all M and α, while later sections will be concerned with special instances.

3.1.1 An Upper Bound for the Number of α-Orientations

A trivial upper bound for the number of α-orientations on M is 2^m as any edge can be directed in two ways. The following easy but useful lemma improves the trivial bound.

Lemma 3.2. *Let M be a planar map, $A \subset E$ a cycle free subset of edges of M, and α a function $\alpha : V \to \mathbb{N}$. Then, there are at most $2^{m-|A|}$ α-orientations of M. Furthermore, M has less than 4^n α-orientations.*

Proof. Let X be an arbitrary but fixed orientation out of the $2^{m-|A|}$ orientations of the edges of $E \setminus A$. It suffices to show that X can be extended to an α-orientation of M in at most one way. We proceed by induction on $|A|$. The base case $|A| = 0$ is trivial. If $|A| > 0$, then, as A is cycle free, there is a vertex v that is incident to exactly one edge $e \in A$. If v has out-degree $\alpha(v)$ respectively $\alpha(v) - 1$ in X, then e must be directed towards v respectively away from v in every α-orientation of M extending X. In either case the direction of e is determined by X, and by induction there is at most one way to extend the resulting orientation of $E\setminus(A-e)$ to an α-orientation of M. If v does not have out-degree $\alpha(v)$ or $\alpha(v) - 1$ in X, then there is no extension of X to an α-orientation of M. The bound $2^{m-n+1} < 4^n$ follows by choosing A to be a spanning forest and applying Euler's formula. $\square$

A bound that improves Lemma 3.2 will be given in Proposition 3.5. The following lemma is needed for the proof.

Lemma 3.3. *Let M be a planar map with n vertices that has an independent set I_2 of n_2 vertices that have degree 2 in M. Then, M has at most $(3n - 6) - (n_2 - 1)$ edges.*

Proof. Consider a triangulation T extending M and let B be the set of additional edges, i.e., of edges of T that are not in M. If $n = 3$, then the conclusion of the lemma is true and we may thus assume $n > 3$ for the rest of the proof. Hence, there are no vertices of degree 2 in T, and every vertex of I_2 must be incident to at least one edge from B. If there is a vertex $v \in I_2$ that is incident to exactly one edge from B, then v and its incident edges can be deleted from I_2, from M and from T, whereby the result follows by induction. The last case is that all vertices of I_2 have at least two incident edges in B. Since every edge in B is incident to at most two vertices from I_2 it follows that $|I_2| \leq |B|$. Therefore,

$$|E(M)| = |E(T)| - |B| \leq |E(T)| - |I_2| = (3n-6) - n_2.$$

$\square$

Remark 3.4. It can be seen from the above proof that K_{2,n_2} plus the edge between the two vertices of degree n_2 is the unique graph to that only $n_2 - 1$ edges can be added. For every other graph at least n_2 edges can be added.

Proposition 3.5. *Let M be a planar map, $\alpha : V \to \mathbb{N}$, and $I = I_1 \cup I_2$ an independent set of M, where I_2 is the subset of vertices of I that have degree 2 in M. Then the number of α-orientations of M is at most*

$$2^{2n-4-|I_2|} \cdot \prod_{v \in I_1} \left(\frac{1}{2^{d(v)-1}} \binom{d(v)}{\alpha(v)} \right). \tag{3.1}$$

Proof. We may assume that M is connected and has minimum degree 2. Let M_i, for $i = 1, \dots c$, be the components of $M - I$. We claim that M has at most $(3n-6) - (c-1) - (|I_2| - 1)$ edges. Note that every component C of $M - I$ must be connected to some other component C' via a vertex $v \in I$ such that the edges vw and vw' with $w \in C$ and $w' \in C'$ form an angle at v. Since w and w' are in different connected components the edge ww' is not in M and we can add it without destroying planarity. We can add at least $c - 1$ edges not incident to I in this fashion. Thus, by Lemma 3.3 we have that $m + (c-1) \leq 3n - 6 - (|I_2| - 1)$.

Let S' be a spanning forest of $M - I$, and let S be obtained from S' by adding one edge incident to every $v \in I$. Then, S is a forest with $n - c$ edges. By Lemma 3.2, M has at most $2^{m-|S|}$ α-orientations and by Lemma 3.3

$$m - |S| \leq (3n-6) - (c-1) - (|I_2| - 1) - (n - c) = 2n - 4 - |I_2|.$$

For every vertex $v \in I_1$ there are $2^{d(v)-1}$ possible orientations of the edges of $M - S$ at v. Only the orientations with $\alpha(v)$ or $\alpha(v) - 1$ outgoing edges at v can potentially be completed to an α-orientation of M. Since I_1 is an independent set it follows that M has at most

$$2^{m-|S|} \cdot \prod_{v \in I_1} \frac{\binom{d(v)-1}{\alpha(v)} + \binom{d(v)-1}{\alpha(v)-1}}{2^{d(v)-1}} \leq 2^{2n-4-|I_2|} \cdot \prod_{v \in I_1} \frac{\binom{d(v)}{\alpha(v)}}{2^{d(v)-1}} \tag{3.2}$$

α-orientations. $\square$

Corollary 3.6. *Let M be a planar map and $\alpha : V \to \mathbb{N}$. Then, M has at most 3.73^n α-orientations.*

Proof. Since M is planar, the Four Color Theorem implies that it has an independent set I of size $|I| \geq n/4$. Let I_1, I_2 be as in Proposition 3.5. Note that for $d(v) \geq 3$

$$\frac{1}{2^{d(v)-1}}\binom{d(v)}{\alpha(v)} \leq \frac{1}{2^{d(v)-1}}\binom{d(v)}{\lfloor d(v)/2 \rfloor} \leq \frac{3}{4}. \tag{3.3}$$

Thus, the result follows from Proposition 3.5.

$$2^{2n-4-|I_2|}\left(\frac{3}{4}\right)^{|I_1|} \leq 2^{2n-4}\left(\frac{3}{4}\right)^{\frac{n}{4}} \leq 3.73^n.$$

□

The best lower bound for general α and M that we can prove, uses Eulerian orientations of the triangular grid, see Section 3.1.3.

3.1.2 Grid Graphs

Enumeration and counting of different combinatorial structures on grid graphs have received a lot of attention in the literature, see e.g. [9, 65, 21]. In Section 3.1.3 we present a family of graphs that have asymptotically at least 2.598^n Eulerian orientations. This family is closely related to the grid graph, and throughout this chapter we will use different relatives of the grid graph to obtain lower bounds. We collect the definitions of these related families in this section.

square grid $G_{k,\ell}$ The *square grid* or *grid graph* $G_{k,\ell}$ with k rows and ℓ columns is defined as follows. The vertex set is

$$V_{k,\ell} = \{(i,j) \mid 1 \leq i \leq k,\ 1 \leq j \leq \ell\}.$$

The edge set $E_{k,\ell} = E^H_{k,\ell} \cup E^V_{k,\ell}$ consists of horizontal edges

$$E^H_{k,\ell} = \Big\{\{(i,j),(i,j+1)\} \mid 1 \leq i \leq k,\ 1 \leq j \leq \ell - 1\Big\}$$

and vertical edges

$$E^V_{k,\ell} = \Big\{\{(i,j),(i+1,j)\} \mid 1 \leq i \leq k-1,\ 1 \leq j \leq \ell\Big\}.$$

V^R_i, V^C_j, E^C_j We denote the ith vertex row by $V^R_i = \{(i,j) \mid 1 \leq j \leq \ell\}$ and the jth vertex column by $V^C_j = \{(i,j) \mid 1 \leq i \leq k\}$. The jth edge column E^C_j is defined as

$E_j^C = \{\{(i,j),(i,j+1)\} \mid 1 \le i \le k\}$. The number of bipolar orientations of $G_{k,\ell}$ is the topic of Section 3.4.1.

The grid on the torus $G^T_{k,\ell}$ is obtained from $G_{k+1,\ell+1}$ as follows. We identify the vertices $(1,i)$ and $(k+1,i)$ as well as $(j,1)$ and $(j,\ell+1)$ for all i and j and delete parallel edges that these identifications create, see Figures 3.1 (a) and (b). Edges of the form $\{(i,1),(i,\ell)\}$ are called horizontal *wrap-around edges* while those of the form $\{(1,j),(k,j)\}$ are the vertical wrap-around edges. Note that $G_{k,\ell}$ can be obtained from $G^T_{k,\ell}$ by deleting the k horizontal and the ℓ vertical wrap-around edges.

$G^T_{k,\ell}$

wrap-around edge

Lieb [65] shows that $G^T_{k,\ell}$ has asymptotically $(8\sqrt{3}/9)^{k\ell}$ *Eulerian orientations*. His analysis involves the calculation of the dominant eigenvalue of a so-called transfer matrix. In Section 3.3 we also use this technique.

Eulerian orientation

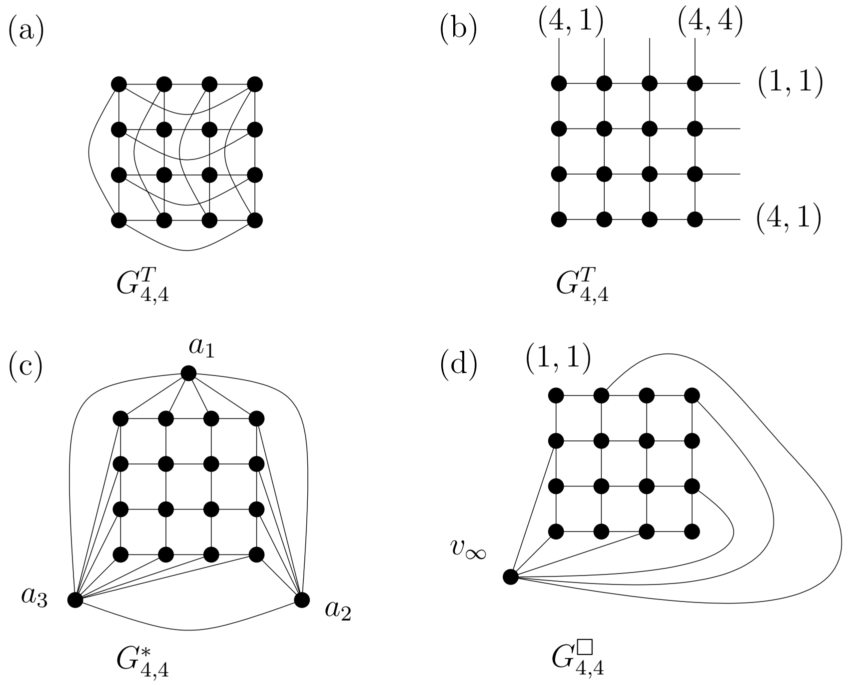

Figure 3.1. The graph $G^T_{4,4}$ and the planar maps $G^*_{4,4}$ and $G^\square_{4,4}$.

We consider the number of Schnyder woods on the augmented grid $G^*_{k,\ell}$ in Section 3.2.2. The augmented grid is obtained from $G_{k,\ell}$ by adding a triangle with vertices $\{a_1,a_2,a_3\}$ to the outer face, see Figure 3.1 (c). The triangle is connected to the boundary vertices of the grid as follows. The vertex a_1 is adjacent to all vertices of V_1^R, a_2 is adjacent the vertices from V_ℓ^C and a_3 to the vertices from $V_k^R \cup V_1^C$.

$G^*_{k,\ell}$

When we consider 2-orientations in Section 3.3 we use the quadrangulation $G^\square_{k,\ell}$, see Figure 3.1 (d). A quadrangulation is a planar map such that all faces have

$G^\square_{k,\ell}$

cardinality 4. This quadrangulation is obtained from the grid $G_{k,\ell}$ by adding a vertex v_∞ to the outer face that is adjacent to every other vertex of the boundary such that $(1,1)$ is not adjacent to v_∞. When k and ℓ are even, this graph is closely related to the torus grid $G^T_{k,\ell}$ which can be obtained from $G^\square_{k,\ell}$ by reassigning end vertices of edge as follows, see Figures 3.1 (a) and (d).

$$\begin{array}{llll} \{(1,j),v_\infty\} \to \{(1,j),(k,j)\}, & \{(k,j),v_\infty\} \to \{(k,j),(1,j)\} & \text{for} & 2 \le j \le \ell \\ \{(i,1),v_\infty\} \to \{(i,1),(i,\ell)\}, & \{(i,\ell),v_\infty\} \to \{(i,\ell),(i,1)\} & \text{for} & 2 \le i \le k \end{array}$$

Since k,ℓ are even, this does not create parallel edges and the resulting graph is $G^T_{k,\ell}$ minus the edges $e_1 = \{(1,1),(1,\ell)\}$ and $e_2 = \{(1,1),(k,1)\}$.

triangular grid $T_{k,\ell}$ In Sections 3.1.3 and 3.4.2 we use the *triangular grid* $T_{k,\ell}$ which we have already encountered in Section 1.3. It is obtained from $G_{k,\ell}$ by adding the diagonal edges $\{(i,j),(i-1,j+1)\}$ for $2 \le i \le k$ and $1 \le j \le \ell-1$, see Figure 3.2 (a). The $T^*_{k,\ell}$ augmented triangular grid $T^*_{k,\ell}$ which we need in Section 3.2.1 is obtained in the same way from $G^*_{k,\ell}$, see Figure 3.3.

The terms vertex row, vertex column, and edge column are used for the triangular grid analogously to the definition above for $G_{k,\ell}$.

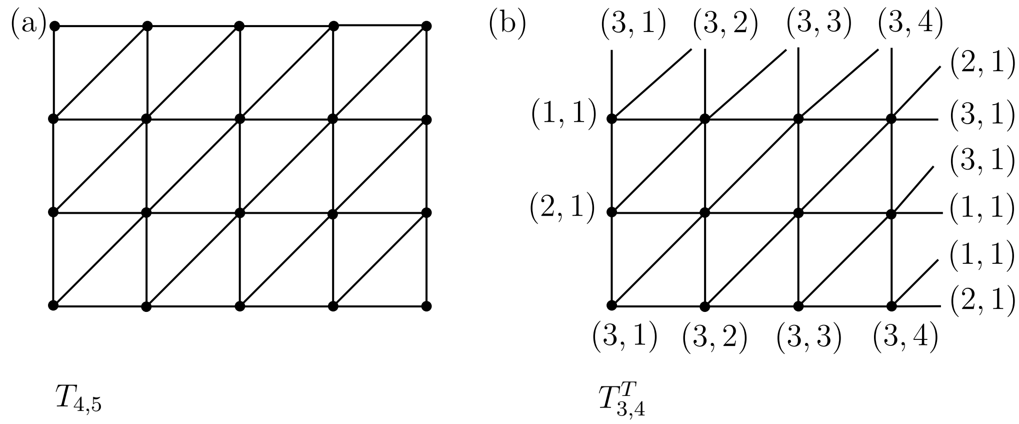

Figure 3.2. The triangular grid $T_{4,5}$, and the torus grid $T^T_{3,4}$. The labels indicate the vertices respectively the end vertices of the pending edges.

$T^T_{k,\ell}$ We also use the triangular grid on the torus $T^T_{k,\ell}$, see Figure 3.2 (b). We adopt the definition from [8], therefore it differs slightly from that of the square grid on the torus. More precisely, instead of identifying vertices $(i,\ell+1)$ and $(i,1)$, we identify vertices $(i,\ell+1)$ and $(i-1,1)$ (and $(1,\ell+1)$ with $(k,1)$) to obtain $T^T_{k,\ell}$ from $T_{k,\ell}$. Baxter [8] calls this boundary condition helical. The wrap-around edges are defined analogously to the square grid case.

Eulerian orientation In [8] Baxter has determined the asymptotic growth of the number of *Eulerian orientations* of $T^T_{k,\ell}$ as $k,\ell \to \infty$. Baxter's analysis uses similar techniques as Lieb's [65] and yields an asymptotic growth rate of $(3\sqrt{3}/2)^{k\ell}$.

3.1.3 A Lower Bound Using Eulerian Orientations

Let M be a planar map such that every $v \in V$ has even degree and let α be defined as $\alpha(v) = d(v)/2$, $\forall v \in V$. Such an α-orientation of M is better known as a *Eulerian orientation*. Eulerian orientations are exactly the orientations that maximize the binomial coefficients in Equation (3.1). The lower bound in the next theorem follows easily from a result by Baxter [8], as we explain below. It is the best lower bound that we have for $\max_{(M,\alpha)\in\mathcal{M}_n} r_\alpha(M)$, where $\mathcal{M}_n$ is the set of all planar maps with n vertices and no restrictions are imposed on α. *Eulerian orientation*

Theorem 3.7. *Let $\mathcal{E}(M)$ the set of Eulerian orientations of $M \in \mathcal{M}_n$. Then, the following bounds hold for n big enough.* $\mathcal{E}(M)$

$$2.59^n \leq (3\sqrt{3}/2)^{k\ell} \leq \max_{M\in\mathcal{M}_n} |\mathcal{E}(M)| \leq 3.73^n$$

Proof. The upper bound is the one from Corollary 3.6. For the lower bound consider the *triangular grid* on the torus $T^T_{k,\ell}$. As mentioned above, Baxter [8] was able to determine the exponential growth factor of Eulerian orientations of $T^T_{k,\ell}$ as $k, \ell \to \infty$. Baxter's analysis uses eigenvector calculations and yields an asymptotic growth rate of $(3\sqrt{3}/2)^{k\ell}$. From this graph a planar map $T^+_{k,\ell}$ can be constructed by introducing a new vertex v_∞ that subdivides every wrap-around edge. Thus, all crossings between wrap-around edges can be substituted by v_∞. As every Eulerian orientation of $T^T_{k,\ell}$ yields a Eulerian orientation of $T^+_{k,\ell}$ this graph has at least $(3\sqrt{3}/2)^{k\ell} \geq 2.598^{k\ell}$ Eulerian orientations for k, ℓ big enough. Note that $T^+_{k,\ell}$ has parallel edges. It can be transformed into a simple graph by subdividing $O(\sqrt{k\ell})$ edges. Thus, the claimed bound also holds for simple planar maps. □ *triangular grid*

3.2 The Number of Schnyder Woods

In this section we give asymptotic bounds for the maximum number of *Schnyder woods* on planar triangulations and 3-connected planar maps. We use definitions and facts from Chapter 1, in particular those from Section 1.3. *Schnyder wood*

We treat triangulations and 3-connected planar maps separately because the more direct bijection from Theorem 1.8 allows us to obtain a better upper bound for Schnyder woods on triangulations than for the general case. We also have a better lower bound for the general case of Schnyder woods on 3-connected planar maps than for the restriction to triangulations.

3.2.1 Schnyder Woods on Triangulations

Bonichon [12] found a bijection between Schnyder woods on triangulations with n vertices and pairs of non-crossing Dyck-paths. This bijection implies that there are $C_{n+2}C_n - C_{n+1}^2$ Schnyder woods on triangulations with n vertices. By C_n we denote the nth Catalan number $C_n = \binom{2n}{n}/(n+1)$. Hence, asymptotically there are about 16^n Schnyder woods on triangulations with n vertices. Tutte's classic result [95] yields that there are asymptotically about 9.48^n plane triangulations on n vertices. See [75] for a proof of Tutte's formula using Schnyder woods. The two results together imply that a triangulation with n vertices has on average about 1.68^n Schnyder woods. The next theorem is concerned with the maximum number of Schnyder woods on triangulations.

$\mathcal{S}(T)$ **Theorem 3.8.** *Let $\mathcal{T}_n$ denote the set of all plane triangulations with n vertices and $\mathcal{S}(T)$ the set of Schnyder woods of $T \in \mathcal{T}_n$. Then,*

$$2.37^n \leq \max_{T\in\mathcal{T}_n} |\mathcal{S}(T)| \leq 3.56^n.$$

Recall that the Schnyder woods of a triangulation are in bijection with its *3-orientations*, see Theorem 1.8. The upper bound follows from Proposition 3.5 by using that for $d(v) \geq 3$ it holds that $\binom{d(v)}{3} \cdot 2^{1-d(v)} \leq 5/8$.

3-orientation

triangular grid

For the proof of the lower bound we use the augmented *triangular grid* $T^*_{k,\ell}$. Figure 3.3 shows a canonical Schnyder wood on $T^*_{k,\ell}$ in which the vertical edges

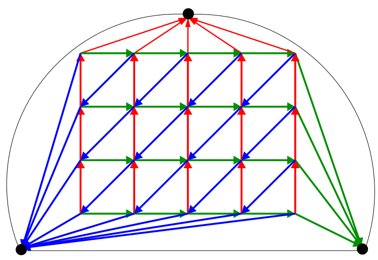

Figure 3.3. The augmented triangular grid $T^*_{4,5}$ with a canonical Schnyder wood.

are directed upwards, the horizontal edges to the right and the diagonal ones downwards. We work with α^*-orientations of $T_{k,\ell}$ instead of 3-orientations of $T^*_{k,\ell}$, where

$$\alpha^*(i,j) = \begin{cases} 3 & \text{if } 2 \leq i \leq k-1 \text{ and } 2 \leq j \leq \ell - 1 \\ 1 & \text{if } (i,j) \in \{(1,1),(1,\ell),(k,\ell)\} \\ 2 & \text{otherwise.} \end{cases}$$

For the sake of simplicity, we refer to α^*-orientations of $T_{k,\ell}$ as 3-orientations.

Intuitively, $T_{k,\ell}$ promises to be a good candidate for a lower bound because the canonical orientation shown in Figure 3.3 has many directed cycles. We formalize this intuition in the next proposition which we restrict to the case $k = \ell$ to keep the notation simple.

Proposition 3.9. *The graph $T^*_{k,k}$ has at least $2^{5(k-1)^2/4}$ Schnyder woods, and for k big enough we have*

$$2.37^{k^2+3} \leq |\mathcal{S}(T^*_{k,k})| \leq 2.599^{k^2+3}.$$

Proof. The face boundaries of the triangles of $T_{k,k}$ can be partitioned into two classes $\mathcal{C}$ and $\mathcal{C}'$ of directed cycles, such that each class has cardinality $(k-1)^2$ and no two cycles from the same class share an edge. Thus, a cycle $C' \in \mathcal{C}'$ shares an edge with three cycles from $\mathcal{C}$ if it does not share an edge with the outer face of $T_{k,k}$ and otherwise C' shares an edge with one or two cycles from $\mathcal{C}$.

For any subset D of $\mathcal{C}$ flipping all the cycles in D yields a 3-orientation of $T_{k,k}$, and we can encode this orientation as a 0-1-sequence of length $(k-1)^2$ that records which cycles have been flipped. After performing the flips of a given 0-1-sequence a, an inner cycle $C' \in \mathcal{C}'$ is directed if and only if either all or none of the three cycles sharing an edge with C' have been reversed. If $C' \in \mathcal{C}'$ is a boundary cycle, then it is directed if and only if none of the adjacent cycles from $\mathcal{C}$ has been reversed. Thus, the number of different cycle flip sets is bounded from below by *flip*

$$\sum_{a \in \{0,1\}^{(k-1)^2}} 2^{\sum_{C' \in \mathcal{C}'} X_{C'}(a)}.$$

Here $X_{C'}(a)$ is an indicator function that takes value 1 if C' is directed after performing the flips of a and 0 otherwise.

We now assume that every $a \in \{0,1\}^{(k-1)^2}$ is chosen uniformly at random. The expected value of the above function is then

$$\mathbb{E}[2^{\sum X_{C'}}] = \frac{1}{2^{(k-1)^2}} \sum_{a \in \{0,1\}^{(k-1)^2}} 2^{\sum_{C' \in \mathcal{C}'} X_{C'}(a)}.$$

Jensen's inequality $\mathbb{E}[\varphi(X)] \geq \varphi(\mathbb{E}[X])$ holds for a random variable X and a convex function φ. Using this we derive that *Jensen's inequality*

$$\mathbb{E}[2^{\sum X_{C'}}] \geq 2^{\mathbb{E}[\sum X_{C'}]} = 2^{\sum \mathbb{P}[C' \text{ flippable}]}.$$

The probability that $C' \in \mathcal{C}'$ is flippable is at least $1/4$. For a cycle C' that does not include a boundary edge, the probability depends only on the three cycles from $\mathcal{C}$ that share an edge with C'. Two out of the eight flip vectors for these three cycles make C' flippable. A similar reasoning applies if C' includes a boundary edge. Altogether this yields that

$$\sum_{a \in \{0,1\}^{(k-1)^2}} 2^{\sum_{C' \in \mathcal{C}'} X_{C'}(a)} \geq 2^{(k-1)^2} \cdot \mathbb{E}[2^{\sum X_{C'}}] \geq 2^{(k-1)^2} \cdot 2^{(k-1)^2/4}.$$

It remains to argue that different cycle flip sequences yield different Schnyder woods. The orientation of an edge is easily determined. The edge direction is reversed with respect to the canonical orientation if and only if exactly one of the two cycles on which it lies has been flipped. We can tell a flip sequence apart from its complement by comparing the boundary edges.

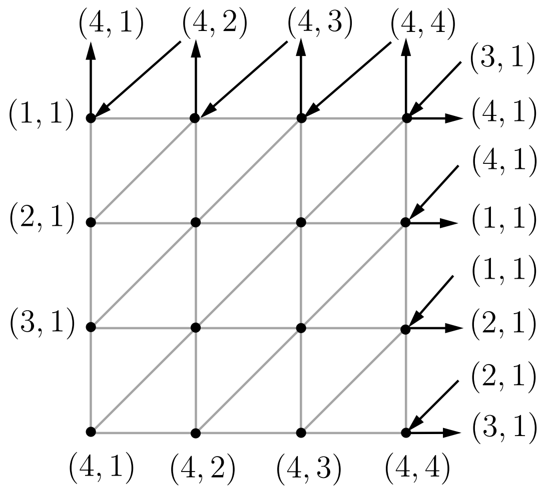

Figure 3.4. The graph $T_{4,4}$ with the additional edges simulating Baxter's boundary conditions for $T^T_{4,4}$.

For the upper bound we use Baxter's result for Eulerian orientations on the torus $T^T_{k,k}$, see Sections 3.1.2 and 3.1.3. Every 3-orientation of $T_{k,k}$ plus the wrap-around edges, oriented as shown in Figure 3.4, yields a Eulerian orientation of $T^T_{k,k}$. We deduce that $T^*_{k,k}$ has at most 2.599^n Schnyder woods. □

Remark 3.10. Let us briefly come back to the number of Eulerian orientations of $T^T_{k,\ell}$ which was mentioned in Sections 3.1.2 and 3.1.3 and in the above proof. There are only $2^{2(k+\ell)-1}$ different orientations of the wrap-around edges, one of them is shown in Figure 3.4. By the pigeon hole principle there is an orientation of these edges which can be extended to a Eulerian orientation of $T^T_{k,\ell}$ in asymptotically $(3\sqrt{3}/2)^{k\ell}$ ways. Thus, there are out-degree functions $\alpha_{k,\ell}$ for $T_{k,\ell}$ such that there are asymptotically $2.598^{k\ell}$ $\alpha_{k,\ell}$-orientations. Note, however, that directing all the wrap-around edges away from the vertex to which they are attached in Figure 3.3 induces a unique Eulerian orientation of $T_{k,\ell}$.

We have not been able to specify orientations of the wrap-around edges which allow to conclude that $T_{k,\ell}$ has $(3\sqrt{3}/2)^{k\ell}$ 3-orientations with these boundary conditions. In particular we have no proof that Baxter's result also gives a lower bound for the number of 3-orientations.

3.2.2 Schnyder Woods on the Grid and 3-Connected Planar Maps

In this section we discuss bounds on the number of Schnyder woods for all 3-connected planar maps. The lower bound comes from the grid. The upper bound for this case is much larger than the one for triangulations. This is due to the encoding of Schnyder woods by *α_S-orientations* on the *primal dual completion map* $\widetilde{M}$ which has more vertices than M. We summarize the results of this section in the following theorem.

α_S-orientation

primal dual completion map

Theorem 3.11. *Let $\mathcal{M}_n^3$ be the set of 3-connected planar maps with n vertices and $\mathcal{S}(M)$ denote the set of Schnyder woods of $M \in \mathcal{M}_n^3$. Then*

$$3.209^n \leq \max_{M \in \mathcal{M}_n^3} |\mathcal{S}(M)| \leq 8^n.$$

The example used for the proof of the lower bound is the augmented *square grid* $G^*_{k,\ell}$.

square grid

Theorem 3.12. *The number of Schnyder woods of the augmented grid $G^*_{k,\ell}$ is asymptotically $|\mathcal{S}(G^*_{k,\ell})| \approx 3.209^{k\ell}$.*

Proof. The graph induced by the non-rigid edges in the primal dual completion map $\widetilde{G}^*_{k,\ell}$ of $G^*_{k,\ell}$ is $G_{2k-1,2\ell-1} - (2k-1,1)$, see Figure 3.5. This is a square grid of roughly twice the size as the original and with the lower left corner removed. The rigid edges can be identified using the fact that $\alpha_S(v_\infty) = 0$, and deleting them induces α'_S on $G_{2k-1,2\ell-1} - (2k-1,1)$. The new α'_S only differs from α_S for vertices that are incident to an outgoing rigid edge, and it turns out that $\alpha'_S(v) = d(v) - 1$ for all primal or dual vertices and $\alpha'_S(v) = 1$ for all edge-vertices of $G_{2k-1,2\ell-1} - (2k-1,1)$. Thus, a bijection between α'_S-orientations and perfect matchings of $G_{2k-1,2k-\ell} - (2k-1,1)$ is established by identifying matching edges with edges directed away from edge-vertices. The closed form expression for the number of perfect matchings of $G_{2k-1,2k-\ell} - (2k-1,1)$ is known (see [62]) to be

$$\prod_{i=1}^{k}\prod_{j=1}^{\ell}\left(4 - 2\cos\frac{\pi i}{k} - 2\cos\frac{\pi j}{\ell}\right).$$

The number of perfect matchings of $G_{2k-1,2\ell-1} - (2k-1,1)$ is sandwiched between that of $G_{2k-2,2\ell-2}$ and that of $G_{2k,2\ell}$. Therefore the asymptotic behavior is the same and in [69], the limit of the number of perfect matchings of $G_{2k,2\ell}$, denoted as $\Phi(2k,2\ell)$, is calculated to be

$$\lim_{k,\ell\to\infty} \frac{\log\Phi(2k,2\ell)}{2k\cdot 2\ell} = \frac{\log 2}{2} + \frac{1}{4\pi^2}\int_0^\pi\int_0^\pi \log(\cos^2(x)+\cos^2(y))dxdy \approx 0.29.$$

This implies that $G^*_{k,\ell}$ has asymptotically $e^{4\cdot 0.29\cdot k\ell} \approx 3.209^{k\ell}$ Schnyder woods. $\square$

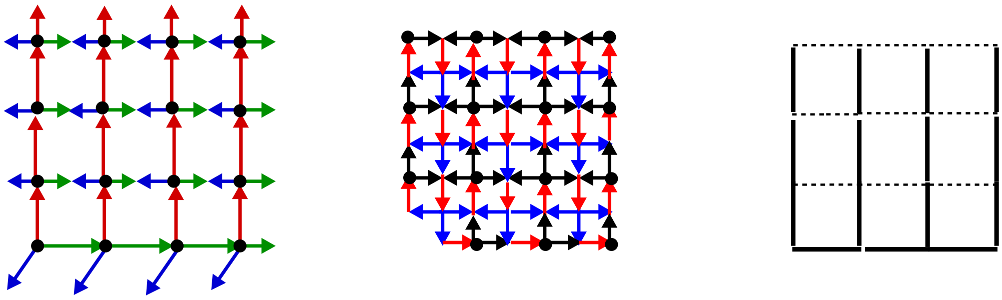

Figure 3.5. A Schnyder wood on $G^*_{4,4}$, the reduced primal dual completion map $G_{7,7} - (7,1)$ with the corresponding α'_S-orientation and the associated spanning tree.

Remark 3.13. In [89] Temperley describes a bijection between spanning trees of $G_{k,\ell}$ and perfect matchings of $G_{2k-1,2\ell-1} - (2k-1,1)$. Thus, Schnyder woods of $G^*_{k,\ell}$ are in bijection with spanning trees of $G_{k,\ell}$, see Figure 3.5. This bijection can be read off directly from the Schnyder wood: the unidirected edges not incident to a special vertex form exactly the related spanning tree. Encoding both, the Schnyder woods and the spanning trees, as α-orientations also gives an immediate proof of this bijection.

Proof of Theorem 3.11. It remains to prove the upper bound stated in Theorem 3.11. For this we use the upper bound for Schnyder woods on plane triangulations, see Theorem 3.8. We define a triangulation T_M such that there is an injective mapping of the Schnyder woods of M to the Schnyder woods of T_M. We use the generic structure of the faces of a Schnyder wood, see Lemma 1.11.

We may assume that M has a triangular outer face. Otherwise we may continue to work with a map M' that is obtained from M by triangulating the outer face with the help of one additional vertex.

The triangulation T_M is obtained from M by adding a vertex v_F to every face F of M with $|F| \geq 4$, see Figure 3.6. A vertex v_F is adjacent to all the vertices of F. A Schnyder wood of M can be mapped to a Schnyder wood of T_M using the generic structure of the bounded faces as shown in Figure 3.6. The green-blue non-special edges of F become green unidirected. Their blue parts are substituted by unidirected blue edges pointing from their original start-vertex towards v_F. Similarly the blue-red non-special edges become blue unidirected and the red-green ones red unidirected. Three of the edges incident to v_F are still undirected at this point. We orient them away from v_F and assign colors in accordance with Rule (W3), see Section 1.1.

Let two different Schnyder woods be given that have different directions or colors on an edge e. That the map is injective can be verified by comparing the edges on the boundary of the two triangles on which the edge e lies in T_M.

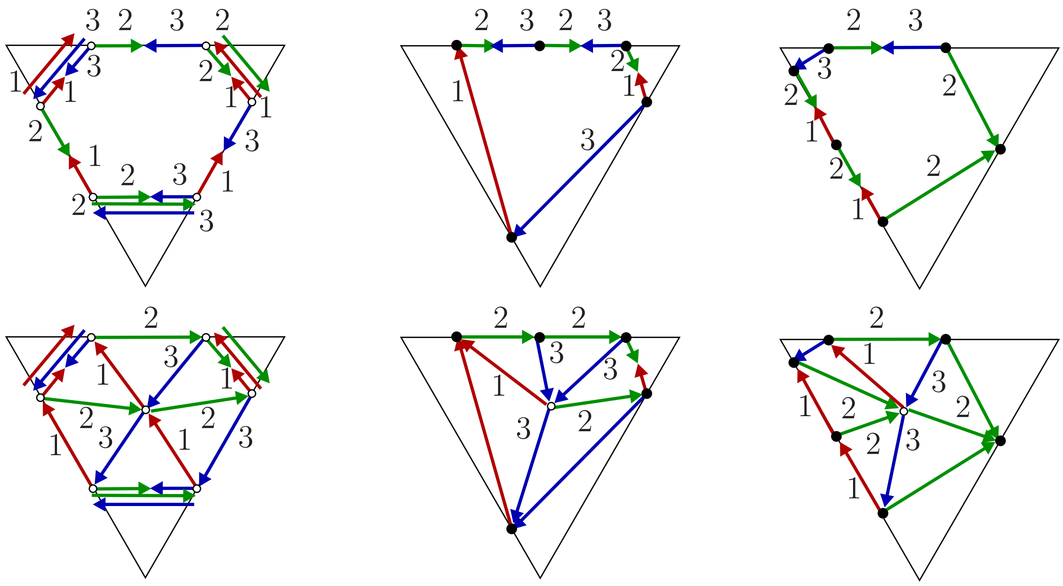

Figure 3.6. A Schnyder wood on a map M^σ induces a Schnyder wood on T_{M^σ}. The three special edges of a face are those that do not lie on the black triangle.

Thus, it suffices to bound the number of Schnyder woods of T_M in order to prove Theorem 3.11. We do this by specializing Proposition 3.5. We denote the set of vertices of T_M that correspond to faces of size 4 in M by F_4 and its size by f_4 and similarly $F_{\geq 5}$ and $f_{\geq 5}$ are defined. Note that $I = F_4 \cup F_{\geq 5}$ is an independent set and that T_M has a spanning tree in which all the vertices from I are leaves. Let n_T denote the number of vertices of T_M. Then, T_M has at most

$$\begin{aligned} 2^{3n_T-6-(n_T-1)} \cdot \prod_{v \in I} \left(\frac{\binom{d(v)}{3}}{2^{d(v)-1}} \right) &\leq 4^{n+f_4+f_{\geq 5}} \cdot \left(\frac{1}{2}\right)^{f_4} \cdot \left(\frac{5}{8}\right)^{f_{\geq 5}} \\ &= 4^n \cdot 2^{f_4} \cdot \left(\frac{5}{2}\right)^{f_{\geq 5}} \end{aligned} \tag{3.4}$$

Schnyder woods. Note that $n + f_4 + f_{\geq 5} + f_4 + 2f_{\geq 5} \leq m + 1 + f_4 + 2f_{\geq 5} \leq 3n - 5$ which implies that $f_4 + \frac{3}{2} f_{\geq 5} \leq n$. Maximizing the right hand side of (3.4) under this condition yields that the maximum 8^n is attained when $f_4 = n$. Thus, M has no more than 8^n Schnyder woods. □

The proof of the lower bound 3.209^n uses knowledge about the number of perfect matchings of the square grid which is obtained using non-combinatorial methods. Therefore, we complement this bound with a result for another graph family that uses a straight-forward analysis, but still yields that these graphs have more Schnyder woods than the triangular grid, see Proposition 3.9.

The graph we consider is the *filled hexagonal grid* $H_{k,\ell}$, see Figure 3.7 (a).

filled hexagonal grid $H_{k,\ell}$

Neglecting boundary effects, the hexagonal grid has twice as many vertices as hexagons. This can be seen by associating with every hexagon the vertices of its northwestern edge. Thus, neglecting boundary effects, the filled hexagonal grid has five vertices per hexagon. The boundary effects will not hurt our analysis because $H_{k,\ell}$ has only $2(k+\ell)$ boundary vertices but $5 \cdot k\ell + 2(k+l)$ vertices in total.

Proposition 3.14. *For k, ℓ big enough the number of Schnyder woods of the filled hexagonal grid $H_{k,\ell}$ can be bounded as follows.*

$$2.63^n \leq |\mathcal{S}(H_{k,\ell})| \leq 6.07^n$$

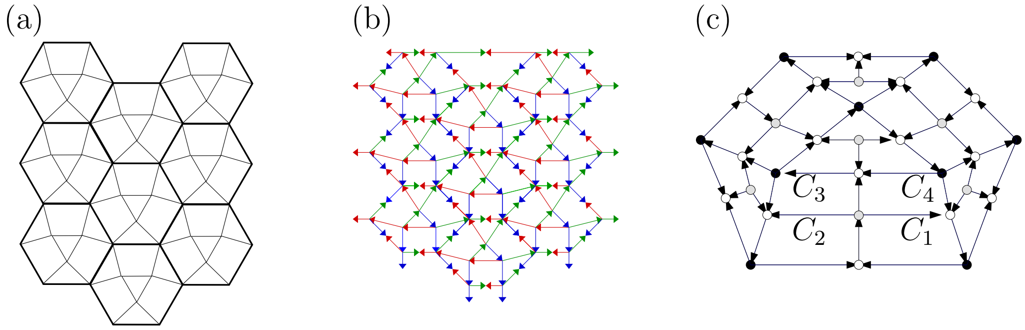

Figure 3.7. The filled hexagonal grid $H_{3,3}$, a Schnyder wood on this grid and the primal-dual suspension of a hexagonal building block of $H_{k,\ell}$. Primal vertices are black, face vertices grey and edge-vertices white.

Proof. Using Theorem 1.10 it suffices to count how many different α_S-orientations a filled hexagon has. Figure 3.7 (c) shows an α_S-orientation of a filled hexagon. Note that this orientation is feasible on the boundary when we glue together the filled hexagons to a grid $H_{k,\ell}$ and add a triangle of three special vertices around the grid, see Figure 3.7 (b). We flip only boundary edges of a hexagon that belong to a 4-face in this hexagon. As these edges belong to a triangle in the hexagon on their other side, the cycle flips in any two filled hexagons can be performed independently.

Let us now count how many orientations a filled hexagon admits, see Figure 3.7 (c) for the definition of the cycles C_1, C_2, C_3 and C_4. If the 6-cycle induced by the central triangle of a filled hexagon is directed as shown in Figure 3.7 (c), then we can flip either C_1 or C_2 and if C_2 is flipped, C_3 can be flipped as well. This yields 4^3 orientations, as the situation is the same at the other two 4-faces of the hexagon. If the 6-cycle is flipped the same calculation can be done with C_3 replaced by C_4. This makes a total of $2 \cdot 4^3 = 128$ orientations per filled hexagon. That is, there are at least $128^{k \cdot \ell} \geq 2.639^{5 \cdot k \cdot \ell}$ orientations of $H_{k,\ell}$.

We start the proof of the upper bound by collecting some statistics about $H_{k,\ell}$, where we neglect summands of the form $\kappa \cdot k$ and $\kappa \cdot \ell$ since they are not relevant asymptotically. As mentioned above, $H_{k,\ell}$ has $n = 5 \cdot k\ell$ interior vertices, $12 \cdot k\ell$ edges and $7 \cdot k\ell$ faces. Thus, the primal-dual completion has $48 \cdot k\ell$ edges.

There is no choice for the orientation of the $3 \cdot 4/7 \cdot f = 12 \cdot k\ell$ edges incident to the face vertices of triangles. We can choose a spanning tree T on the remaining $5 \cdot k\ell + 12 \cdot k\ell + 3 \cdot k\ell$ vertices such that all face vertices are leafs and proceed as in the proof of Proposition 3.5, using that we know the exact number of edges. Since in the independent set of the remaining face vertices, all of them have degree 4 and required out-degree 3, they contribute a factor of $\binom{4}{3}2^{-3} = 1/2$ each. Thus, there are at most $2^{(48-12-20)k\ell} \cdot 2^{-3 \cdot k\ell} = 2^{13 \cdot k\ell} \leq 6.07^n$ Schnyder woods on $H_{k,\ell}$. $\square$

3.3 The Number of 2-Orientations

Felsner et al. [43] present a theory of 2-orientations of plane quadrangulations which shows many similarities with Schnyder woods of triangulations. Recall that a quadrangulation is a planar map such that all faces have cardinality 4. A *2-orientation of a quadrangulation* Q is an orientation of the edges such that all vertices but two non-adjacent ones on the outer face have out-degree 2. *2-orientation of a quadrangulation*

In [42] it is shown that 2-orientations of quadrangulations with $f-1$ inner quadrangles are counted by the Baxter-number B_f. Hence the number of 2-orientations on quadrangulations with n vertices is asymptotically 8^n, since $f = n - 2$ for quadrangulations. Tutte gave an explicit formula for the number of rooted quadrangulations. A bijective proof of Tutte's formula is contained in the thesis of Fusy [51]. The formula implies that asymptotically there are about 6.75^n quadrangulations on n vertices. The two results together yield that a quadrangulation with n vertices has on average about 1.19^n 2-orientations. Figure 3.8 (a) shows a 2-orientation of $G^{\square}_{6,6}$, and we will show that the $G^{\square}_{k,\ell}$ have considerably more 2-orientations than the average quadrangulation.

We now give a lower bound for the number of 2-orientations of the quadrangulation $G^{\square}_{k,\ell}$ that is obtained from the *square grid*. The proof method via transfer matrices and eigenvalue estimates comes from Calkin and Wilf [21]. There it is used for asymptotic enumeration of independent sets of the grid graph. Let $\mathcal{Z}(Q)$ denote the set of all 2-orientations of a quadrangulation Q with fixed sinks. *square grid* *$\mathcal{Z}(Q)$*

Proposition 3.15. *For k, ℓ big enough the number of 2-orientations of $G^{\square}_{k,\ell}$ can be bounded as follows.*

$$1.537^{k\ell} \leq |\mathcal{Z}(G^{\square}_{k,\ell})| \leq (8 \cdot \sqrt{3}/9)^{k\ell} \leq 1.5397^{k\ell}$$

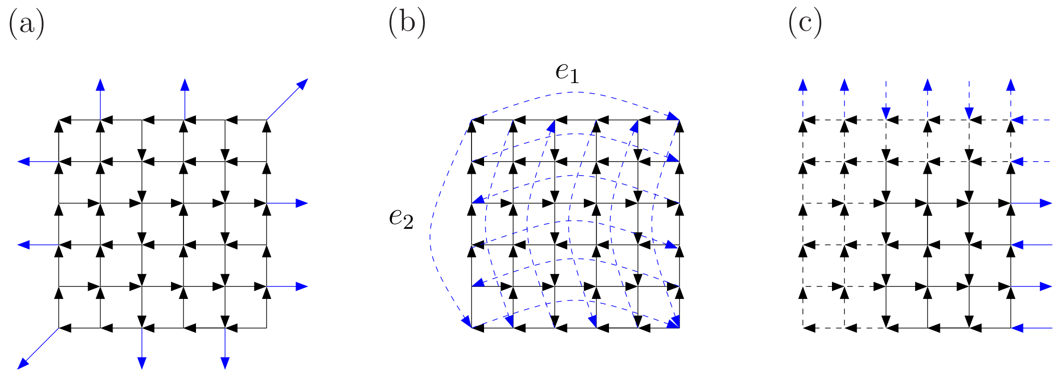

Figure 3.8. A 2-orientation of $G^{\square}_{6,6}$, the corresponding Eulerian orientation X of $G^T_{6,6}$ and an alternating orientation of $G^T_{4,4}$ that can be extended to X.

Proof. We consider 2-orientations of $G^{\square}_{k,\ell}$ with sinks $(1,1)$ and v_∞, see Figure 3.8 (a). These 2-orientations induce Eulerian orientations of $G^T_{k,\ell}$ when k and ℓ are even. The wrap-around edges inherit the direction of the respective edges incident to v_∞ (see Section 3.1.2) and e_1, e_2 are directed away from $(1,1)$, see Figure 3.8 (b). Therefore $G^{\square}_{k,\ell}$ has at most as many 2-orientations as $G^T_{k,\ell}$ has Eulerian orientations, which implies the claimed upper bound.

Conversely a Eulerian orientation of $G^T_{k,\ell}$ in which the wrap-around edges have these prescribed orientations induces a 2-orientation of $G^{\square}_{k,\ell}$. Such Eulerian orientations are called almost alternating orientations in the sequel, see Figures 3.8 (b) and (c). Proving a lower bound for the number of almost alternating Eulerian orientations yields a lower bound for the number of 2-orientations of $G^{\square}_{k,\ell}$.

For the sake of simplicity we will work with alternating orientations of $G^T_{k-2,\ell-2}$ instead of almost alternating ones of $G^T_{k,\ell}$. In these Eulerian orientations of $G^T_{k-2,\ell-2}$, the wrap-around edges are directed alternatingly up and down respectively left and right. It is easy to see that this gives a lower bound for the number of almost alternating orientations of $G^T_{k,\ell}$. The solid edges in Figure 3.8 (c) show an alternating orientation and the dashed ones how it can be augmented to an almost alternating orientation. Since we are interested in an asymptotic lower bound, there is no difference in counting alternating orientations of $G^T_{k-2,\ell-2}$ and $G^T_{k,\ell}$ from our point of view. Therefore, we will continue working with alternating orientations of $G^T_{k,\ell}$ to keep the notation simple.

Consider a vertex column V^C_j of $G^T_{k,\ell}$ and the edge columns E^C_{j-1} and E^C_j. Let X_1 and X_2 be orientations of E^C_{j-1} respectively E^C_j. Let $\delta(X_1, X_2) = 1$ if and only if the edges induced by V^C_j can be oriented such that all the vertices of V^C_j have out-degree 2. Let $\delta_U(X_1, X_2) = 1$ respectively $\delta_D(X_1, X_2) = 1$ if and only if $\delta(X_1, X_2) = 1$ and the wrap-around edge induced by V^C_j is directed upwards respectively downwards.

Note that

$$\delta_U(X_1,\ X_1) = 1 = \delta_D(X_1,\ X_1)$$

and

$$\delta_U(X_1,\ X_2) = 1 \Longleftrightarrow \delta_D(X_2,\ X_1) = 1.$$

We define two transfer matrices $T_U(2k)$ and $T_D(2k)$. These are square 0-1-matrices with the rows and columns indexed by the $\binom{2k}{k}$ orientations of an edge column of size $2k$, that have k edges directed to the right. The transfer matrices are defined by $(T_U(2k))_{X_1,X_2} = \delta_U(X_1,\ X_2)$ and $(T_D(2k))_{X_1,X_2} = \delta_D(X_1,\ X_2)$. Hence $T_U(2k) = T_D(2k)^T$ and $T_{2k} = T_U(2k) \cdot T_D(2k)$ is a real symmetric non-negative matrix with positive diagonal entries, see Figure 3.9.

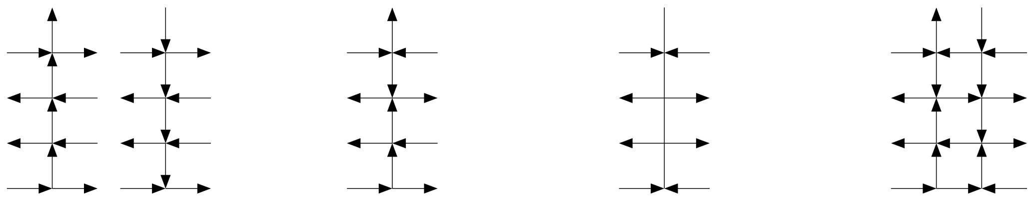

$(T_U(4))_{X_1,X_1} = 1$ $(T_D(4))_{X_1,X_1} = 1$ $(T_U(4))_{X_1,X_2} = 1$ $(T_U(4))_{X_1,X_3} = 0$ $(T_D(4))_{X_1,X_3} = 0$ $(T_4)_{X_1,X_3} \geq 1$

Figure 3.9. An illustration of the different transfer matrices.

From the combinatorial interpretation, it can be seen that T_{2k} is primitive, that is there is an integer $\ell \geq 1$ such that all entries of T_{2k}^{ℓ} are positive and thus the Perron-Frobenius Theorem can be applied, see [57]. Hence, T_{2k} has a unique eigenvalue Λ_{2k} with largest absolute value, its eigenspace is 1-dimensional and the corresponding eigenvector is positive.

Let X_A be one of the two edge column orientations that have alternating edge directions and e_A the vector of dimension $\binom{2k}{k}$ that has all entries 0 except the one that stands for X_A, which is 1. The number $c_A(2k, 2\ell)$ of alternating orientations of $G^T_{2k,2\ell}$ is $\left(T_{2k}^{\ell}\right)_{X_A,X_A} = \langle e_A, T_{2k}^{\ell} e_A\rangle$. Since the eigenvector belonging to Λ_k is positive, it is not orthogonal to any column of T_{2k}, and we obtain

$$\lim_{\ell\to\infty} c_A(2k, 2\ell)^{1/\ell} = \lim_{\ell\to\infty} \left(\left(T_{2k}^{\ell}\right)_{X_A,X_A}\right)^{1/\ell} = \Lambda_{2k}.$$

The last equality is justified by an argument known as the power method. It follows from [65] that the limit $\lim_{k\to\infty} \Lambda_{2k}^{1/k}$ exists, but for the sake of completeness we provide an argument from [21]. We use that $\Lambda_{2k}^{p} \geq \langle v, T_{2k}^{p} v\rangle / \langle v, v\rangle$ for any vector

v and that $\langle e_A, T_{2k}^p e_A\rangle = \langle e_A, T_{2p}^k e_A\rangle$ since both expressions count the number of alternating orientations of $G^T_{2k,2p}$.

$$\left(\Lambda_{2k}^{1/k}\right)^p = \left(\Lambda_{2k}^p\right)^{1/k} \geq \left(\langle e_A, T_{2k}^p e_A\rangle\right)^{1/k} = \left(\langle e_A, T_{2p}^k e_A\rangle\right)^{1/k}$$

Taking limits with respect to k on both sides yields

$$\left(\liminf_{k\to\infty} \Lambda_{2k}^{1/k}\right)^p \geq \liminf_{k\to\infty} \left(\langle e_A, T_{2p}^k e_A\rangle\right)^{1/k} = \Lambda_{2p},$$

which implies $\liminf\limits_{k\to\infty} \Lambda_{2k}^{1/k} \geq \limsup\limits_{p\to\infty} \Lambda_{2p}^{1/p}$. It follows that $\lim\limits_{k\to\infty} \Lambda_{2k}^{1/k}$ exists. Similar arguments as above yield the following.

$$\Lambda_{2k}^p \geq \frac{\langle e_A T_{2k}^q, T_{2k}^p T_{2k}^q e_A\rangle}{\langle T_{2k}^q e_A, T_{2k}^q e_A\rangle} = \frac{\langle e_A, T_{2k}^{p+2q} e_A\rangle}{\langle e_A, T_{2k}^{2q} e_A\rangle} = \frac{\langle e_A, T_{2p+4q}^k e_A\rangle}{\langle e_A, T_{4q}^k e_A\rangle}$$

Taking limits with respect to k on both sides yields

$$\lim_{k\to\infty} \Lambda_{2k}^{1/k} \geq \left(\frac{\Lambda_{4q+2p}}{\Lambda_{4q}}\right)^{1/p}.$$

We are interested in $\lim\limits_{k\to\infty}\lim\limits_{\ell\to\infty} c_A(2k,2\ell)^{1/4k\ell} = \lim\limits_{k\to\infty} \Lambda_{2k}^{1/4k}$ since $4k\ell$ is the number of vertices of $G^T_{2k,2\ell}$. Using a *Mathematica* program we have computed Λ_{10} and Λ_8 with the result that

$$\left(\frac{\Lambda_{10}}{\Lambda_8}\right)^{1/4} \geq \left(\frac{2335.8714}{418.2717}\right)^{1/4} \geq 1.537.$$

$\square$

Remark 3.16. We return to the correspondence between 2-orientations of $G^{\square}_{k,\ell}$ and Eulerian orientations of $G^T_{k,\ell}$, that was mentioned at the beginning of the last proof. By the pigeon hole principle, there must be a sequence of orientations $X_{k,\ell}$ of the wrap-around edges that extends asymptotically to $(8\cdot\sqrt{3}/9)^{k\ell}$ Eulerian orientations of $G^T_{k,\ell}$. This implies that for k,ℓ big enough, there is an $\alpha_{k,\ell}$ on $G_{k,\ell}$ such that there are $(8\cdot\sqrt{3}/9)^{k\ell}$ $\alpha_{k,\ell}$-orientations of $G_{k,\ell}$. This $\alpha_{k,\ell}$ satisfies $\alpha_{k,\ell}(v) = 2$ for every inner vertex v and $\alpha_{k,\ell}(w) \in \{0,1,2\}$ for every boundary vertex w. We call α-orientations that have $\alpha(v) = 2$ for every inner vertex v *inner 2-orientations*. *inner 2-orientation*

We think that $G^{\square}_{k,\ell}$ has asymptotically $(8\cdot\sqrt{3}/9)^{k\ell}$ 2-orientations. But we were not able to show this, see also Remark 3.10.

Theorem 3.17. *Let $\mathcal{Q}_n$ denote the set of all plane quadrangulations with n vertices and $\mathcal{Z}(Q)$ the set of 2-orientations of $Q \in \mathcal{Q}_n$. Then for n big enough*

$$1.53^n \leq \max_{Q \in \mathcal{Q}_n} |\mathcal{Z}(Q)| \leq 1.91^n.$$

Proof. The lower bound is that from Proposition 3.15. An upper bound of 2^n follows immediately from Lemma 3.2. Note that we may assume that Q does not have vertices of degree 2, because their incident edges would be rigid. We use Theorem 4.8 to conclude that Q has a spanning tree T with at least $n/3$ leaves. As Q is bipartite, T has a set I of at least $n/6$ leafs that is an independent set of Q. As in Proposition 3.5, this yields that there are at most $2^n \cdot (3/4)^{n/6} \leq 1.91^n$ 2-orientations of Q. □

3.4 The Number of Bipolar Orientations

We first give an overview of facts about bipolar orientations that we need in this section. A good starting point for further reading about bipolar orientations is [32].

Let G be a connected graph and s, t two distinguished vertices of G. An orientation X of the edges of G is a *bipolar orientation* of G if it is acyclic, s is the only vertex without incoming edges, and t is the only vertex without outgoing edges. We call s and t the source respectively sink of X. There are many equivalent definitions of bipolar orientations, see [32]. The following characterization of plane bipolar orientations will be useful to keep some proofs in the sequel short. When working with bipolar orientations, we always assume that s and t lie on the outer face. *bipolar orientation*

Proposition 3.18. *An orientation X of a planar map M with two special outer vertices s and t is a bipolar orientation if and only if it has the following properties.*

(1) Every vertex other than the source s and the sink t has incoming as well as outgoing edges.

(2) There is no directed facial cycle.

Furthermore, the following stronger versions of the above properties hold for every bipolar orientation.

(1') At every vertex other than the source and the sink, the incoming and outgoing edges form two non-empty bundles of consecutive edges.

(2') The boundary of every face has exactly one sink and one source, i.e. consists of two directed paths.

We omit the proof that Properties (1) and (2) imply that X is a bipolar orientation. The proof that every bipolar orientation has Properties (1′) and (2′) (and thus Properties (1) and (2) as well) can be found in [98, 88].

Given a planar map M, two bipolar orientations of M can have different out-degree sequences as the example from Figure 3.10 shows.

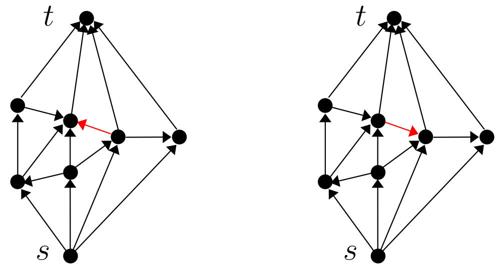

Figure 3.10. Two bipolar orientations of the same graph with different out-degree sequences.

angle graph $\widehat{M}$

Nevertheless, the bipolar orientations of a map M are in bijection with α-orientations of the *angle graph* $\widehat{M}$ of M. Let $\mathcal{F}$ be the set of faces of M. The angle graph $\widehat{M}$ is the bipartite graph on the vertex set $V \cup \mathcal{F}$ where two vertices $v \in V$ and $F \in \mathcal{F}$ are adjacent if and only if v lies on the boundary of F in M. The following theorem is due to Rosenstiehl [80]. A proof can also be found in [32], where the α-orientation of $\widehat{M}$ comes in the disguise of an "e-angle colouration".

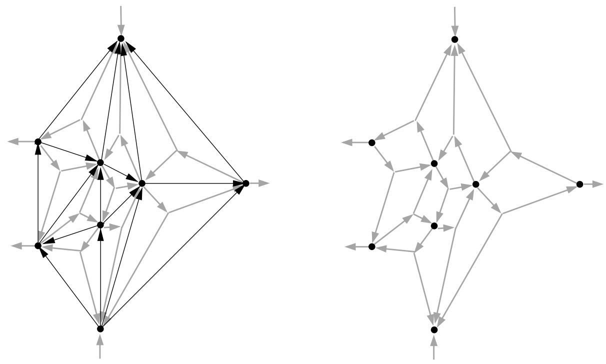

Figure 3.11. A bipolar orientation and the corresponding $\hat{\alpha}$-orientation of the angle graph. The vertex for the unbounded face of the angle graph is omitted.

Theorem 3.19. *Let M be a planar map and $\widehat{M}$ its angle graph. Let $\hat{\alpha} : V \cup \mathcal{F} \to \mathbb{N}$ be defined as follows. All $F \in \mathcal{F}$ and $v \in V \setminus \{s, t\}$ have $\hat{\alpha}(F) = 2$ respectively $\hat{\alpha}(v) = 2$. The source s and sink t have $\hat{\alpha}(s) = \hat{\alpha}(t) = 0$. Then, the bipolar orientations of M are in bijection with the $\hat{\alpha}$-orientations of $\widehat{M}$.*

Figure 3.11 illustrates Theorem 3.19. Below in Theorem 3.31 we give another encoding of bipolar orientations which will turn out to be useful for approximate counting.

Note that the angle graph $\widehat{M}$ is a quadrangulation. If the edge $\{s,t\}$ is in the planar map M, then the $\hat{\alpha}$-orientations are the same as the 2-orientations defined in Section 3.3. Since bipolar orientations and 2-orientations of quadrangulations are in bijection, we explain now why none of the bounds from Theorems 3.17 and 3.23 are redundant. A triangulation with n vertices has an angle graph with roughly $3n$ vertices. Hence the upper bound of 1.91^{3n}, that Theorem 3.17 yields for the number of bipolar orientation, is worse than the upper bound 3.97^n from Theorem 3.23. Conversely, every quadrangulation Q with n vertices is the angle graph of a map Q'. One of the partition classes of Q is the vertex set of Q' and two vertices of Q' are adjacent if and only if they lie on a common 4-face of Q. This might yield a multi-graph if Q has degree 2 vertices. But we may neglect this, since parallel edges must have the same direction in every bipolar orientation. One of the partition classes of Q has size at most $n/2$, and thus the upper bound from Theorem 3.23 yields that Q has at most $3.97^{n/2}$ 2-orientations, which is worse than the bound 1.91^n from Theorem 3.17. The grid graphs $G^{\square}_{k,\ell}$, which have asymptotically 1.53^n 2-orientations are angle graphs of graphs with roughly $n/2 =: n'$ vertices. Therefore this yields only an example with $1.53^{2n'}$ bipolar orientations, which is far away from the lower bound 2.91^n given in Theorem 3.23. Conversely, the triangular grid $T_{k,\ell}$, which has at least 2.91^n bipolar orientations, has an angle graph with roughly $3n = n'$ vertices. This yields a quadrangulation with $2.91^{n'/3}$ 2-orientations which is worse than the bound 1.53^n for the number of 2-orientations that we obtained in Proposition 3.15.

3.4.1 Bipolar Orientations of the Grid

We now turn to analyzing the number of bipolar orientations of $G_{k,\ell}$, with source $(1,1)$ and sink (k,ℓ) if k is odd and sink $(k,1)$ if k is even. For the proof of the Theorem, we need sparse sequences. A *sparse sequence* is a 0-1-sequences without consecutive 1s and it is well known that there are F_{n+2} such sequences of length n, where F_{n+2} denotes the $(n+2)$th *Fibonacci number*. Let $\mathcal{B}(M)$ denote the set of bipolar orientations of the map M.

sparse sequence F_{n+2} *Fibonacci number* $\mathcal{B}(M)$

Theorem 3.20. *For k,ℓ big enough, the number of bipolar orientations of the square grid $G_{k,\ell}$ is bounded by*

square grid

$$2.18^{k\ell} \leq |\mathcal{B}(G_{k,\ell})| \leq 2.619^{k\ell}.$$

Proof. We first prove the lower bound with an argument using directed cycles in a canonical orientation, as in Proposition 3.9. Therefore we do not spell out all the details of the proof but only sketch it. We work on the angle graph $\hat{G}_{k,\ell}$ and use the bijection from Theorem 3.19. The graph $\hat{G}_{k,\ell}$ has $2k\ell - 3(k+\ell)+4$ square faces not incident to v_∞. Figure 3.12 shows the angle graph $\hat{G}_{4,5}$. All edges that are dotted in Figure 3.12 (b) are rigid, just like the four edges which are adjacent to a degree 2 vertex. Therefore, we may neglect all these edges in the rest of the proof and work with graphs like the one in Figure 3.12 (c).

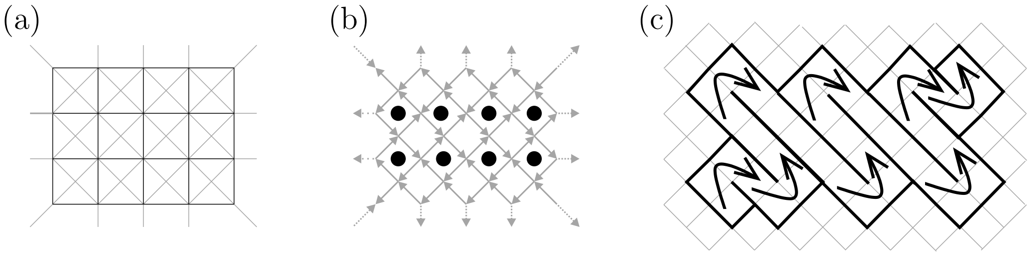

Figure 3.12. Part (a): The grid $G_{4,5}$ with its angle graph in gray. Part (b): A canonical 2-orientation on $\hat{G}_{4,5}$, the dotted edges all connect to an additional vertex v_∞ and the dots mark an edge disjoint set of directed cycles. Part (c): The central part of $\hat{G}_{6,9}$ and the traversal used in the proof of the upper bound.

The set I of edge disjoint directed cycles in the canonical orientation is marked by black dots in Figure 3.12 (b) and includes approximately half of all squares. The set I' consists of all squares that are not in I. Members of I' can be flipped if either the two cycles of I above it, or the two cycles of I below it are flipped, that is in 2 out of 16 cases (1 out of 4 for boundary squares). Roughly half of all squares are in I'. Thus, there are at least $2^{|I|+|I'|/8}$ bipolar orientations of $G_{k,\ell}$, which leads to an asymptotic lower bound of $2^{9k\ell/8} \approx 2.18^{k\ell}$.

face 3-coloring

inner 2-orientation

For the proof of the upper bound we use a bijection that Lieb describes in [65]. The bijection relates *face 3-colorings* where no two squares sharing an edge have the same color and *inner 2-orientations* of the square grid as shown in Figure 3.13 (a). Figure 3.13 (b) shows the face 3-coloring corresponding to the canonical 2-orientation, that we used for the proof of the lower bound.

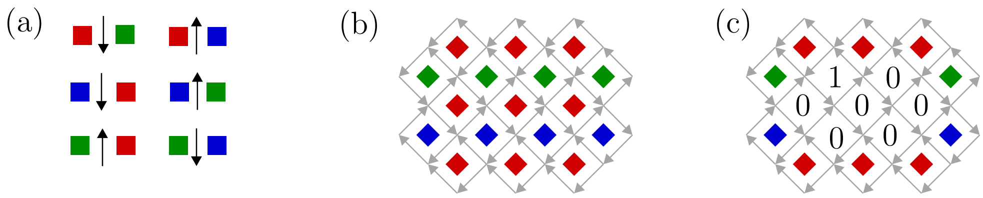

Figure 3.13. Part (a): Lieb's bijection between inner 2-orientations and face 3-colorings on the grid. The face 3-coloring for a particular orientation in Part (b) and its encoding in Part (c).

Here we use Lieb's bijection on $\hat{G}_{k,\ell}$ and prove an upper bound for the number of face 3-colorings of $\hat{G}_{k,\ell}$. Figure 3.12 (c) shows the *central part* of $\hat{G}_{k,\ell}$ bounded by a thick polygon. We will encode the 3-coloring on the faces of the central part of $\hat{G}_{k,\ell}$ as a sparse sequence a, where a_i represents the ith square on the path P indicated by the arrows in Figure 3.12 (c). The idea for this encoding is due to Graham Brightwell [18].

The set $\mathcal{D}$ of faces that are not in the central part has less than $3^{|\mathcal{D}|}$ 3-colorings. In the encoding described next, the code for the ith face of the path P depends only on faces in $\mathcal{D}$ and faces of P with index smaller than i. Figure 3.14 shows how the color of the highlighted face is encoded by a 0 or a 1 and Figure 3.13 (c) shows an example. The arrows indicate the direction in which we traverse the central part of the graph. There are three cases, one for a face where the path makes no turn and two for the two different types of turn faces. The variables X, Y, Z represent an arbitrary permutation of R, G, B.

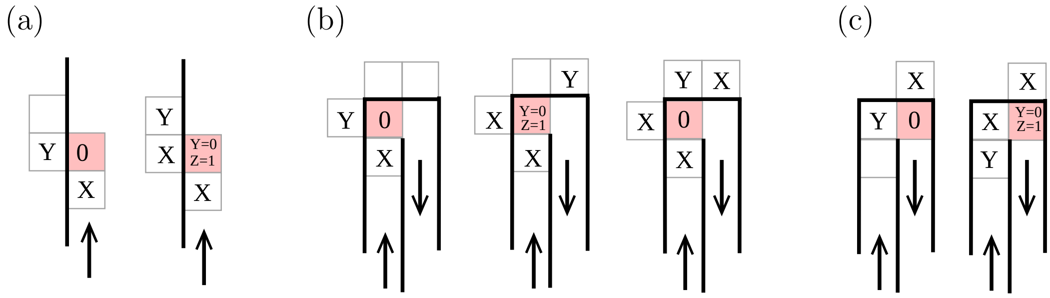

Figure 3.14. Encoding a 3-coloring by a sparse 0-1−sequence. In Part (a) the encoding for a square where the path makes no turn, in Parts (b) and (c) for the two different kinds of turn faces.

Concerning the decoding it is implied by Figure 3.14 that the faces marked with an X or Y plus the 0-1 encoding uniquely determine the color of the face in question. Thus, the encoding of a 3-face coloring using the colors of the squares from $\mathcal{D}$ and the 0-1 sequence for the central part of the grid is injective. It remains to show that there cannot be consecutive 1s in this sequence. This follows from the observation that writing a 1 means that the two faces that will be used for the encoding of the next face on the path have different colors. Thus, this face will be encoded by a 0.

We bound the number of such encodings from above. The set $\mathcal{D}$ can be covered by four horizontal plus four vertical rows of faces, thus $|\mathcal{D}| \leq 4(k+\ell)$. The length of the path is bounded by the number of bounded faces of $\hat{G}_{k,\ell}$ which is less than $2k\ell$. Therefore, there are at most $3^{4(k+\ell)} \cdot F_{2k\ell+2}$ encodings. Using the asymptotics for the Fibonacci numbers this implies that there are at most $2.619^{k\ell}$ encodings for k, ℓ big enough. $\square$

Lieb's analysis of the number of Eulerian orientations of $G^T_{k,\ell}$, see Proposition 3.15 and Remark 3.16, is of interest in this case as well. It allows to improve the upper bound for grids with side lengths ratio one to two.

Proposition 3.21. *For k big enough, the number of bipolar orientations of the grid $G_{k,2k}$ is bounded by*

$$2.18^{2k^2} \leq |\mathcal{B}(G_{k,2k})| \leq 2.38^{2k^2}.$$

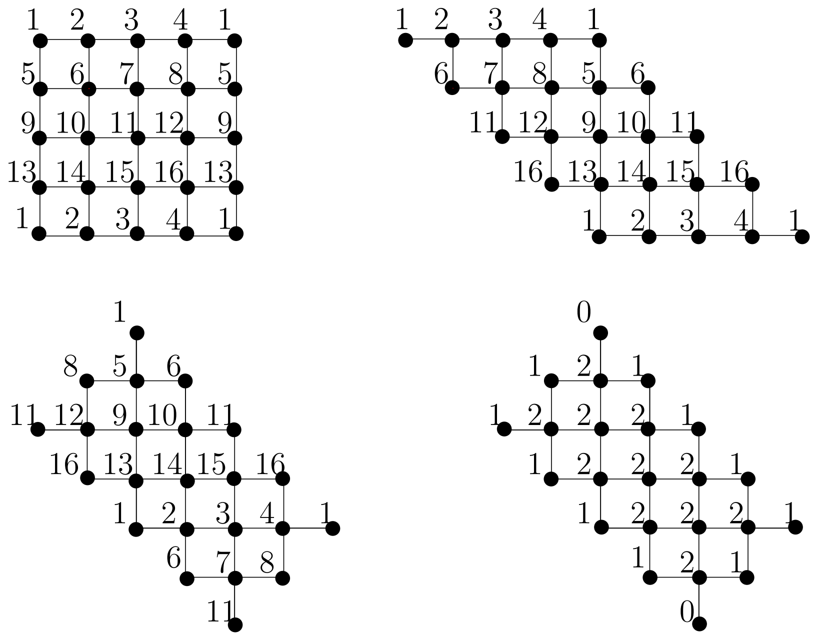

Figure 3.15. Obtaining the tilted grid $\hat{G}_{5,3}$ from $G^T_{4,4}$ with two cuts. The numbers in the first three drawings indicate vertex labels, in the last one they indicate $\hat{\alpha}$.

Proof. By $\hat{G}'_{k,\ell}$ we denote the graph obtained from $\hat{G}_{k,\ell}$ by deleting v_∞ and all incident edges. These edges are dotted in Figure 3.12. Figure 3.15 shows how to cut $G^T_{4,4}$ in two steps such that the grid looks like $\hat{G}'_{3,5}$ (if we do not identify vertices). The last drawing shows that every $\hat{\alpha}$-orientation of $\hat{G}'_{3,5}$ yields a Eulerian orientation of $G^T_{4,4}$ when we do the appropriate identifications. In general, this approach yields an injection from the bipolar orientations of $G_{k+1,2k+1}$ to the Eulerian orientations of $G^T_{2k,2k}$. As Lieb [65] has shown that $G^T_{2k,2k}$ has asymptotically $(8 \cdot \sqrt{3}/9)^{4k^2}$ Eulerian orientations, this yields an upper bound of $(64/27)^{2k^2}$ for the number of Eulerian orientations of $G_{k+1,2k+1}$. Every bipolar orientation of $G_{k,2k}$ can be complemented to a bipolar orientation of $G_{k+1,2k+1}$, thus $G_{k,2k}$ has at most as many bipolar orientation as $G_{k+1,2k+1}$. The lower bound follows from Theorem 3.20. □

Remark 3.22. The same problems as described in Remarks 3.10 and 3.16 arise here when trying to show that $G_{k,2k}$ actually has $(64/27)^{2k^2}$ bipolar orientations by using Lieb's result for the torus.

3.4.2 Bipolar Orientations of Planar Maps

Note that adding edges to the faces of size at least 4 of a planar map M can only increase the number of bipolar orientations by Proposition 3.18. Thus, we can restrict our considerations to plane inner triangulations in this section.

Theorem 3.23. *Let $\mathcal{M}_n$ denote the set of all planar maps with n vertices and $\mathcal{B}(M)$ the set of all bipolar orientations of $M \in \mathcal{M}_n$. Then for n big enough*

$$2.91^n \leq \max_{M \in \mathcal{M}_n} |\mathcal{B}(M)| \leq 3.97^n.$$

For the proof we need a couple of facts about *Fibonacci numbers* which are summarized in the following lemmas. The Fibonacci numbers are the integer series defined by the recursion *Fibonacci number*

$$F_0 = 0,\ F_1 = 1,\ F_n = F_{n-1} + F_{n-2} \text{ for } n \geq 3.$$

Let $\phi = \frac{1+\sqrt{5}}{2}$ be the Golden Ratio. The first two formulas in the next lemma are standard results from the vast theory of Fibonacci numbers, the last one is attributed to Shiwalkar and Deshpande in [84, A001629].

Lemma 3.24. *The Fibonacci numbers have the following properties.*

- $F_n = (\phi^n - (1-\phi)^n)/\sqrt{5}$
- $\lim_{n\to\infty} \left(F_n - \phi^n/\sqrt{5}\right) = 0$
- $\sum_{i=0}^{n} F_i F_{n-i} = \left(n(F_{n+1} + F_{n-1}) - F_n\right)/5$

The next lemma summarizes facts about *sparse sequences.* *sparse sequence*

Lemma 3.25. *The number of sparse sequences of length n is F_{n+2}. Let $r_n(i)$ be the number of sparse sequences of length n whose ith entry is 1. Then*

- $r_n(i) = F_i \cdot F_{n+1-i}$
- $\sum_{i=1}^{n} r_n(i) = \left(2(n+1)F_n + nF_{n+1}\right)/5$
- $\lim_{n\to\infty} \frac{\sum_{i=1}^{n} r_n(i)}{nF_{n+2}} = (\sqrt{5}\phi)^{-1} \approx 0.2764$

The first identity follows from a construction of sparse sequences of length n from sparse sequences of length $n-1$ plus the string "0" and sparse sequences of length $n-2$ plus "01". The second and third identity then follow using the facts from Lemma 3.24.

Before proving Theorem 3.23 we give two results for the number of bipolar orientations of special classes of planar maps.

stacked triangulation

Proposition 3.26. *A stacked triangulation with n vertices has 2^{n-3} bipolar orientations.*

Proof. The K_4 has two bipolar orientations for fixed source and sink. We proceed by induction and assume that a stacked triangulation with n vertices has 2^{n-3} bipolar orientations. Now let T be a stacked triangulation with $n+1$ vertices and v a vertex of degree 3 in T. Then, $T-v$ has 2^{n-3} bipolar orientations by induction. Now stacking v into T again, there are exactly two ways to complete a given bipolar orientation on $T-v$ without violating Properties (1) or (2) from Proposition 3.18. Thus, there are $2^{(n+1)-3}$ bipolar orientations of T. $\square$

Proposition 3.27. *Let $\mathcal{O}_n$ be the set of all outerplanar maps with n vertices. Then*

$$\max_{M\in\mathcal{O}_n} |\mathcal{B}(M)| = F_{n-1} \approx 1.618^{n-1}.$$

triangular grid

Proof. We show first that there are indeed outerplanar maps with F_{n-1} bipolar orientations. Let $T := T_{2,\ell}$ be the *triangular grid* with two rows. We consider bipolar orientations of T with source $(1,1)$ and sink $(2,\ell)$. In every such bipolar orientation the boundary edges form two directed paths from $(1,1)$ to $(2,\ell)$. We start by defining the standard bipolar orientation B_0 of T that is shown in Figure 3.16.

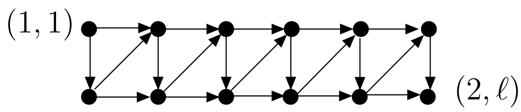

Figure 3.16. The standard bipolar orientation B_0 on $T_{2,\ell}$.

In B_0 the vertical inner edges are directed downwards and the diagonal ones upwards. Now we encode any other orientation of the inner edges by a sequence $(a_i)_{i=1\dots n'}$ of length $n'=n-3$, where $a_i=1$ if the corresponding edge has the opposite direction as in B_0 and $a_i=0$ otherwise. The entries come in the natural left to right order in $(a_i)_{i=1\dots n'}$. We show that all sparse sequences of length n' produce bipolar orientations. In a sparse sequence there are no consecutive 1s, thus out of the two inner edges incident to a vertex, at most one is reversed with respect to B_0. This guarantees that there is no directed facial 3-cycle. As all vertices except $(1,1)$ and $(2,\ell)$ have an incoming and an outgoing outer edge, the resulting orientation is bipolar, according to Proposition 3.18.

It remains to show that F_{n-1} is an upper bound for the number of bipolar orientations of any outerplanar map M with n vertices. We may assume that M is a plane inner triangulation. The proof uses induction on the number of vertices and the claim is trivial for $n = 3$. Now, let M have $n+1$ vertices and let s be the source vertex. If M has a vertex $x \neq s, t$ of degree 2 with neighbors v, w, then the direction of the edge $\{v, w\}$ determines the directions of the edges $\{x, v\}$ and $\{x, w\}$. Therefore, M has at most as many bipolar orientations as $M - x$, that is at most F_{n-1}. If all vertices but s and t have degree at least 3, then s and t have degree 2 and the vertices of every inner edge of M are separated by s and t on the outer cycle. This is because the interior of the boundary cycle on $n+1$ vertices is partitioned into $n-1$ triangles, and thus two of these triangles must share two edges with the boundary, which yields two degree 2 vertices.

So s is incident to only two vertices v and w, and we may assume that v has degree 3 in M, that is the inner edge $e = \{v, w\}$ is the only inner edge incident to v. Now, let X be some bipolar orientation of M in which e is directed from v to w. Then, the orientation of $M - s$ induced by X is a bipolar orientation with source v. For a bipolar orientation Y in which e is oriented from w to v, the orientation of $M - s$ induced by X is a bipolar orientation with source w and v is a vertex of degree 2 in $M - s$. This mapping is injective, and thus M has at most as many bipolar orientations as $M - s$ and $M - \{s, v\}$ together, that is $F_{n-1} + F_{n-2} = F_n$. □

Remark 3.28. From the above proof it also follows that $T_{2,\ell}$ is the only outerplanar map on 2ℓ vertices that has $F_{2\ell-1}$ bipolar orientations.

The example that gives the lower bound for the number of bipolar orientations of planar maps is the *triangular grid* $T_{k,k}$ with source $(1,1)$ and sink (k,k). *triangular grid*

Proposition 3.29. *Let $T_{k,k}$ be the triangular grid and k big enough. Then,*

$$|B(T_{k,k})| \geq 2.91^n.$$

Proof. We first show that $T_{k,k}$ has at least 2.618^{k^2} bipolar orientations. To see this we glue together $k-1$ copies of $T_{2,k}$ with bipolar orientations. Every bipolar orientation of $T_{k,k}$ obtained in this way corresponds to a concatenation of $k-1$ sparse sequences of length $2k-3$, as the proof of Proposition 3.27 shows. We call such a concatenation of sparse sequences an almost sparse sequence. We denote the set of all such sequences of length $2k^2 - 5k + 3$ by $S(k)$, the cardinality of $S(k)$ is F_{2k-1}^{k-1} which is bounded from below by $F_{2k^2-5k+3} \geq 2.618^{k^2}$ for k big enough. That each $s \in S(k)$ corresponds to a bipolar orientation of $T_{k,k}$ can be checked using Proposition 3.18.

Now we improve this to the claimed bound of 2.91^n. The horizontal edge $e_{i,j} := (i,j) \to (i,j+1)$ lies on the boundary of two triangles for $2 \le i \le k-1$. As Figure 3.17 (a) shows, the other four edges of these triangles are

$$\{(i,j),(i-1,j+1)\},\ \{(i,j+1),(i-1,j+1)\},$$
$$\{(i,j),(i+1,j)\},\ \{(i,j+1),(i+1,j)\}.$$

The crucial observation for improving the above bound is that we can reorient $e_{i,j}$ if and only if the entries belonging to these four edges show one of the two patterns $10\ldots01$ or $01\ldots10$, see Figures 3.17 (b) and (c).

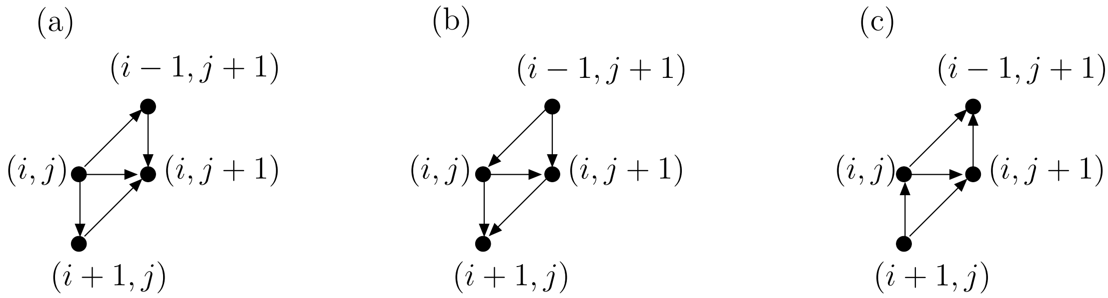

Figure 3.17. The standard bipolar orientation on the triangles incident to $e_{i,j}$ and the two orientations that allow to reorient $e_{i,j}$.

We now choose $k-1$ sparse sequences of length $2k-3$ independently uniformly at random. We then concatenate them to obtain a random almost sparse sequence $s \in S(k)$. It follows from the first identity from Lemma 3.25 with $n = 2k-3$ and $i = 2j-1$ that for $\{(i,j),(i,j+1)\}$ there are $F_{2j-1}F_{2k-2j-1}F_{2k-1}^{k-2}$ sequences that have $\{(i,j),(i-1,j+1)\}$ marked 1 out of the total F_{2k-1}^{k-1} sequences. This and the second identity from Lemma 3.24 are used to calculate the probability that the entry for $\{(i,j),(i-1,j+1)\}$ is 1 as

$$\begin{aligned}\lim_{k\to\infty} \frac{F_{2j-1}F_{2k-2j-1}F_{2k-1}^{k-2}}{F_{2k-1}^{k-1}} &= \lim_{k\to\infty} \frac{1}{\sqrt{5}} \cdot \phi^{2j-1} \cdot \phi^{2k-2j-1} \cdot \phi^{-2k+1} \\ &= \frac{1}{\sqrt{5}\phi}. \end{aligned} \tag{3.5}$$

Taking the limit is only justifiable if $(2j-1) \to \infty$ and $(2k-2j-1) \to \infty$ for $k \to \infty$. Therefore we introduce $\delta > 0$ and denote the set of horizontal edges with $\delta(k-1) \le j \le (1-\delta)(k-1)$ and $2 \le i \le k-1$ by E_δ. Taking the limit in Equation (3.5) is justified for all $e_{i,j} \in E_\delta$. The size of this set is $|E_\delta| = (1-2\delta)(k-1)^2$.

Analogously we can calculate the probability that the edge $\{(i,j+1),(i+1,j)\}$ is flipped and these events are independent. Thus, the probability of the pattern

$10\ldots01$ is $(5\phi^2)^{-1}$ in the limit. The pattern $01\ldots10$ has the same probability and the patterns mutually exclude each other. Thus, for every $\epsilon > 0$ the probability that the edge $\{(i,j),(i,j+1)\} \in E_\delta$ can be flipped is

$$\mathbb{P}[\mathbb{1}_{i,j}(s) = 1] \geq \frac{2}{5\phi^2} - \epsilon$$

for k big enough.

We now analyze how many of the flip-patterns we expect for a sparse sequence s. Let $Q(s) = \sum_{i,j} \mathbb{1}_{i,j}(s)$ be a random variable counting the number of flippable edges in s. We use *Jensen's inequality* to estimate the number of orientations which is *Jensen's inequality*

$$|S(k)| \cdot \mathbb{E}_{s\in\mathcal{S}}[2^{Q(s)}] \geq F_{2k^2-5k+3} \cdot 2^{\mathbb{E}_{s\in\mathcal{S}}[Q(s)]} \geq 2.618^{k^2} \cdot 2^{(1-2\delta)(k-1)^2\left(\frac{2}{5\phi^2}-\epsilon\right)} \geq 2.91^{k^2}$$

for k big enough. □

Remark 3.30. We have included the third identity of Lemma 3.25 because it emphasizes in conjunction with Equation (3.5) that the expected number of 1s at a fixed entry of a sparse sequence does not depend strongly on the choice of the entry.

The following relation is useful to prove an upper bound for the number of bipolar orientations of plane inner triangulations. It has been presented with a different proof in [70]. Let $\mathcal{F}_b$ be the set of bounded faces of M and $\mathcal{B}(M)$ the set of bipolar orientations of M. Fix a bipolar orientation B. The boundary of every triangle $\Delta \in \mathcal{F}_b$ consists of a path of length 2 and a direct edge from the source to the sink of Δ. We say that Δ is a $+$ *triangle* of B if looking along the direct source-sink edge the triangle is on the left. Otherwise, if the triangle is on the right of the edge we speak of a $-$ *triangle*, see Figure 3.18. We use this notation to define a function $G_B : \mathcal{F}_b \to \{-,+\}$. *+ triangle* *− triangle*

Theorem 3.31. *Let M be a plane inner triangulation and B a bipolar orientation of M. Given s, t, and G_B, i.e. the signs of bounded faces, it is possible to recover B. In other words the function $B \to G_B$ is injective on $\mathcal{B}(M)$.*

Proof. Given G_B we construct B. We start by orienting all edges on the boundary of the outer face such that s and t are the unique source and sink of this face. We extend this partial orientation Y with two rules. The *vertex rule* is applied to a vertex v that already has incoming and outgoing edges. It takes a bundle of consecutive edges of v that is bounded by two outgoing edges. It orients all the edges of the bundle such that they are outgoing at v. Note that these edge orientations are forced by Property (1′) of bipolar orientations in Proposition 3.18.

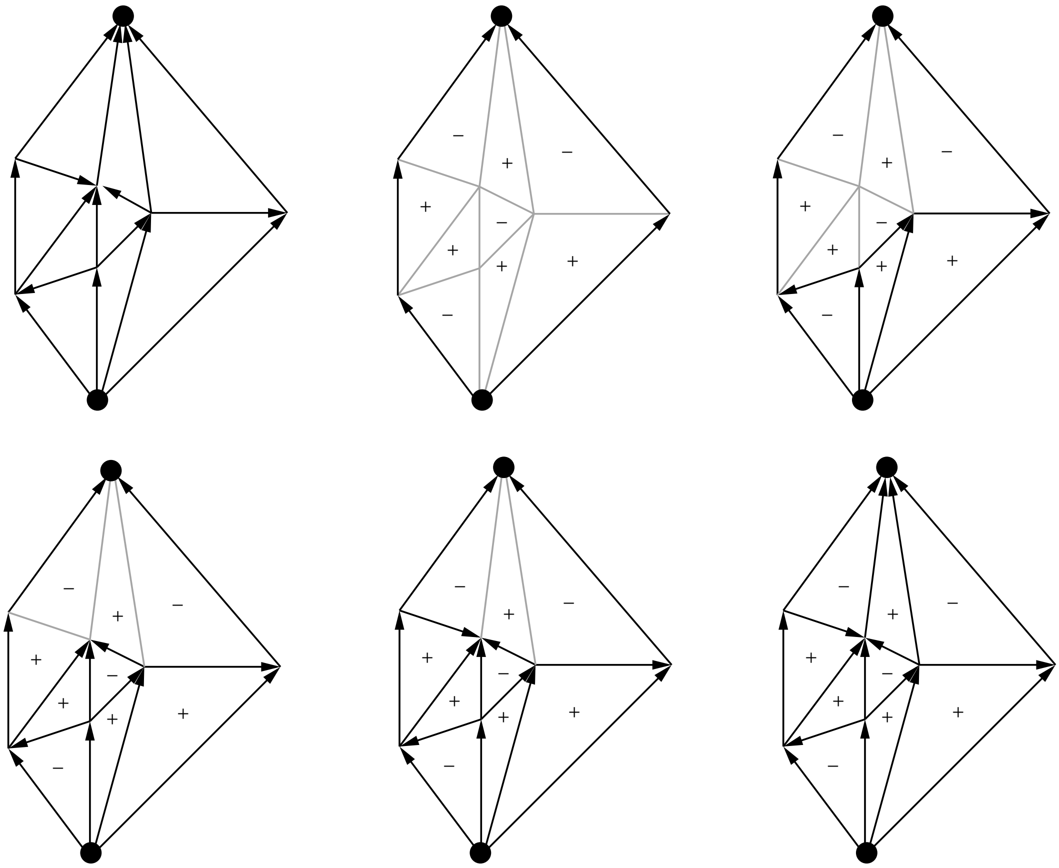

Figure 3.18. A bipolar orientation, the corresponding $+/-$ encoding and an illustration of the decoding algorithm.

The *face rule* is applied to a facial triangle Δ that has two oriented edges. The sign $G_B(\Delta)$ is used to deduce the orientation of the third edge.

Note that these two rules preserve the property that every vertex v that is incident to an oriented edge in Y can be reached from s along an oriented path. In particular, v has an incoming edge.

Let A_Y be the union of all faces that have all boundary edges oriented. Initially, A_Y consists of the outer face. Since B is acyclic, the boundary of A_Y is acyclic as well. Consequently, as long as there are faces that do not belong to A_Y, there is a vertex v on the boundary of A_Y that has two outgoing edges that belong to the boundary of A_Y. Either v is a candidate to extend the orientation using the vertex rule or there is a face incident to v that becomes an element of A_Y by applying the face rule to it.

We have thus shown that the rules can be applied until A_Y is the whole plane, i.e., all edges are oriented. They have to be oriented as in B, by construction. $\square$

The next theorem gives a necessary and sufficient condition for a vector in $\{-,+\}^{|\mathcal{F}_b|}$ to induce a bipolar orientation. For the sake of simplicity we state it only for triangulations, but the generalization to inner triangulations is straight forward. In order to obtain a more elegant formulation, we adopt the convention that the unbounded face is signed $+$ if the bounded face adjacent to sink and source is signed $-$. Otherwise the unbounded face is signed $-$. Thus, we work now with signings of the set $\mathcal{F}$ of all faces. We say that a $+$ triangle is the *right knee* of the vertex at which it has an incoming and an outgoing edge. Similarly a $-$ triangle is the *left knee* of exactly one of its vertices. For a vertex v of an inner triangulation T we denote by $\Delta^+(v)$ and $\Delta^-(v)$ the triangles that are the right respectively left knee of v.

right knee

left knee

$\Delta^+(v)$, $\Delta^-(v)$

Theorem 3.32. *Let T be a triangulation, $x \in \{-,+\}^{|\mathcal{F}|}$, and $\mathcal{F}^-$ and $\mathcal{F}^+$ the sets of faces that have negative respectively positive sign in x. Let $\widehat{M}(T)^+$ and $\widehat{M}(T)^-$ denote the subgraphs of the reduced angle graph $\widehat{M}(T) - \{s,t\}$ induced by $V \cup \mathcal{F}^-$ respectively $V \cup \mathcal{F}^+$.*

Then x induces a bipolar orientation on T if and only if both $\widehat{M}(T)^+$ and $\widehat{M}(T)^-$ have a unique perfect matching.

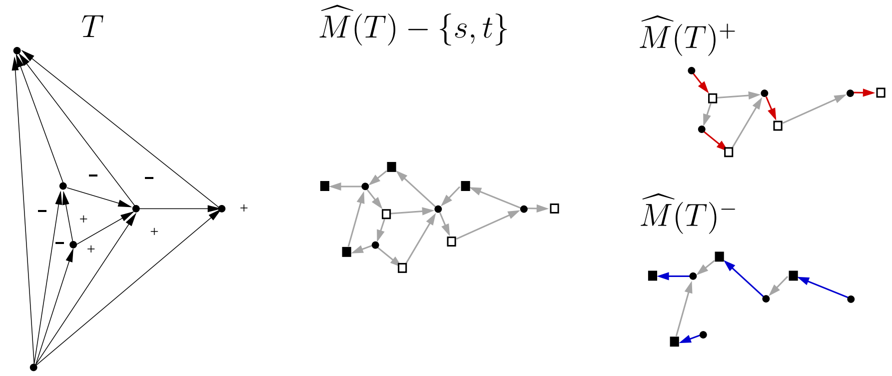

Figure 3.19. The two perfect matchings induced by a bipolar orientation.

Theorem 3.32 is illustrated in Figure 3.19. It implies that every vertex of T other than s, t must be adjacent to at least one $+$ triangle and one $-$ triangle. This fact can be proved with less effort than Theorem 3.32. Therefore we delay the rather long proof of Theorem 3.32 to Section 3.4.3.

Proposition 3.33. *Let T be a plane inner triangulation. Then, T has at most 3.97^n bipolar orientations.*

Proof. By Euler's formula, there are $2^{f-1} \leq 2^{2n-2-f_\infty} = 4^{n-1} \cdot 2^{-f_\infty}$ many binary vectors of length $f-1$. By the bijection from Theorem 3.31, $4^{n-1} \cdot 2^{-f_\infty}$ is also an upper bound for the number of bipolar orientations of T.

To squeeze the bound below 4^n we use the above observation that every vertex of T must be adjacent to at least one $+$ triangle and at least one $-$ triangle. Thus, out of the $2^{d(v)}$ possible $+/-$ vectors at a vertex v at least two are not feasible. Similarly, at an outer vertex $v \neq s, t$, there is exactly one angle forming a knee at v. The sign of this angle depends on which of the two oriented paths of the outer boundary v lies on, but it is fixed either way. Thus, out of the $2^{d(v)-1}$ possible sign patterns at v at least one is not feasible. We summarize, that at most a fraction of $(1-2^{1-d(v)})$ of all sign vectors is potentially feasible at every vertex but s and t. We denote the set $V \setminus \{s,t\}$ by V' and its cardinality by n', i.e. $n' = n-2$.

Jensen's inequality

We apply *Jensen's inequality* which says that for a concave function φ the inequality $\varphi\left(\sum x_i/n\right) \geq \left(\sum \varphi(x_i)\right)/n$ holds. As $\log x$ is concave we obtain

$$\log\left(\left(\prod_{v\in V'}(1-2^{1-d(v)})\right)^{1/n'}\right) = \frac{1}{n'}\sum_{v\in V'}\log\left(1-2^{1-d(v)}\right) \leq \log\left(\frac{1}{n'}\sum_{v\in V'}\left(1-2^{1-d(v)}\right)\right).$$

By the monotonicity of the logarithm this implies

$$\prod_{v\in V'}(1-2^{1-d(v)}) \leq \left(\frac{1}{n'}\sum_{v\in V'}\left(1-2^{1-d(v)}\right)\right)^{n'}.$$

The function $1-2^{1-x}$ is concave and we apply Jensen's inequality again which yields

$$\frac{1}{n'}\sum_{v\in V'}\left(1-2^{1-d(v)}\right) \leq 1-2^{1-\sum_{v\in V'} d(v)/n'}.$$

Since we deal with simple planar maps, $\sum_{v\in V'} d(v) \leq 2(3n-6) = 6n'$ and we conclude

$$\prod_{v\in V'}(1-2^{1-d(v)}) \leq \left(1-2^{1-6}\right)^{n'} = \left(\frac{31}{32}\right)^{n'}.$$

By the Four Color Theorem T can be partitioned into at most four independent sets I_k, $k=1,\ldots,4$. Thus,

$$\left(\frac{31}{32}\right)^{n'} \geq \prod_{v\in V'}(1-2^{1-d(v)}) = \prod_{k=1}^{4}\prod_{v\in I_k}(1-2^{1-d(v)})$$

and for at least one of the independent sets it must hold that

$$\left(\frac{31}{32}\right)^{n'/4} \geq \prod_{v\in I_k}(1-2^{1-d(v)}).$$

We are ready to conclude that there are at most

$$4^{n-1} \cdot 2^{-f_\infty} \cdot \left(\frac{31}{32}\right)^{(n-2)/4} < 3.97^n \cdot 2^{-f_\infty} \cdot \left(\frac{32}{31}\right)^{1/2} < 3.97^n$$

bipolar orientations of T. □

3.4.3 Bipolar Orientations and Face Signings

This section is devoted to the proof of Theorem 3.32 which we repeat here for convenience.

Theorem 3.32. *Let T be a triangulation, $x \in \{-,+\}^{|\mathcal{F}|}$, and $\mathcal{F}^-$ and $\mathcal{F}^+$ the sets of faces that have negative respectively positive sign in x. Let $\widehat{M}(T)^+$ and $\widehat{M}(T)^-$ denote the subgraphs of the reduced angle graph $\widehat{M}(T) - \{s,t\}$ induced by $V \cup \mathcal{F}^-$ respectively $V \cup \mathcal{F}^+$.*

Then x induces a bipolar orientation on T if and only if both $\widehat{M}(T)^+$ and $\widehat{M}(T)^-$ have a unique perfect matching.

Proof. We may assume that the outer face F_∞ is signed $+$ and that the bounded face F_0 that is incident to both s and t is signed $-$. It is easy to see that F_∞ is the only degree 1 vertex of $\widehat{M}^+(T)$ and F_0 is the only degree 1 vertex of $\widehat{M}^-(T)$. The degree 2 vertices of both graphs correspond to triangles that are incident to either s or t.

We first show that if x indeed induces a bipolar orientation B, then the graphs $\widehat{M}^+(T)$ and $\widehat{M}^-(T)$ both have a unique perfect matching. By symmetry it suffices to show this for $\widehat{M}^+(T)$. Clearly, $\widehat{M}^+(T)$ must have a perfect matching $\Sigma^+(T)$, since every vertex $v \neq s,t$ has a right knee in B and every $+$ triangle is a right knee for exactly one vertex. We now show that $\Sigma^+(T)$ is the unique perfect matching of $\widehat{M}^+(T)$. We use the well known fact, that a perfect matching is unique if and only if it has no alternating cycle.

Since B is a bipolar orientation, v lies on a directed s-t-path P_v in B that uses the two edges incident to v which lie on $\Delta^+(v)$. This path forms two bounded regions R_1 and R_2 with the boundary of the outer face F_∞, and we may assume that R_1 is the region that contains $\Delta^+(v)$ and no other triangle incident to v. We observe that all inner vertices of P_v also have their right knee in R_1. An alternating cycle of $\Sigma^+(T)$ in $\widehat{M}^+(T)$ through v corresponds to a sequence of vertices and triangles $v = v_1, T_1, \ldots, v_k, T_k, v_{k+1} = v$, such that T_i is the right knee of v_i and T_{i-1} is not a knee of v_i. By the above observation all T_i lie in R_1. The only face from R_1 incident to v is T_1. This is a contradiction since v must be reentered from $T_k \notin R_1$.

We have thus shown that if x induces a bipolar orientation, then $\widehat{M}^+(T)$ and $\widehat{M}^-(T)$ have unique perfect matchings, since no vertex lies on an alternating cycle.

We now prove that if $\widehat{M}^+(T)$ and $\widehat{M}^-(T)$ each have a unique perfect matching $\Sigma^+(T)$ respectively $\Sigma^-(T)$, then x induces a bipolar orientation. The proof uses induction and the two reduction rules shown in Figure 3.20.

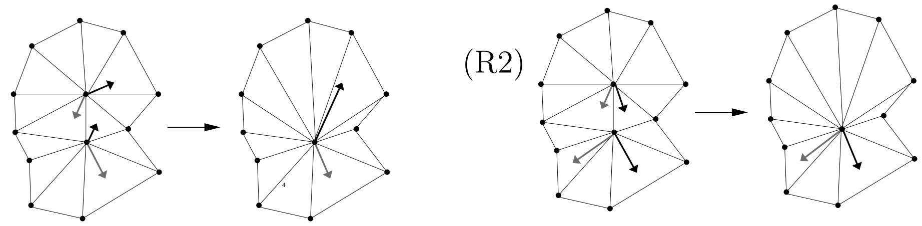

Figure 3.20. The reductions used for the induction step. A black arrow indicates a $+$ triangle and a gray arrow indicates a $-$ triangle.

Reduction (R1) is *applicable* if there is an edge such that the corresponding 4-cycle in $\widehat{M}(T) - \{s, t\}$ contains one edge from every matching and these two edges do not share a vertex. Reduction (R2) is *applicable* if there is an edge such that the corresponding 4-cycle in $\widehat{M} - \{s, t\}$ contains one edge from every matching and these two edge do share a vertex.

We first describe how one would intuitively build an orientation of T from $\Sigma^+(T)$ and $\Sigma^-(T)$. Since T is embedded, we have a clockwise adjacency list $L(v)$ of incident edges at every vertex v. In a bipolar orientation, the edges that lie between the angle of the edge $e^-(v) \in \Sigma^-(T)$ and the angle of the edge $e^+(v) \in \Sigma^+(T)$ in $L(v)$ are directed away from v and those which lie between $e^+(v)$ and $e^-(v)$ are directed towards v. We have to show that this is well defined, that is if an edge $e = \{v, w\}$ lies between $e^-(v)$ and $e^+(v)$ then it lies between $e^+(w)$ and $e^-(w)$. An edge is called a *good edge* if it satisfies this condition and *bad* otherwise. In the above situation e is called an *out-edge* of v and an *in-edge* of w. Thus, an edge is bad if it is an out-edge of both its end vertices or an in-edge of both its end vertices. A reduction is *admissible* if it is applicable and the resulting triangulation T' satisfies the following conditions.

(i) The graph $\widehat{M}^+(T')$ has a unique perfect matching if and only if $\widehat{M}^+(T)$ does.

(ii) The graph $\widehat{M}^-(T')$ has a unique perfect matching if and only if $\widehat{M}^-(T)$ does.

(iii) The triangulation T' has bad edges if and only if T has bad edges.

We will show that if T has at least four vertices and unique perfect matchings $\Sigma^+(T)$ and $\Sigma^-(T)$, then one of the reductions (R1), (R2) is admissible. We then apply admissible reductions until T is reduced to just one triangle. A triangle with given source s and sink t has a unique bipolar orientation and no inner edge, in particular no bad edge. This implies that T cannot have bad edges, since the presence of bad edges is preserved by the admissible reductions. We then explain how to construct a bipolar orientation of T by reversing the reductions, where we use that all edges of T are good.

We now show that if T has at least four vertices and $\Sigma^+(T)$ and $\Sigma^-(T)$ are unique, then there is quadrangle in $\widehat{M}(T) - \{s, t\}$ such that two of its four edges are matching edges. Since two edges from the same matching on a 4-cycle would form an alternating cycle, this implies that (R1) or (R2) is applicable. We will then show that if (R1) or (R2) is applicable, then there is also a reduction which is admissible. To show that some reduction is applicable we use a counting argument. Let n denote the number of vertices of T. Then, $\widehat{M}(T) - \{s, t\}$ has $3n - 6 - (d(s) + d(t) - 1)$ quadrangular faces, one for every edge of T that is neither incident to s nor to t. There are $2n-4$ matching edges. Two of these matching edges lie on no quadrangle, namely those incident to F_0 and F_∞. There are $d(s) - 2 + d(t) - 2$ matching edges that lie on one quadrangle, namely those at faces that are incident to either s or t. All $2n - 4 - 2 - (d(s) + d(t) - 4)$ other matching edges lie on two quadrangles. Thus, there are

$$2(2n - d(s) - d(t) - 2) + d(s) + d(t) - 4 = 4n - d(s) - d(t) - 8$$

incidences between matching edges and quadrangles. Since

$$4n - d(s) - d(t) - 8 > 3n - d(s) - d(t) - 5$$

for $n \geq 4$, the pigeon hole principle implies that at least one quadrangle must have two matching edges, and thus some reduction is applicable.

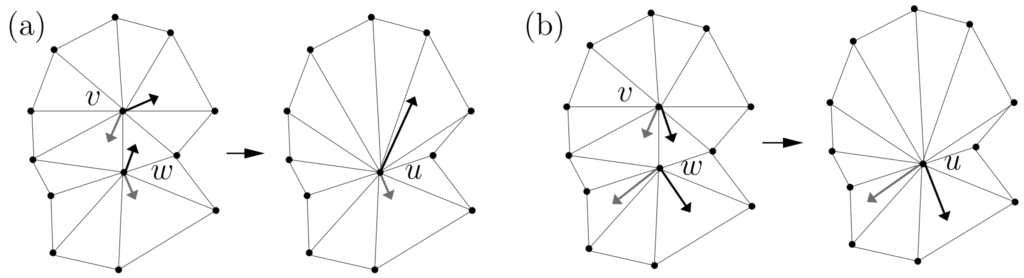

Figure 3.21. Reductions of good edges.

We now show that if there is an applicable reduction, then there also is an admissible one. Suppose (R1) is applicable to contract the edge $e = \{v, w\}$ and denote the new vertex of T' by u, see Figure 3.21 (a). It is easy to see that if $\Sigma^+(T)$

has an alternating cycle, then so does $\Sigma^+(T')$. The converse is also true, since v is reachable from w in $\widehat{M}^+(T)$ on an alternating path of $\Sigma^+(T)$. Since w is reachable from v in $\widehat{M}^-(T)$ on an alternating path of $\Sigma^-(T)$ it follows that $\widehat{M}^-(T')$ has an alternating cycle if and only if $\widehat{M}^-(T')$ has an alternating cycle. Furthermore every out-edge of v or w is an out-edge of u and the same is true for in-edges. Since e is a good edge, and no other edges change their in-edge respectively out-edge status at any vertex, we conclude that T' has bad edges if and only T does.

Next, suppose that (R2) is applicable to an edge $e = \{v, w\}$ and let u be defined as above. If e is a good edge, then proof is analogous to that of the case when (R1) is applicable, see Figure 3.21 (b).

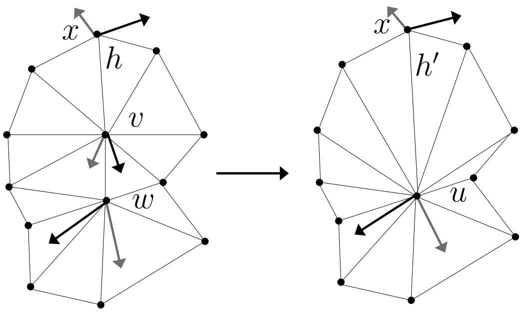

Figure 3.22. Reduction of a bad edge, when v is incident to a good edge h.

Now suppose that e is a bad edge. Out-edges of v become in-edges of u and in-edges of v become out-edges of u. Thus, if v is incident to a good edge $h = \{v, x\}$, then the edge $h' = \{u, x\}$ of T' will be a bad edge, see Figure 3.22. It is again easy to see that $\Sigma^+(T')$ respectively $\Sigma^-(T')$ has an alternating cycle if and only if $\Sigma^+(T)$ respectively $\Sigma^-(T)$ does.

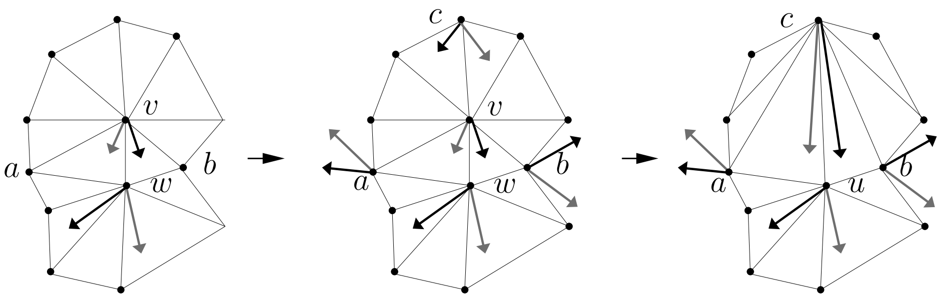

Figure 3.23. Reduction of a bad edge, when all edges incident to v are bad.

It remains to consider the case that all edges incident to v are bad, see Figure 3.23. Let a and b be the vertices that form the triangles incident to e together with v and w. Since $\{v, a\}$ and $\{v, b\}$ are bad, and $\Sigma^+(T)$ and $\Sigma^-(T)$ do not have alternating cycles, none of a, b can have a knee at a face incident to v. There are $d(v) - 2$ triangles incident to v that are not a knee of v and only $d(v) - 3$ neighbors of v for which these triangles can be a knee. Hence, there must be one neighbor c of v that has both its knees at triangles incident to v, see Figure 3.23. We contract the edge $\{v, c\}$ with an (R2) reduction. Then, the edge e remains bad and thus T' satisfies the third condition for admissibility. It is again easy to see that $\Sigma^+(T')$ respectively $\Sigma^-(T')$ has an alternating cycle if and only if $\Sigma^+(T)$ respectively $\Sigma^-(T)$ does.

Thus, we have shown that (R1) or (R2) is always admissible. Since the reduction process yields a triangulation that has no bad edges, T does not have bad edges either, i.e. we have only been performing reductions as shown in Figure 3.21.

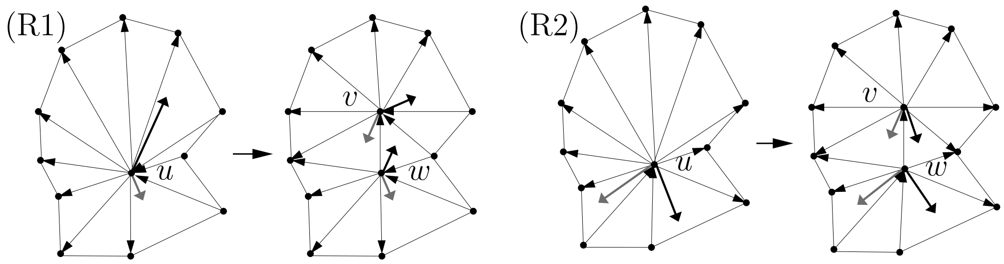

Figure 3.24. Constructing a bipolar orientation by reversing the reductions.

Figure 3.24 shows how to construct an orientation of T from a given orientation of T' by reversing an (R1) respectively (R2) reduction of a good edge. Using Proposition 3.18 it is easy to see that if the original orientation of T' is bipolar, then so is the resulting one of T. It is also immediate that the face signing G_B of the constructed bipolar orientation B of T is indeed x. This concludes the proof of the theorem. □

Remark 3.34. Let $\widehat{M}^+(T)$ and $\widehat{M}^-(T)$ be given along with perfect matchings $\Sigma^+(T)$ and $\Sigma^-(T)$ and the information that all edges of T are good. It is then easy to show that both matchings are the unique perfect matchings of the respective graph. That is, the conditions that the matchings are unique and that all edges are good are equivalent.

3.5 The Complexity of Counting α-Orientations

Given a planar map M and a function $\alpha : V \to \mathbb{N}$, what is the complexity of computing the number of α-orientations of M?

The class $\#P$ plays a prominent role in the complexity theory of counting problems. For the sake of completeness we briefly introduce the class $\#P$. Loosely speaking, $\#P$ is the class of counting problems naturally associated with the decision problems from $\mathcal{NP}$. We now give a formal definition. Let Σ be a finite alphabet and $R \subseteq \Sigma^* \times \Sigma^*$ a relation. Furthermore, we assume that there exists a polynomial p with $\langle x, y\rangle \in R \Rightarrow |y| \leq p(|x|)$ and that the language $L = \{\langle x, y\rangle \in R\}$ is decidable in polynomial time. Let $R(x) = \{y \mid \langle x, y\rangle \in L\}$. Then, problem of deciding for $x \in X$ whether $R(x)$ is nonempty is in $\mathcal{NP}$ and the counting problem to determine $|R(x)|$ is in $\#P$. A counting problem is *$\#P$-complete* if it is in $\#P$ and every problem in $\#P$ can be polynomially reduced to it.

$\#P$-complete

In some instances the number of α-orientations can be computed efficiently, e.g. for perfect matchings and spanning trees of planar maps. In Section 3.5.1 we show that counting is $\#P$-complete for other α-orientations, and in Section 3.5.2 we discuss how to take advantage of an existing fully polynomial randomized approximation scheme by applying it to α-orientations.

3.5.1 $\#P$-Completeness

Recently, Creed [29] has proved Theorem 3.35. As already mentioned, *Eulerian orientations* are α-orientations, and hence Theorem 3.35 says that counting α-orientations is $\#P$-complete. Since we use Creed's proof technique in the sequel, we sketch the proof here. It uses a reduction from counting Eulerian orientations which has been proven to be $\#P$-complete by Mihail and Winkler in [71].

Eulerian orientation

Theorem 3.35. *It is $\#P$-complete to count Eulerian orientations of planar graphs.*

Proof. We aim to show that the number of Eulerian orientations of a graph G can be computed in polynomial time with the aid of polynomially many calls to an oracle for the number of Eulerian orientations of a planar graph.

In order to count the Eulerian orientations of a graph G with n vertices, a drawing of this graph in the plane with ℓ crossings is produced. We may assume that no three edges cross in the same point and that ℓ is of order $O(n^4)$.

From this drawing a family of graphs G_k for $k = 0, \ldots, \ell$ is produced. In G_k every crossing of two edges $\{u, v\}$ and $\{x, y\}$ is replaced by the *crossover box* H_k on $4k + 1$ vertices, see Figure 3.25. In G_0 for example every crossing is simply replaced by a vertex. We call the edges $\{u, w_k^u\}$, $\{v, w_k^v\}$, $\{x, w_k^x\}$, $\{y, w_k^y\}$ the *connection edges of H_k*.

Every Eulerian orientation of G induces a Eulerian orientation of G_0, but there are Eulerian orientations of G_0 that do not come from a Eulerian orientation of G. Given a Eulerian orientation of G_0, we call the orientation of the edges incident to a vertex w_0 that replaces a crossing *valid* if exactly one of the edges $\{v, w_0\}$, $\{u, w_0\}$ is directed away from w_0, and invalid otherwise.

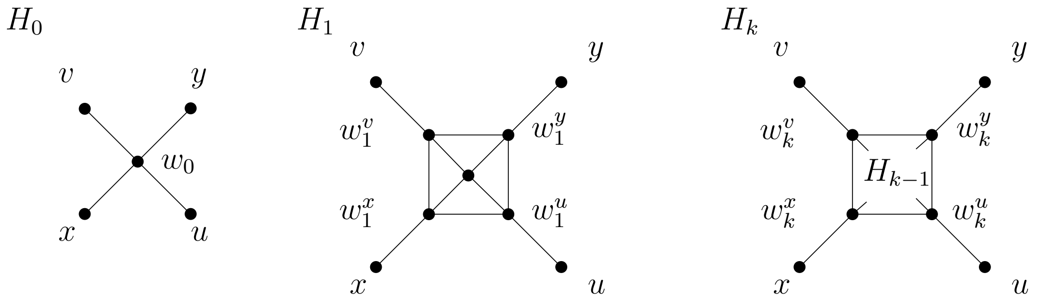

Figure 3.25. The crossover boxes defined in [29].

We call a configuration of the connection edges of H_k valid if exactly one of the edges $\{u, w_k^u\}$, $\{v, w_k^v\}$ is directed towards $\{u, v\}$ and exactly one of the edges $\{x, w_k^x\}$, $\{y, w_k^y\}$ is directed towards $\{x, y\}$. We call a configuration of the connection edges of H_k invalid if exactly two of the connection edges are directed towards $\{u, v, x, y\}$, but it is not a valid configuration. By x_k respectively y_k we denote the number of ways that a valid respectively invalid configuration of the connection edges of H_k can be extended to an orientation of the edges of H_k, such that all vertices of H_k have out-degree 2.

It is clear that $x_0 = 1 = y_0$ and Creed observes the following recursion

$$\begin{aligned} x_{k+1} &= 4x_k + 2y_k \\ y_{k+1} &= 4x_k + 3y_k. \end{aligned}$$

This recursion formula can be verified by a simple enumeration. In [29] a lemma from [96] is used to argue that the sequence x_k/y_k is non-repeating. In Lemma 3.36 we provide an easy argument to show from first principles that x_k/y_k is strictly monotonically decreasing.

Let N_i denote the number of Eulerian orientations of G_0 which have i valid cross over boxes, that is N_ℓ is the number of Eulerian orientations of G. The number of Eulerian orientations of G_k is

$$EO(G_k) = \sum_{i=0}^{\ell} N_i x_k^i y_k^{\ell-i}.$$

Hence the number $EO(G_k)/y_k^\ell$ is the value of the polynomial $p(z) = \sum_{i=0}^{\ell} N_i z^i$ at the point $z = x_k/y_k$. Since computing y_k^ℓ is easy, the polynomial p of degree ℓ can be evaluated at $\ell+1$ different points with $\ell+1$ calls to an oracle for counting Eulerian orientations of planar graphs. Hence the coefficients of p can be determined using polynomial interpolation. In particular this yields a way to compute N_ℓ. $\square$

Lemma 3.36. *Let $x_0 = 1 = y_0$ and $x_{k+1} = 4x_k + 2y_k$ while $y_{k+1} = 4x_k + 3y_k$. Then the sequence x_k/y_k is strictly monotonically decreasing.*

Proof. It follows directly from the recursion that $x_{k+1} = y_{k+1} - y_k$ and $y_{k+1} = 7y_k - 4y_{k-1}$. In order to show that $x_{k+1}/y_{k+1} < x_k/y_k$ we use the following equivalence.

$$\frac{x_{k+1}}{y_{k+1}} = \frac{4x_k + 2y_k}{4x_k + 3y_k} < \frac{x_k}{y_k} \iff 4\left(\frac{x_k}{y_k}\right)^2 - \frac{x_k}{y_k} - 2 > 0$$

Since $x_k/y_k > 0$ this inequality is satisfied if and only if $x_k/y_k > (1+\sqrt{33})/8 =: c$. Note that $1/(1-c) = 3 + 4c$ can be easily derived since c solves $4t^2 - t - 2 = 0$.

It remains to show that $x_k/y_k = 1 - (y_{k-1}/y_k) > c$ which we do by induction. Since $x_0/y_0 = 1$ and $x_1/y_1 = 6/7 > c$ the induction base holds and the induction hypothesis guarantees that $1 - c > y_{k-2}/y_{k-1}$ for $k \geq 2$. We have

$$\frac{x_k}{y_k} = 1 - \frac{y_{k-1}}{y_k} > c \iff \frac{y_k}{y_{k-1}} > \frac{1}{1-c}.$$

Using the induction hypothesis we obtain

$$\frac{y_k}{y_{k-1}} = \frac{7y_{k-1} - 4y_{k-2}}{y_{k-1}} = 7 - 4\frac{y_{k-2}}{y_{k-1}} > 3 + 4c = \frac{1}{1-c}$$

□

In [71] counting perfect matchings in bipartite graphs is reduced to counting Eulerian orientations. This reduction creates vertex degrees that grow linearly with the number of vertices of the reduced graph. It remains open whether counting Eulerian orientations of graphs with bounded maximum degree is #P-complete. We could not settle this question, but the next theorem shows that counting α-orientations is #P-complete even when the vertex degrees are restricted.

Theorem 3.37. *For the following graph classes and out-degree functions α the counting of α-orientations is #P-complete.*

1. *Planar maps with $d(v) = 4$ and $\alpha(v) \in \{1, 2, 3\}$ for all $v \in V$.*

2. *Planar maps with $d(v) \in \{3, 4, 5\}$ and $\alpha(v) = 2$ for all $v \in V$.*

The proof uses the planarization method from the proof of Theorem 3.35 in conjunction with the following theorem from [30].

Theorem 3.38. *For $k \geq 3$ counting perfect matchings of k-regular bipartite graphs is #P-complete.*

Proof of Theorem 3.37. Perfect matchings of a bipartite graph G with vertex set $V = A \cup B$ are in bijection with α-orientations of G with $\alpha(v) = 1$ for $v \in A$ and $\alpha(v) = d(v) - 1$ for $v \in B$. This bijection is established by identifying matchings edges with edges directed from A to B. Hence in k-regular bipartite graphs counting perfect matchings is equivalent to counting what we call 1-$(k-1)$-orientations in the sequel.

We observe that the planarization method from the proof of Theorem 3.35 can be used in a more general setting. Let $\mathcal{G}_D$ be the set of all graphs with vertex degrees in $D \subset \mathbb{N}$ and $\mathcal{P}_D$ the set of all planar graphs with degrees in D. Let $I \subset \mathbb{N}$ and associate with every $G \in \mathcal{G}_D$ an out-degree function α_G whose image is contained in I. Then, the proof of Theorem 3.35 shows that counting the α_G-orientations of graphs in $\mathcal{G}_D$ can be reduced to counting α'_G-orientations of the graphs in $\mathcal{P}_{D\cup\{4\}}$ where the image of α'_G is contained in $I \cup \{2\}$ for all $G' \in \mathcal{P}_{D\cup\{4\}}$.

When we apply this to 4-regular bipartite graphs with 1-3-orientations, that is perfect matchings, it yields the first claim of the theorem since counting perfect matchings of bipartite 4-regular graphs is $\#P$-complete by Theorem 3.38.

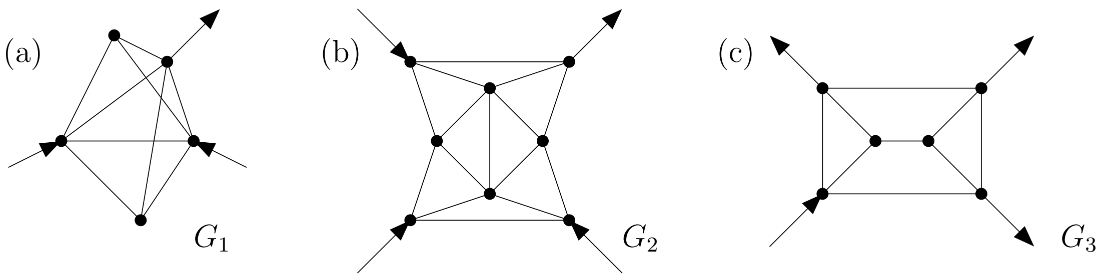

Figure 3.26. The gadgets that translate (a) $\alpha(v) = 1$ to $\alpha \equiv 2$ for $d(v) = 3$, (b) $\alpha(v) = 1$ to $\alpha \equiv 2$ for $d(v) = 4$, and (c) $\alpha(v) = 3$ to $\alpha \equiv 2$ for $d(v) = 4$.

We give two different proofs for the second claim. Let G be a graph with a degree 3 vertex v and α_G an associated out-degree function with $\alpha_G(v) = 1$. We substitute v by the gadget G_1 from Figure 3.26 (a) to obtain a graph G'. The gadget has five vertices that induce nine edges and 3 edges connect it with the neighbors of v in G. Let $\alpha_{G'} = \alpha_G$ on $V(G) - v$ and be $\alpha_{G'}(u) = 2$ for $u \in V(G_1)$. Exactly one of the connection edges must be directed away from G_1 in every $\alpha_{G'}$-orientation of G'. Note that G_1 is symmetric in the three connection vertices. It is easy to check that G_1 has ten orientations with out-degree 2 at every vertex once the outgoing connection edge has been chosen. Thus, every α_G-orientation is associated with ten α'_G-orientations of G' and since the underlying α_G-orientation can be reconstructed from every $\alpha_{G'}$-orientation, we obtain that $r_{\alpha_{G'}}(G') = 10 \cdot r_{\alpha_G}(G)$.

Let a 3-regular bipartite graph G with vertex set $V = A \cup B$ be given and G' be obtained from G by substituting every vertex $v \in A$ by a copy of G_1. The number of 1-2-orientations of G is $10^{-|A|} \cdot r_{\alpha_{G'}}(G')$. Using the planarization method this yields that counting orientations with $\alpha_{G'} \equiv 2$ for graphs from $\mathcal{P}_{\{3,4,5\}}$ is $\#P$-complete.

Similarly counting perfect matchings of 4-regular bipartite graphs can be reduced to counting 2-orientations of graphs from $\mathcal{P}_{\{3,4,5\}}$. We mention this second proof since it uses planar gadgets. More precisely G' is obtained from G by substituting the vertices of one partition class of a 4-regular bipartite graph G by the gadget G_2 shown in Figure 3.26 (b) and the vertices from the other partition class by the gadget G_3 shown in Figure 3.26 (c). Using a similar reasoning as above one obtains that the number of 1-3-orientations of G is $(26 \cdot 6)^{-|A|} r_{\alpha_{G'}}(G')$, where $\alpha_{G'} \equiv 2$, since G_2 and G_3 give a blow-up factor 26 respectively 6. $\square$

Having proved Theorem 3.37 it is natural to ask whether it is $\#P$-complete to count α-orientations for k-regular planar graphs and constant α. This setting implies that $\alpha \equiv k/2$. Planar graphs have average degree less than 6. Furthermore, a 2-regular, connected graph is a cycle and therefore has two Eulerian orientations. Hence the above question should be asked for $k = 4$. Note that the planarization method yields vertices with $d(v) = 4$ and $\alpha(v) = 2$, so we do not need to restrict our considerations to planar graphs when trying to answer the following question. Is it $\#P$-complete to count Eulerian orientations of 4-regular graphs? To the best of our knowledge the following problem is also open. Is it $\#P$-complete to count Eulerian orientations of graphs with degrees in $\{1, \ldots, k\}$, for some arbitrary but fixed $k \in \mathbb{N}$?

We present one more $\#P$-completeness result since it has a nice connection with the first question stated above. On every 4-regular bipartite graph there is a bijection between 2-factors and Eulerian orientations. For 3-regular bipartite graphs 2-factors are in bijection with their complements the 1-factors, i.e. perfect matchings. Hence it is $\#P$-complete to count 2-factors of 3-regular bipartite graphs. The next theorem generalizes this observation.

Theorem 3.39. *For every $i \geq 3$, $i \neq 4$ counting 2-factors of i-regular bipartite graphs is $\#P$-complete.*

Proof. The case $i = 3$ follows from Theorem 3.38 as explained above. The proof for $i \geq 5$ is a reduction from counting 2-factors of 3-regular bipartite graphs, i.e. the case $i = 3$. The method comes from the proof of Theorem 3.38 in [30].

The following preliminary considerations will be needed later. We fix some edge e_0 of the complete bipartite graph $K_{i,i}$ on $2i$ vertices. This graph has i^2 edges and every 2-factor of $K_{i,i}$ has $2i$ edges. Let c_i be the number of 2-factors of $K_{i,i}$. We want to find the ratio between 2-factors of $K_{i,i}$ that contain e_0 and

2-factors that do not contain e_0. Consider pairs of 2-factors and edges (F, e). Obviously, there are $i^2 \cdot c_i$ such pairs and we have $|\{(F,e) \mid e \in F\}| = 2i \cdot c_i$ while $|\{(F,e) \mid e \notin F\}| = (i^2 - 2i) \cdot c_i$. It is obvious that

$$\{(F,e) \mid e \in F\} = \bigcup_{j=1}^{i^2} \{(F, e_j) \mid e_j \in F\}$$

and symmetry implies that all sets $\{(F, e_j) \mid e_j \in F\}$ have the same cardinality. We conclude that there are $a_i = 2c_i/i$ 2-factors including e_0. It follows similarly that there are $b_i = (1 - 2/i)c_i$ 2-factors not including e_0. We infer that $a_i/b_i = 2/(i-2) \leq 2/3$ for $i \geq 5$.

The bridge gadget $P_i(k)$ is a concatenation of k disjoint copies of $K_{i,i} - e_0$, with $k-1$ connection edges as shown in Figure 3.27 for $i = 5$ and $k = 4$. The gadget is connected to the rest of the graph via two edges at the degree $i-1$ vertices.

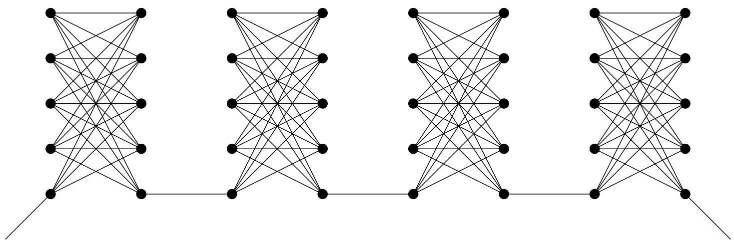

Figure 3.27. The bridge gadget $P_5(4)$

Let G be a 3-regular bipartite graph with $2n$ vertices and $G'(k)$ be obtained from G by augmenting it with $n(i-3)$ disjoint bridge gadgets $P_i(k)$ such that $G'(k)$ is i-regular and bipartite. Let P be some fixed bridge. Note that every 2-factor of $G'(k)$ includes all or none of the connection edges of P. This is because a 2-factor is a partition into cycles and therefore intersects every edge cut of cardinality 2 in either none or both of the edges.

Let c_G be the number of 2-factors of G. Among the 2-factors of $G'(k)$ let S denote the set of those that induce a 2-factor of G, and let S^c be the complement of S. We have $|S| = c_G \cdot b_i^{k \cdot n(i-3)}$ since the 2-factors in S cannot include any connection edges and thus can be partitioned into a 2-factor of G that is augmented by one of the b_i^k possible 2-factors on every bridge.

Next we want an upper bound for $|S^c|$. Every 2-factor in S^c includes the connection edges of at least one bridge P. Since 2^{3n} is the number of subsets of $E(G)$ each of which can be augmented to a 2-factor in at most a_i^k ways on P, we have that $|S^c| \leq a_i^k b_i^{k \cdot (n(i-3)-1)} 2^{3n}$. The number of 2-factors of $G'(k)$ is $c_{G'(k)} = |S| + |S^c|$ and we can bound it as follows.

$$\begin{array}{rlllll} & c_G \cdot b_i^{k\cdot n(i-3)} & \leq & c_{G'(k)} & \leq & c_G \cdot b_i^{k\cdot n(i-3)} + a_i^k b_i^{k\cdot(n(i-3)-1)} 2^{3n} \\ \iff & c_G & \leq & c_{G'(k)} b_i^{-k\cdot n(i-3)} & \leq & c_G + 2^{3n} \left(\frac{a_i}{b_i}\right)^k \end{array}$$

Since $a_i/b_i \leq 2/3$ we conclude that $k > (3n+1)/(\log 3 - 1)$ implies $2^{3n}(a_i/b_i)^k < 1/2$. Note that this bound for k is linear in n. We finally conclude that $c_G = \lfloor c_{G'(k)} b_i^{-k\cdot n(i-3)} \rfloor$ for k large enough. Since $b_i^{-k\cdot n(i-3)}$ is easy to compute and $G'(k)$ has size polynomial in n, this proves the theorem. □

Remark 3.40. We would like to point out that the missing case $k = 4$ in Theorem 3.39 cannot be fixed by substituting the bridge gadget by another gadget. It is crucial that $a_i/b_i < 1$ and symmetry implies that in every 4-regular graph the number of 2-factors including a fixed edge e_0 is equal to the number of 2-factors not containing e_0.

3.5.2 Approximation

Counting α-orientations can be reduced to counting f-factors in bipartite planar graphs and to counting perfect matchings in bipartite graphs. We next describe these transformations. They are useful because bipartite perfect matchings have been the subject of extensive research, see for example [69, 79, 58].

First, note that the α-orientations of M are in bijection with the α'-orientations of the bipartite planar map M'. This map M' is obtained from M by subdividing every edge once and we define $\alpha'(v) = \alpha(v)$ for the original vertices of M and $\alpha'(v) = 1$ for all subdivision vertices. The α'-orientations of M' are in bijection with the f-factors of M' where $f(v) = \alpha'(v)$ for all vertices of M'. The bijection identifies factor edges with edges directed from a vertex of M to an edge-vertex.

The idea for the next transformation is due to Tutte [94]. The graph M' is blown up to a graph M'' such that M'' has $\prod_{v\in M}(d(v) - f(v))!$ times as many perfect matchings, i.e. 1-factors, as there are f-factors of M'. To obtain M'' from M' substitute $v \in V(M)$ by a $K_{d(v),d(v)-f(v)}$, such that each of the $d(v)$ edges incident to v in M' connects to one of the vertices from the partition class of cardinality $d(v)$.

In [58] Jerrum, Sinclair, and Vigoda give a fully polynomial randomized approximation scheme for counting perfect matchings of bipartite graphs. Thus, the above transformation yields a fully polynomial randomized approximation scheme for α-orientations as well.

The number of perfect matchings of a bipartite graph with a Pfaffian orientation can be computed in polynomial time. Little [68] gives a complete characterization of graphs with a Pfaffian orientation and in [79] a polynomial time algorithm to test whether a given graph is Pfaffian is introduced. Little characterizes Pfaffian graphs as those graphs with no central even subdivision of a $K_{3,3}$. In an even subdivision every edge is subdivided by an even number of vertices. An induced subgraph of a graph is central if the rest of the graph has a perfect matching.

As a special case of Little's characterization it follows that all planar graphs are Pfaffian. Hence, in all cases where M'' is planar the counting is easy. For spanning trees the above transformations yields planar graphs, while for Eulerian orientations it does not, as Theorem 3.35 implies. Although we do not have a hardness result for Schnyder woods or bipolar orientations on planar maps, there are in both cases instances for which the transformation yields a non-Pfaffian graph. One such example is the augmented *triangular grid* from Section 3.2.1. *triangular grid*

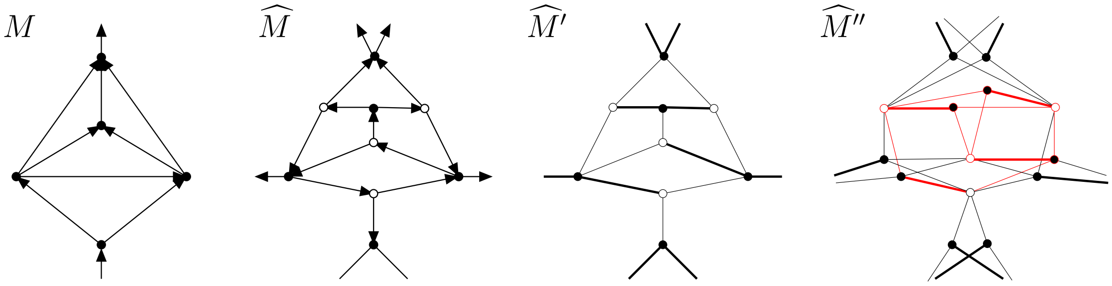

Figure 3.28. An oriented subgraph of a triangulation with a bipolar orientation that induces a central and even subdivision of $K_{3,3}$.

With respect to bipolar orientations Figure 3.28 shows a subgraph of a map M with five vertices and four faces that implies that $\widehat{M}''$ cannot be Pfaffian. The figure shows simplified versions of $\widehat{M}'$ and $\widehat{M}''$. We can choose $\widehat{M}' = \widehat{M}$ since the angle graph is bipartite. The Tutte transformation substitutes face vertices by $K_{1,3}$ and primal vertices of degree d by $K_{2,d}$. Instead one can simply create a copy of every primal vertex with the same neighborhood as the original and leave the face vertices unchanged to obtain a simplified version of $\widehat{M}''$.

3.6 Conclusions

In this chapter we have studied the maximum number of α-orientations for different classes of planar maps and different α. In most cases we have proved exponential upper and lower bounds c_L^n and c_U^n for this number.

The obvious problem is to improve on the constants c_L and c_U for the different classes. We think that in particular improving the upper bound of 8^n for the number of Schnyder woods on 3-connected planar maps is worth further efforts.

For bipolar orientations the more efficient encoding from Theorem 3.31 helps to improve the upper bound. We think that finding a more efficient encoding for Schnyder woods might be needed to substantially improve on the 8^n bound.

Problem 3.41. Improve the upper bound of 8^n for the number of Schnyder woods of a planar map with n vertices.

Results by Lieb [65] and Baxter [8] yield the exact asymptotic behavior of the number of Eulerian orientations for the square and triangular grid on the torus. We have used Baxter's result in Section 3.1.3 to construct a family of planar maps with asymptotically 2.59^n Eulerian orientations. We have also shown that Lieb's and Baxter's results yield upper bounds for the number of 2-orientations on the square grid respectively the number of Schnyder woods on the triangular grid. We have not been able to take advantage of these results for improving the lower bounds for the number of 2-orientations respectively Schnyder woods.

Problem 3.42. Show that the quadrangulation of the grid $G^{\square}_{k,\ell}$ has asymptotically $(8 \cdot \sqrt{3}/9)^{k\ell}$ 2-orientations.

Problem 3.43. Show that the augmented triangular grid $T^*_{k,\ell}$ has asymptotically $(3\sqrt{3}/2)^{k\ell}$ Schnyder woods.

For some instances of α-orientations there are $\#P$-completeness results. This contrasts with spanning trees and planar bipartite perfect matchings for which polynomial algorithms are known. It remains open to determine the complexity of counting Schnyder woods and bipolar orientations on planar maps and of counting Eulerian orientations of graphs with bounded maximum degree.

Problem 3.44. Is it $\#P$-complete to count

(i) Schnyder woods?

(ii) bipolar orientations of planar maps?

(iii) Eulerian orientations of graphs with degrees in $\{1, \ldots, k\}$, for some arbitrary but fixed $k \in \mathbb{N}$?

(iv) Eulerian orientations of 4-regular graphs?

Chapter 4

Spanning Trees with Many Leaves

Spanning trees are one of the topics that appear in every textbook about graph theory. Nevertheless, there are still interesting open problems related to spanning trees and one such problem is the topic of this chapter. We are concerned with finding a spanning tree with the maximum number of leaves for a given graph. The precise formulation as a decision problem is the following.

Max-Leaves Spanning Tree (MaxLeaf): MaxLeaf
INSTANCE: A graph G and an integer k.
QUESTION: Does G have a spanning tree T with at least k leaves?

Different approaches have been proposed for the problem MaxLeaf which is known to be $\mathcal{NP}$-complete, see [53]. One method yields constructive lower bounds for the number of leaves that guarantee a certain fraction of the vertices to become leaves. This method has been applied to various graph classes, see [54, 63, 15]. The results from [54, 63, 15] all yield *extremal* lower bounds, in the sense that examples exist which show that the bounds are tight. When choosing k as a parameter, an algorithm for MaxLeaf is called an *FPT algorithm* (short for *fixed parameter tractable*) if its complexity is bounded by $f(k)g(n)$, where $g(n)$ is a polynomial in the number of vertices of G. FPT algorithms for MaxLeaf have received much attention, see [11, 16, 35, 15]. In [11] Bodlaender gave the first FPT algorithm with a parameter function $f(k)$ of roughly $(17k^4)!$. In [16, 35, 15] the abovementioned extremal lower bounds are used to improve Bodlaender's result. The third approach for MaxLeaf that should be mentioned are approximation algorithms. A 2-approximation is given in [85], and when the input is restricted to cubic graphs, the current best approximation is a 5/3-approximation [27]. Our main contribution in this chapter is a new constructive and extremal lower bound for a very broad graph class. This result can be used to improve the best known time complexity for an FPT algorithm.

We now discuss previous work on finding extremal lower bounds in more detail and explain our contribution. Throughout this chapter all graphs are assumed to be simple unless otherwise stated. In fact the only multi-graph that we use is the K_2 plus one additional edge. The minimum vertex degree of a graph G is denoted by $\delta(G)$. A vertex of degree 1 is called a *leaf*. Since we do not work with $\delta(G)$ edge orientations in this chapter we use the simpler notation uv instead of $\{u, v\}$ to *leaf* denote an edge between vertices u and v. Throughout this chapter we will denote the number of vertices of a graph G by n instead of $n(G)$ whenever it is clear from the context which graph is meant.

Let G be a connected graph with at least two vertices. Linial and Sturtevant [67] and Kleitman and West [63] have shown that every graph G with $\delta(G) \geq 3$ has a spanning tree with at least $n(G)/4 + 2$ leaves, and that this bound is best possible. In [63] Kleitman and West also improve on this bound for graphs of higher minimum degree.

diamond

cubic diamond

The examples showing that the bound $n/4 + 2$ is best possible for graphs of minimum degree 3 all consist of a cycle in which every vertex is replaced by a cubic diamond. A *diamond* is the graph K_4 minus one edge, and an induced diamond subgraph of a graph G is a *cubic diamond* if its four vertices all have degree 3 in G, see Figure 4.1 (a).

(a) (b) (c)

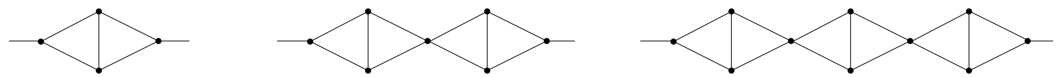

Figure 4.1. A cubic diamond, and concatenations of two respectively three diamonds forming a 2-necklace.

cubic graph

Since these examples are very restricted it is natural to ask if better bounds can be obtained when diamonds are forbidden as subgraphs. This question was answered by Griggs, Kleitman and Shastri [54] for *cubic graphs*, that is graphs where every vertex has degree 3. They show that a cubic graph G without diamonds always admits a spanning tree with at least $n/3 + 4/3$ leaves. This result is best possible, as we will see in Section 4.4.

For minimum degree 3 Bonsma [15] obtains the following bound. A graph G with $\delta(G) \geq 3$, without cubic diamonds, contains a spanning tree with at least $2n/7 + 12/7$ leaves. This bound is also best possible and the examples showing that the linear factor cannot be improved are very similar to those for the $n/4+2$ bound. More precisely these examples can be obtained by replacing every vertex of a cycle by the graph shown in Figure 4.1 (b).

2-necklace

With regard to these latter two results, the following conjecture from [15] seems natural. Every graph G with $\delta(G) \geq 3$ and without 2-necklaces contains a spanning tree with at least $n/3 + 4/3$ leaves. Informally speaking, a *2-necklace* is a concatenation of $k \geq 1$ diamonds with only two outgoing edges, see Figure 4.1. If true, this conjecture would improve the bound $2n/7+12/7$ with only a minor extra restriction, and would also generalize the result for cubic graphs from [54]. In Section 4.1 we disprove this conjecture by constructing graphs with $\delta(G) = 3$ without 2-necklaces that do not admit spanning trees with more than $4n/13+24/13$ leaves.

On the positive side, we prove that the statement is true after only excluding one more very specific structure, called a *2-blossom*, see Figure 4.2 (a). Precise definitions of 2-necklaces and 2-blossoms are given in Section 4.1. Our result is more general than the above mentioned results from [63, 54, 15] since it does not need any restriction on the minimum degree. The exact statement is given in Theorem 4.1.

2-blossom

Let $V_{\geq 3}(G)$ denote the set of vertices in G with degree at least 3 and $n_{\geq 3}(G)$ its cardinality. Let $\ell(T)$ be the number of leaves of a graph T. By Q_3 we denote the graph of the 3-dimensional cube, see Figure 4.2 (b). Furthermore, G_7 is the graph on seven vertices shown in Figure 4.2 (c).

$V_{\geq 3}(G)$, $n_{\geq 3}(G)$
$\ell(T)$
Q_3, G_7

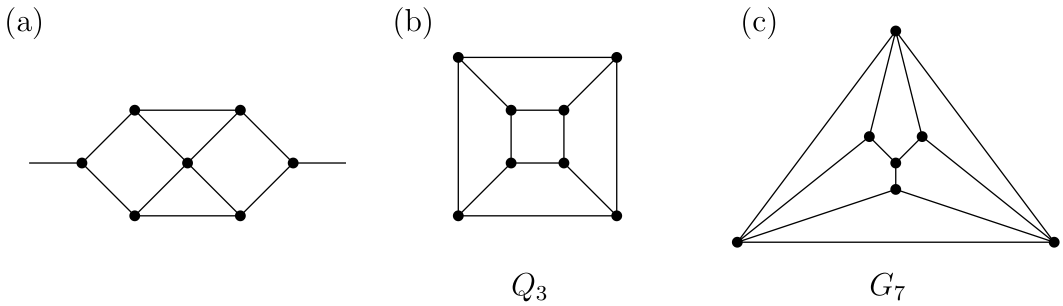

Figure 4.2. A 2-blossom and the graphs Q_3 and G_7.

Theorem 4.1. *Let G be a simple, connected graph on at least two vertices which contains neither 2-necklaces nor 2-blossoms. Then, G has a spanning tree T with*

$$\ell(T) \geq n_{\geq 3}(G)/3 + \begin{cases} 4/3 & \text{if } G = Q_3, \\ 5/3 & \text{if } G = G_7 \text{ or } G \neq Q_3 \text{ is cubic}, \\ 2 & \text{otherwise.} \end{cases}$$

At the beginning of Section 4.4 we show that Theorem 4.1 is best possible. The proof of Theorem 4.1 in Section 4.4 is constructive and can be turned into a polynomial time algorithm for the construction of a spanning tree. Our proof methods extends that of Griggs et al. [54] and makes use of their results. This enables us to strengthen and streamline known results and to shorten the proof of Theorem 4.1 by incorporating a lemma from [54] into our proof. In Section 4.4.2 we argue that the long case study in [54] actually proves this strong new lemma, which we then use as an important step in the proof of Theorem 4.1. We share the opinion expressed in [54] that a shorter proof for the bound for cubic graphs might not exist. Therefore using that result in order to prove the more general statement seems appropriate.

Graph class	Bound	Reference
$\delta(G) \geq 3$	$n/4 + 2$	[63]
$\delta(G) \geq 4$	$2n/5 + 8/5$	[63]
$\delta(G) \geq 5$	$n/2 + 2$	[55]
unrestricted	$n_{\neq 2}/4 + 3/2$	Theorem 4.7
$\delta(G) \geq 3$, no triangles	$n/3 + 4/3$	[15]
no triangles	$n_{\neq 2}/3 + 2/3$	Theorem 4.8
$\delta(G) \geq 3$, no cubic diamonds	$2n/7 + 12/7$	[15]
no cubic diamonds	$2n_{\geq 3}/7 + 12/7$	Theorem 4.9
cubic graphs, no cubic diamonds	$n/3 + 4/3$	[54]
no 2-necklaces	$4n_{\geq 3}/13 + 20/13$	Theorem 4.5
no 2-necklaces, no 2-blossoms	$n_{\geq 3}/3 + 4/3$	Theorem 4.1

Table 4.1. An overview of known lower bounds for different graph classes.

$n_{\neq 2}(G)$ Table 4 shows an overview of known extremal results and the new bounds that we prove in this chapter. We denote by $n_{\neq 2}(G)$ the number of vertices of G that do not have degree 2. With regard to the first three result from Table 4 we mention a conjecture which is attributed to N. Linial in [54]. For every $k \geq 3$ and every graph G with $\delta(G) \geq k$ there is a spanning tree T with $\ell(T) \geq (k-2)n/(k+1)+d$ for some $d \geq 0 \in \mathbb{R}$. In [22] it is remarked that a result obtained by Alon [5] with probabilistic methods disproves this conjecture for large k.

We explain the basic proof method that we use to prove Theorem 4.1 in Section 4.2 and demonstrate it by proving a generalization of the $n/4 + 2$ bound from [63]. Similarly we generalize a bound from [15] for graphs without triangles.

In Section 4.3 we give a new proof of Bonsma's abovementioned $2n/7 + 12/7$ result from [15]. This result can also be obtained as a corollary of Theorem 4.1, see Section 4.4.4. We include Section 4.3 to demonstrate the use of graph reductions in a relatively simple setting by giving a short self-contained proof of this bound. Graph reductions will also be used to prove Theorem 4.1.

Section 4.4 is devoted to the proof of Theorem 4.1. In Section 4.4.4 we show how Theorem 4.1 can be strengthened, in the sense that we do not ask for the graph to have no 2-necklaces and no 2-blossoms, but the bound becomes weaker depending on the number of 2-necklaces and 2-blossoms. We also show that every graph G without 2-necklaces, but possibly with 2-blossoms, has a spanning tree with at least $4n_{\geq 3}(G)/13 + 20/13$ leaves.

We have mentioned above that Theorem 4.1 can be applied to obtain a fast FPT algorithm for MaxLeaf. We now provide some background and state the result. However we omit the details and instead refer the reader to [17]. Recall that an algorithm for MaxLeaf with k as a parameter, is called an *FPT algorithm* if its complexity is bounded by $f(k)g(n)$, where $g(n)$ is a polynomial. See [49] and [34] for introductions to FPT algorithms. The function $f(k)$ is called the *parameter function* of the algorithm. Usually, $g(n)$ will turn out to be a low degree polynomial. Hence, to assess the speed of the algorithm it is mainly important to consider the growth rate of $f(k)$. Since MaxLeaf is $\mathcal{NP}$-complete, we content ourselves with exponential bounds for $f(k)$. Bodlaender [11] constructed the first FPT algorithm for MaxLeaf with a parameter function of roughly $(17k^4)!$. Since then, considerable effort has been put in finding fast FPT algorithms for this problem, see e.g. [33, 36, 16, 35, 15]. In [16, 35, 15] a strong connection between extremal graph-theoretic results and fast FPT algorithms is established. In [16], the bound of $n/4+2$ from [63] mentioned above is used to find an FPT algorithm with parameter function $O^*(\binom{4k}{k}) \subset O^*(9.49^k)$. Here the O^* notation ignores polynomial factors. With the same techniques the bound of $2n/7+12/7$ is turned into the so far fastest algorithm, with a parameter function in $O^*(\binom{3.5k}{k}) \subset O^*(8.12^k)$, see [15]. Similarly Theorem 4.1 yields a new FPT algorithm for MaxLeaf.

Theorem 4.2. *There is an $O(m) + O^*(6.75^k)$ FPT algorithm for* MaxLeaf, *where m denotes the size of the input graph and k the desired number of leaves.*

This algorithm is the fastest FPT algorithm for MaxLeaf at the moment, both optimizing the dependency on the input size and the parameter function. It simplifies the ideas introduced by Bonsma, Brueggemann and Woeginger [16] and is also significantly simpler than the other recent fast FPT algorithms. Hardly any preprocessing of the input graph is needed, since Theorem 4.1 is already formulated for a very broad graph class. For further details concerning Theorem 4.2 we refer the reader to [17].

4.1 Obstructions for Spanning Trees with Many Leaves

As mentioned above, 2-necklaces have been identified as an obstruction for the existence of spanning trees with $n/3+c$ leaves in graphs with minimum degree 3, see [63] and [15]. In this section we show that they are not the only such obstruction. We start by precisely defining 2-necklaces and 2-blossoms.

In order to avoid confusion we sometimes denote the degree of a vertex v in a graph G by $d_G(v)$ in this chapter. If ambiguities can be excluded we use $d(v)$ as in earlier chapters. A vertex v of a subgraph H of G with $d_H(v) < d_G(v)$ is called a *terminal* of H.

$d_G(v)$

terminal

diamond
diamond necklace N_k
connection vertex
inner vertex of N_k

As mentioned above, the graph K_4 minus one edge is called a *diamond* and denoted by N_1. For $k \geq 2$ the *diamond necklace* N_k is obtained from the graph N_{k-1} and a vertex disjoint N_1 by identifying a degree 2 vertex of N_1 with a degree 2 vertex of N_{k-1}. The *connection vertices* of the graph N_k are its two degree 2 vertices. A vertex of degree 3 or 4 is an *inner vertex of* N_k. Diamond necklaces will also be called *necklaces* for short.

2-necklace
cubic diamond

An N_k subgraph of G is a *2-necklace* if only its connection vertices are terminals, and they both have degree 3 in G, see Figure 4.1. We have already mentioned that if G contains an N_1 this way, then this N_1 subgraph is also called a *cubic diamond* of G.

blossom graph

In the course of studying the leafy tree problem we found that the subgraphs that we define next are also an obstacle for the existence of spanning trees with $n/3 + c$ leaves. The graph B on seven vertices shown in Figure 4.3 (a) is the *blossom graph*.

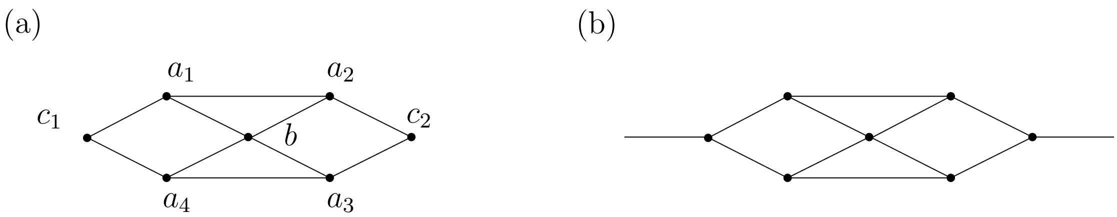

Figure 4.3. A blossom graph and a 2-blossom.

connection vertex
2-blossom

The vertices named c_1 and c_2 in the figure are the *connection vertices* of B. A blossom subgraph B of G is a *2-blossom* if the connection vertices are its only terminals, and they both have degree 3 in G, see Figure 4.3 (b). If G contains a 2-blossom B, only the vertex b has degree 4 in G, and the remaining vertices of B have degree 3 in G.

The two outgoing edges of a 2-necklace respectively a 2-blossom may in fact be the same edge. In that case G is just a 2-necklace respectively 2-blossom plus one additional edge. The next lemma shows how many leaves can be gained within a blossom.

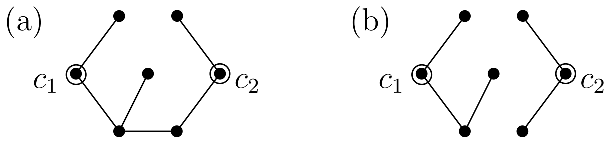

Figure 4.4. Spanning trees restricted to a blossom.

Lemma 4.3. *Let G be a graph with a blossom subgraph B that has c_1 and c_2 as its only terminals, see Figure 4.3 (a). Then, a spanning tree T of G with the maximum number of leaves exists, such that $E(T) \cap E(B)$ has one of the forms shown in Figure 4.4.*

Proof. Consider a spanning tree T of G with maximum number of leaves. We may distinguish the following two cases for $E(T) \cap E(B)$, since c_1 and c_2 are the only terminals of B. Either $E(T) \cap E(B)$ induces a tree, or it induces a forest with two components, one containing c_1 and the other containing c_2.

In the first case, at most three non-terminal vertices of B can be leaves of T, since a path from c_1 to c_2 contains at least two internal vertices. In addition, if one of c_1 and c_2 is a leaf of T, then T can be seen to have at most two non-terminal vertices of B among its leaves. Since c_1 and c_2 together form a vertex cut of G, one of them is not a leaf in T. It follows that replacing $E(T) \cap E(B)$ by the edge set in Figure 4.4 (a) does not decrease the number of leaves. Since this edge set forms again a spanning tree of B, the resulting graph is a spanning tree of G.

Now suppose $E(T) \cap E(B)$ has two components. At most four non-terminal vertices of B can be leaves of T. If one of c_1 and c_2 is a leaf of T, then T can have at most three non-terminal vertices of B among its leaves. One of c_1 and c_2 is not a leaf in T, and thus it follows again that replacing $E(T) \cap E(B)$ by the edge set in Figure 4.4 (b) does not decrease the number of leaves, while maintaining a spanning tree of G. □

We now present a family of graphs with minimum degree 3 that do not contain 2-necklaces but do not have spanning trees with $n/3 + c$ leaves.

The *flower graph* is the graph on ten vertices shown in Figure 4.5 (a). A flower subgraph F of G is a *1-flower* if c is its only terminal and $d_G(c) = 3$ in G, see Figure 4.5 (b). From every ternary tree we can obtain a *flower tree* by substituting every inner vertex by a cubic triangle and every leaf by a 1-flower. Figure 4.5 (c) shows a flowertree and the solid edges form a spanning tree with $4n/13 + 24/13$ leaves which we will show to be optimal.

flower graph

1-flower

flower tree

Proposition 4.4. *A flowertree G with n vertices has no spanning tree with more than $4n/13 + 24/13$ leaves.*

Proof. Let F be a 1-flower in G containing a 2-blossom B, and f_1, f_2, c the other three vertices of F, see Figure 4.5 (b). We will argue that no spanning tree T of G has more than four leaves among $V(B) \cup \{f_1, f_2\}$.

Using Lemma 4.3 we may assume without loss of generality that $E(T) \cap E(B)$ has one of the two forms in Figure 4.4. If it has the first form, then either f_1 or f_2 may be a leaf of T, but not both since together they form a vertex cut of G. If $E(T) \cap E(B)$ has the second form, then f_1 and f_2 are both cut vertices of T, so neither can be a leaf.

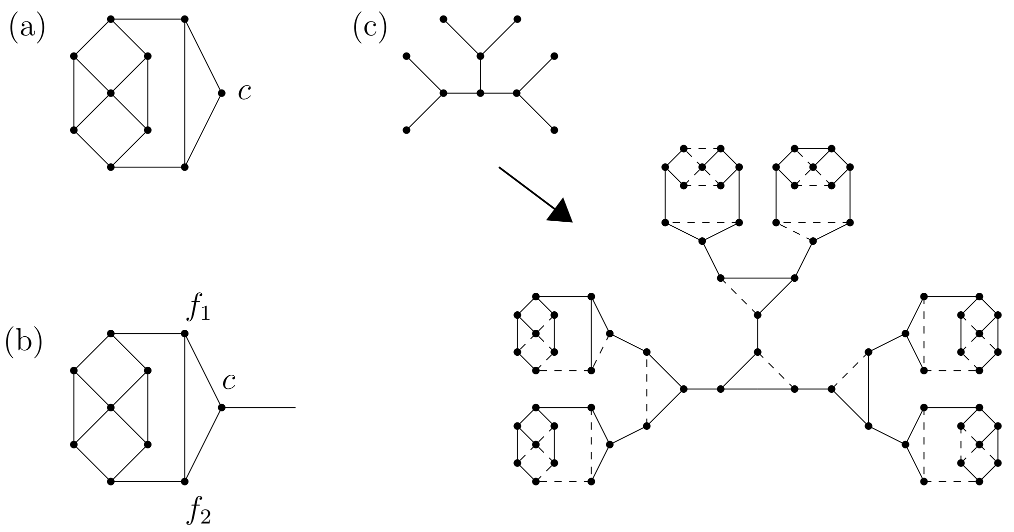

Figure 4.5. A flower graph, a 1-flower, and a flower tree with six 1-flowers. The solid edges show a tree with maximum number of leaves.

Note that the connection vertices of the 1-flowers and all vertices of G that are not part of a 1-flower are cut vertices of G. Hence, none of them can be a leaf of T. Thus, T has no more than $4i$ leaves, where i is the number of leaves of the ternary tree corresponding to G. Since a ternary tree with i leaves has $i-2$ inner vertices, G has $n = 10i + 3(i-2) = 13i - 6$ vertices. This proves the claim since $4i = 4(n+6)/13$. $\square$

The following theorem shows that every graph without 2-necklaces has a spanning tree with at least $4/13n + 20/13$ leaves. We will prove it in Section 4.4.4 as a corollary of Theorem 4.1.

Theorem 4.5. *Let G be a simple, connected graph on at least two vertices, which contains no 2-necklaces. Then, G has a spanning tree T with*

$$\ell(T) \geq 4n_{\geq 3}(G)/13 + \begin{cases} 20/13 & \text{if } G \text{ is cubic} \\ 24/13 & \text{otherwise.} \end{cases}$$

4.2 Introduction to the Proof Method

The purpose of this section is to introduce the proof method that we use throughout this chapter. Theorem 4.6 is proven for cubic graphs by Griggs, Kleitman and Shastri in [54]. The stronger statement given in Theorem 4.6 was proved by Kleitman and West in [63].

Theorem 4.6. *Every graph G with minimum degree 3 has a spanning tree T with*

$$\ell(T) \geq n(G)/4 + 2.$$

The proofs in [54] and in [63] both use the same method. Our proofs of Theorems 4.1, 4.7, 4.8, and 4.9 are extensions of this method. These extensions enable us to obtain some new results and strengthen some old ones. For example, Theorem 4.6 can be strengthened to the following claim, using the method we describe below.

Theorem 4.7. *Every graph G with at least two vertices has a spanning tree T with*

$$\ell(T) \geq n(G)_{\neq 2}/4 + 3/2.$$

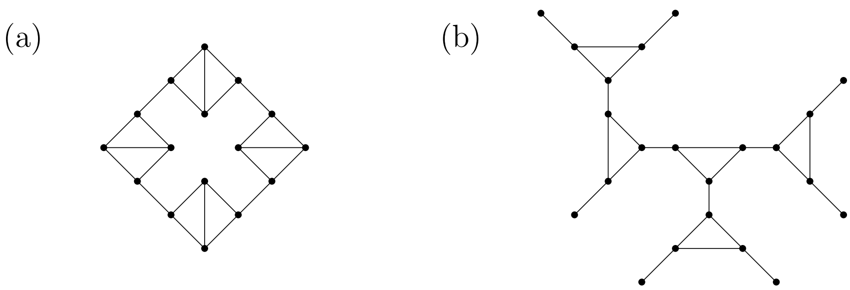

Figure 4.6. Examples for the tightness of Theorems 4.6 and 4.7.

The example from Figure 4.6 (a) shows that the bound from Theorem 4.6 is tight. Infinitely many examples can be obtained by replacing every vertex of a cycle by a cubic diamond. The additive term in Theorem 4.7 is worse than that in Theorem 4.6. The example in Figure 4.6 (b) shows that an additive term better than $3/2$ cannot be achieved in the setting of Theorem 4.7. Infinitely many examples can be obtained by replacing every inner vertex of a ternary tree by a triangle with three degree 3 vertices.

We now describe the generic proof method which we use for many proofs in this chapter, and then demonstrate it by giving a proof for Theorem 4.7. The description we give now is meant to explain the unifying idea behind the proofs in this section and thereby help the reader navigate through the proofs, since some of them are rather long and technical. But since this is only a sketch of the method, each proof will need some additional tricks and techniques which we neglect here.

Let $\mathcal{G}$ be a graph family and let $\mathcal{G}_n$ the graphs in $\mathcal{G}$ that have n vertices. In the instances we consider later the graph families will be defined by excluding certain subgraphs. For a finite set of integers I we define

$$N_I(G) = \{v \in V(G) \mid d_G(v) \notin I\}.$$ $N_I(G)$

$n_I(G)$ The cardinality of $N_I(G)$ is denoted by $n_I(G) = |N_I(G)|$. We will use two different instances of the set I in this chapter, namely $I = \{2\}$ and $I = \{1, 2\}$. For $n_{\{2\}}(G)$ we also use the notation $n_{\neq 2}(G)$, and $n_{\{1,2\}}(G)$ we denote by $n_{\geq 3}(G)$. Let $a, c > 0$ be constants. The theorems we prove in this chapter are of the following form.

Every $G \in \mathcal{G}_n$ has a spanning tree T with $\ell(T) \geq n_I(G)/a + c$.

Before we give a generic proof outline for a theorem of this type, we introduce the terminology that we use in the proofs throughout the chapter.

$F \subseteq G$, $F \subset G$ When F and G are graphs, $F \subseteq G$ and $F \subset G$ denote the subgraph and proper subgraph relation, respectively. We denote the set of vertices of G which are not in F by

$\overline{V(F)}$
$$\overline{V(F)} = V(G) \setminus V(F).$$

For a subgraph $F \subseteq G$ we define the subgraph of G outside of F as an edge induced subgraph

F^C
$$F^C = G[\{uv \in E(G) : u \notin V(F)\}].$$

boundary The *boundary of a graph* F is $V(F) \cap V(F^C)$. Note that no edges between two vertices that are both in $V(F)$ are included in F^C. If G is clear from the context graph outside F we call F^C the *graph outside* F. Figure 4.7 shows an example, the subgraph F of the Petersen graph P is indicated by solid edges. The encircled vertices are in $\overline{V(F)}$ and the boxed ones are on the boundary of F.

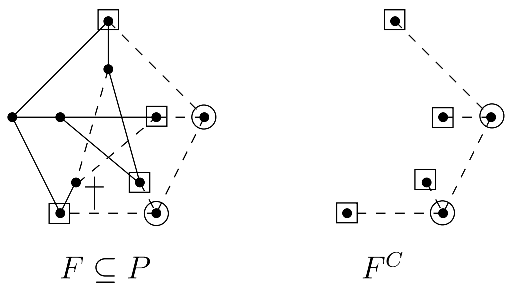

Figure 4.7. An example of a subgraph $F \subseteq P$ with a dead leaf marked by a cross and F^C the graph outside F.

dead leaf A leaf of F is a *dead leaf* if it has no neighbor in $\overline{V(F)}$, and the number of dead $\ell_d(F)$ leaves of F is denoted by $\ell_d(F)$. The subgraph F of P in Figure 4.7 has one dead leaf which is indicated by a cross. We denote the number of vertices of F which $n_{I,G}(F)$ are in $N_I(G)$ by $n_{I,G}(F)$, e.g. in Figure 4.7 $n_{\{1\},P}(F) = 8$ since no vertex of F $cc(G)$ has degree 1 in P. By $cc(G)$ we denote the number of connected components of a graph G.

We now start the actual description of our proof method and choose constants $a_1, a_2 \geq 0$ with $a_1 + a_2 = a$. The *leaf potential* of F is defined as *leaf potential*

$$\mathcal{P}(F,G) = a_1\ell(F) + a_2\ell_d(F) - n_{I,G}(F) - 2a \cdot cc(F).$$

$\mathcal{P}(F,G)$

In the context of the proofs it will always be clear what the values of a_1 and a_2 are. Counting the fraction a_2 of the total weight a of a leaf not until it has no more neighbors in $\overline{V(F)}$ can be interpreted as an amortization technique. This amortized counting was used in [54] and it is one of the main advantages of this proof method.

The subgraph F of G is *extendible* if there exists another subgraph F' with $F \subset F' \subseteq G$ and $\mathcal{P}(F',G) \geq \mathcal{P}(F,G)$. Given $F \subset F' \subseteq G$ we denote by $\Delta n_{I,G}$ the difference $n_{I,G}(F') - n_{I,G}(F)$ and similarly $\Delta\ell$ and $\Delta\ell_d$ are defined. *extendible* $\Delta n_{I,G}$, $\Delta\ell$, $\Delta\ell_d$

We now sketch the proof method. First we find an initial subgraph $F \subset G$ with $\mathcal{P}(F,G) \geq a(c-2)$. Then, we show that every subgraph F of G is extendible. It will be convenient to use the following abbreviation.

$$\Delta(x,y,z) = a_1 y + a_2 z - x$$

$\Delta(x,y,z)$

In order to show that F is extendible we explicitly construct F' such that

$$\begin{aligned}
&\Delta(\Delta n_{I,G}, \Delta\ell, \Delta\ell_d) \geq 0 \quad &&\text{if} \quad cc(F') = cc(F) \qquad &&\text{and}\\
&\Delta(\Delta n_{I,G}, \Delta\ell, \Delta\ell_d) \geq 2a \quad &&\text{if} \quad cc(F') = cc(F) + 1.
\end{aligned}$$

This implies $\mathcal{P}(F',G) \geq \mathcal{P}(F,G)$. The constructions for F' are the technical part of the proofs and typically require the most work. Using these constructions, we extend F until we obtain a spanning subgraph $F \subseteq G$. In a spanning subgraph we have $\ell(F) = \ell_d(F)$, and since $\mathcal{P}(F) \geq a(c-2)$ we obtain

$$a \cdot \ell(F) \geq n_{I,G}(F) + 2a \cdot cc(F) + a(c-2).$$

A spanning subgraph with $cc(F)$ components can be turned into a connected spanning subgraph by adding $cc(F) - 1$ edges. Each edge addition destroys at most two leaves. We may transform the resulting graph into a spanning tree T, since any cycles can be broken by edge deletions without destroying leaves. Thus, we have obtained a spanning tree T with

$$\begin{aligned}
a \cdot \ell(T) &\geq n_{I,G}(F) + 2a \cdot cc(F) + a(c-2) - 2a(cc(F)-1)\\
\Rightarrow \qquad \ell(T) &\geq n_{I,G}(F)/a + c.
\end{aligned}$$

We want to apply this method to prove Theorem 4.7. Before we do so, we introduce some more notation. The statement 'a is adjacent to b' is denoted by $a \sim b$. We write $a \sim F$ if there exists $b \in F$ with $a \sim b$. $a \sim b$, $a \sim F$

$N(v)$, $N[v]$ The neighborhood $N(v)$ of a vertex v is the set of all vertices adjacent to v, and the closed neighborhood of v is $N[v]= N(v) \cup \{v\}$. The operation of *expanding a vertex* $v \in V(G)$ is defined on a subgraph $F \subset G$ and yields a new subgraph with vertex set $V(F) \cup N[v]$, and edge set $E(F) \cup \{uv : u \in N(v) \setminus V(F)\}$. So all newly added neighbors of v become leaves, and v may lose leaf status. The number of components increases by one if and only if $v \notin V(F)$. Expanding a list of vertices means expanding the vertices in the given order.

expanding a vertex

In order to keep the notation in the proofs simple, instead of writing for example $\Delta(\Delta n_{I,G}, \Delta\ell, \Delta\ell_d) \geq \Delta(4, 3, 1) = 4$, we will simply write $\Delta(4, 3, 1) = 4$. Hence the three parameter values need not be exactly $\Delta n_{I,G}$, $\Delta\ell$ and $\Delta\ell_d$ but reflect the worst case scenario. That is, the change in the leaf potential that we prove is to be read as a lower bound for the actual change.

Proof of Theorem 4.7. The proof is based on that of Theorem 4.6 in [63]. The theorem obviously holds for graph with maximum degree 2, that is paths and cycles, and the K_4. So we may assume that the graph G has maximum degree at least 3 and is not the K_4.

In the theorem we have $a = 4$, $c = 3/2$ and $I = \{2\}$. We choose $a_1 = 3$ and $a_2 = 1$, thus the definition of the *leaf potential* that we use here is

leaf potential

$$\mathcal{P}(F, G) = 3\ell(F) + \ell_d(F) - n_{\neq 2,G}(F) - 8cc(F)$$

and we define $\Delta(x, y, z) = 3y + z - x$.

We now construct the initial graph F which must have $\mathcal{P}(F, G) \geq -2$, i.e. $\Delta(n_{\neq 2,G}(F), \ell(F), \ell_d(F)) \geq 6$. First assume that G has a vertex v of degree 1. If v has a neighbor of degree 2, then expanding v yields $\Delta(1, 2, 1) = 6$, which suffices. If v has a neighbor w of degree at least 3, then expanding w yields $\Delta(4, 3, 1) = 6$. Thus, we may assume that $\delta(G) \geq 2$. If $\delta(G) = 2$, then there is a vertex v with $d(v) = 2$ that has a neighbor w with $d(w) \geq 3$. Thus, expanding w yields $\Delta(3, 3, 0) = 6$. We may therefore assume that $\delta(G) \geq 3$. If there is a vertex v with $d(v) \geq 4$, then expanding it yields $\Delta(5, 4, 0) = 7$, so the remaining case is that G is cubic. Since $G \neq K_4$, it has an edge $e = uv$ that is in no triangle. Expanding u, v yields $\Delta(6, 4, 0) = 6$. This concludes the initialization phase.

We now show that every non-empty subgraph F of G is extendible. If F has a non-leaf vertex v that is adjacent to a vertex from $\overline{V(F)}$, then expanding v implies $\Delta(1, 1, 0) = 2$. Thus, we may assume that all vertices on the boundary of F are leaves of F. First assume that there exists $v \in \overline{V(F)}$ with $v \sim F$ and $d(v) \geq 3$. We show that in this case one of the three operations from the proof in [63] can be applied to extend F. If there is $u \in F$ with two neighbors in $\overline{V(F)}$, then expanding u yields $\Delta(2, 1, 0) = 1$. So now we assume that every $v \in F$ has at most one neighbor in $\overline{V(F)}$. If there is $v \in \overline{V(F)}$ with two neighbors in F, then expanding one of those neighbors yields $\Delta(1, 0, 1) = 0$. If $v \in \overline{V(F)}$

has one neighbor u in F and two neighbors in $\overline{V(F)}$, then expanding u, v yields $\Delta(3,1,0) = 0$. One of these operations must be applicable if there is $v \in \overline{V(F)}$ with $v \sim F$ and $d(v) \geq 3$.

Now assume that there is $v \in \overline{V(F)}$ with $d(v) = 2$ and $v \sim u \in F$. Since v is the only neighbor of u in $\overline{V(F)}$, expanding u yields $\Delta(0,0,0) = 0$. Finally, if there is $v \in \overline{V(F)}$ with $d(v) = 1$ and $v \sim u \in F$, then expanding u yields $\Delta(1,0,1) = 0$.

Thus, we have shown that every non-empty $F \subseteq G$ is extendible. This yields the existence of a spanning subgraph F of G with $\mathcal{P}(F,G) \geq -2$. The rest of the proof proceeds as described above. □

Theorem 4.8 is a version of a theorem from Bonsma's thesis [15]. He applies a completely different proof method which uses connected dominating sets. Note that the non-leaf vertices of every spanning tree form a connected dominating set and vice versa. For more about connected dominating sets, we refer the reader to [22, 56]. Our proof of Theorem 4.8 is substantially shorter than that in [15].

Theorem 4.8. *Let G be a connected graph without triangles. Then, G has a spanning tree T with*

$$\ell(T) \geq n_{\neq 2}(G)/3 + 2/3.$$

Bonsma's result is that every graph without triangles and with minimum vertex degree 3 has a spanning tree with at least $n/3 + 4/3$ leaves. Thus, Theorem 4.8 generalizes it in a similar way that Theorem 4.7 generalizes Theorem 4.6.

The additive term 4/3 that Bonsma proves is best possible. The additive term that we prove is weaker, and in this case we do not have an example which shows that it is not possible to improve it to 4/3. Indeed we think, that $n_{\neq 2}(G)/3 + 4/3$ might be the right bound. Besides the examples that Bonsma gives, there are others which show that the bound $n_{\neq 2}(G)/3 + 4/3$ would be best possible. They can be obtained from quaternary trees by substituting every inner vertex by a 4-cycle. The reason that we cannot prove an additive constant of 4/3 might be that we use $a_1 = 2$ and $a_2 = 1$ in the definition of the leaf potential. The small value of a_1 makes it harder to obtain a good initial graph. On the other hand choosing $a_2 = 1$ enables us to prove a bound in terms of $n_{\neq 2}$ instead of $n_{\geq 3}$. We explain now, why we think that this strengthening of the bound is interesting.

Theorem 4.1 shows that for graphs of minimum degree 3 not all triangles have to be forbidden in order to obtain a bound of the form $n_{\geq 3}/3 + 4/3$. Only triangles that appear in 2-necklaces and 2-blossoms have to be forbidden. In contrast to that consider the example from Figure 4.6 (b) which shows that Theorem 4.7 is tight. Since this example contains no 2-necklaces and 2-blossoms, it shows that it is not sufficient to exclude these structures in order to obtain a bound of the form $n_{\neq 2}/3 + c$.

Proof of Theorem 4.8. We may assume that G is neither a path nor a cycle, since then the statement is obviously true. The theorem uses $a = 3$, $c = 2/3$, and $I = \{2\}$.

leaf potential We choose $a_1 = 2$ and $a_2 = 1$, thus the definition of the *leaf potential* that we use here is

$$\mathcal{P}(F, G) = 2\ell(F) + \ell_d(F) - n_{\neq 2,G}(F) - 6cc(F)$$

and we define $\Delta(x, y, z) = 2y + z - x$.

The construction of the initial graph F with $\mathcal{P}(F, G) \geq -4$ is easy. Since G has at least one vertex v with $d(v) \geq 3$, expanding v yields $\Delta(4, 3, 0) = 2$.

We now show that every non-empty $F \subset G$ is extendible. If F has a non-leaf vertex v that is adjacent to a vertex from $\overline{V(F)}$, then expanding v gives $\Delta(1, 1, 0) = 1$. Thus, we may assume that all vertices on the boundary of F are leaves of F. Each of the augmentation rules (A1)–(A9) from Figure 4.8 shows a possibility of extending F with $\Delta(\Delta n_{\neq 2,G}, \Delta\ell, \Delta\ell_d) \geq 0$. White vertices have degree 2, vertices marked with a cross are dead leaves. For all other vertices the shown degrees are to be understood as lower bounds. The encircled vertices already belong to F, the solid edges show how F is extended, while dashed edges are not in the tree.

We assume from now on, that none of (A1) – (A9) is applicable. Hence, every $v \in F$ has at most one neighbor that is not in F by rule (A3). Using the rules (A1) – (A5), we deduce that every $v \in \overline{V(F)}$, $v \sim F$ has degree 3 and one neighbor in F and two neighbors in $\overline{V(F)}$. Furthermore, the rules (A6) – (A9) imply that the neighbors $u, w \in \overline{V(F)}$ of $v \sim F$, both have degree 3 and are not adjacent to F. Since G is triangle free $u \not\sim w$ and both vertices have two neighbors outside $\{u, v, w\}$.

We now explain the intuition behind the rest of the proof. Suppose v, u, w are as defined above and $v \sim s \in F$. Then we may expand s, v thereby obtaining a graph F_1 that has $\mathcal{P}(F_1, G) = \mathcal{P}(F, G) - 1$. If u has a neighbor of degree 2, then we expand u as well and thereby obtain F_2 with $\mathcal{P}(F_2, G) \geq \mathcal{P}(F, G)$. The expansion of u very much resembles (A1) above, and similar considerations can be made for the other rules. In this way it can be shown that $\mathcal{P}(F_2, G) \geq \mathcal{P}(F, G)$, unless u has two neighbors other than v in $\overline{V(F)}$ that have degree 3 and are not adjacent to F. And in the remaining case we have that $\mathcal{P}(F_2, G) = \mathcal{P}(F_1, G) = \mathcal{P}(F, G) - 1$, that is F_2 has the same leaf potential as F_1. In this fashion we may continue expanding vertices without further decreasing the leaf potential, and since G is finite we must eventually terminate with some F_i. Then, either $\mathcal{P}(F_i, G) \geq \mathcal{P}(F, G)$, or $\mathcal{P}(F_i, G) = \mathcal{P}(F, G) - 1$ and, since G is triangle-free, there are additional edges between the vertices of $V(F_i)$. We will then use these edges to show that we can obtain a graph F' with $\mathcal{P}(F', G) \geq \mathcal{P}(F, G)$. We now express this idea more formally, and include the necessary case distinctions.

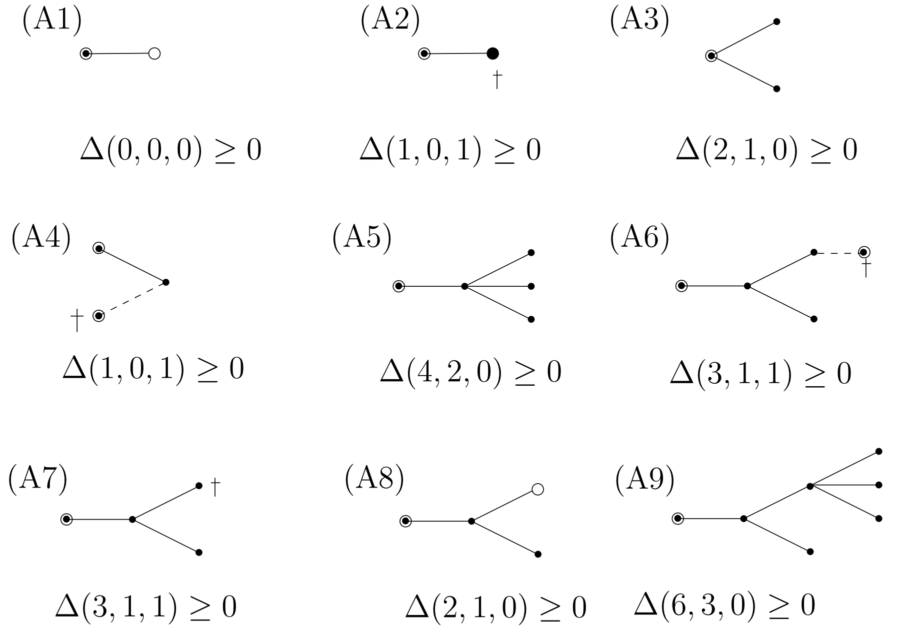

Figure 4.8. The augmentation rules for the proof of Theorem 4.8.

In the rest of the proof we use the notation $V_i = \{v_1, \dots, v_i\}$, $V_i^+ = \{v_0\} \cup V_i$ and $W_i = \{w_1, \dots, w_i\}$. Let $P = v_1, \dots v_{k+1}$ be a path in F^C with $v_1, \dots, v_{k+1} \notin F$ and $v_1 \sim v_0 \in F$ that is maximal under the condition that it has the following properties. The vertex v_1 has a neighbor $w_1 \notin V(F) \cup V_{k+1}$, and v_i has a neighbor $w_i \notin V(F) \cup V_{k+1} \cup W_{i-1}$ for all $i = 1, \dots, k$. Since none of (A1) – (A9) is applicable such a path exists, and $k \geq 2$. Since the path is maximal, v_{k+1} and w_k both have at most one neighbor that is not in $V(F) \cup V_{k+1} \cup W_k$.

If some v_i has a neighbor $w \notin V(F) \cup V_{i+1} \cup W_i$, then we can expand $v_0, \dots v_i$ to obtain $\Delta(2i+2, i+1, 0) = 0$. We will say in the sequel, that we expand V_i^+, instead of giving the list $v_0, \dots v_i$. If some w_i or v_{k+1} has degree 2, we obtain $\Delta(2k, k, 0) = 0$ by expanding V_k^+. If some w_i has at least three neighbors not in $V(F) \cup V_{i+1} \cup W_i$ when we expand V_i^+, w_i to obtain $\Delta(2i+4, i+2, 0) \geq 0$. If some $v_i \in V_{k+1} \backslash \{v_1\}$ or some $w_i \in W_k$ is adjacent to F, then we obtain $\Delta(2k+1, k, 1) \geq 0$ by expanding V_k^+. This is also possible if some $w_i \in W_k$ or v_{k+1} has no neighbor outside $V(F) \cup V_{k+1} \cup W_k$. Hence, no vertex from $(V_{k+1} \setminus \{v_1\}) \cup W_k$ is adjacent to F. The vertices v_i with $i \leq k$ have no neighbor outside $V_{k+1} \cup W_k$ and all w_i have one or two neighbors outside $V(F) \cup V_{k+1} \cup W_k$.

From the considerations in the last paragraph it also follows that v_{k+1} and w_k each have one neighbor z respectively z' that is not in $V_{k+1} \cup W_k$. Throughout the rest of the proof $z = z'$ is allowed. Since G is triangle-free $v_{k+1} \not\sim w_k$ and the vertices v_{k+1} and w_k must each have at least one neighbor in W_{k-1}. Suppose v_{k+1} has at least two neighbors $w_s, w_t \in W_{k-1}$, where $s < t$. As mentioned above w_s has

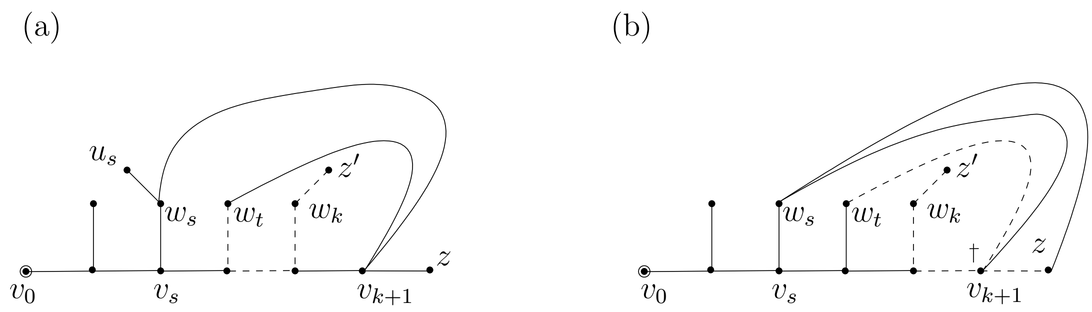

Figure 4.9.

a neighbor $u_s \notin V(F) \cup V_{k+1} \cup W_k$. If $u_s \neq z$, then expanding V_s^+, w_s, v_{k+1} yields $\Delta(2s+6, s+3, 0) = 0$, see Figure 4.9 (a). If $u_s = z$, then expanding V_{k-1}^+, w_s yields $\Delta(2k+1, k, 1) = 0$, since v_{k+1} becomes a dead leaf, see Figure 4.9 (b). Hence, v_{k+1} and w_k each have exactly one neighbor in W_{k-1}. We may assume that $v_{k+1} \sim w_t$ and $w_k \sim w_s$ with $s \leq t$.

We first treat the case $s = t$, i.e. we assume that v_{k+1} and w_k have a common neighbor $w_s = w_t$. We have seen that w_s has a neighbor u_s outside $V(F) \cup V_{k+1} \cup W_k$, thus expanding V_s^+, w_s yields $\Delta(2s+4, s+2, 0)$, see Figure 4.10 (a). We may therefore assume $s < t$.

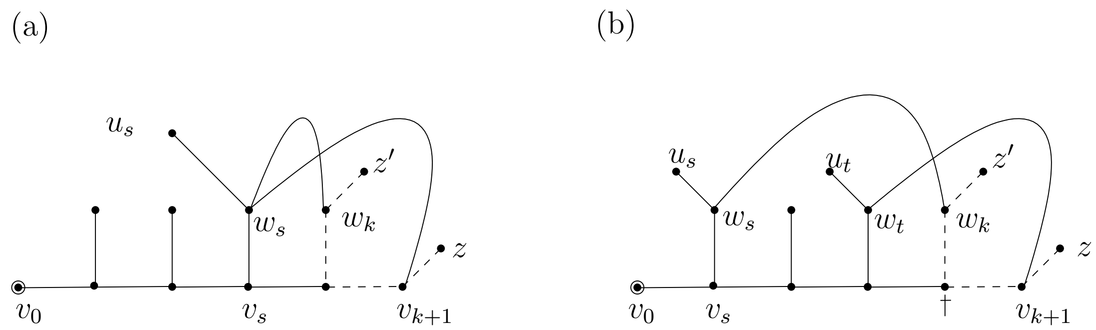

Figure 4.10.

As mentioned above w_s has a neighbor $u_s \notin V(F) \cup V_{k+1} \cup W_k$ and w_t has a neighbor $u_t \notin V(F) \cup V_{k+1} \cup W_k$. Then, u_s is the only neighbor of w_s which is not in $V(F) \cup V_{k+1} \cup W_k$ since we have already considered the case that w_s has at least three neighbors not in $V(F) \cup V_{s+1} \cup W_s$. Similarly u_t is the only neighbor of w_t which is not in $V(F) \cup V_{k+1} \cup W_k$. Assume first, that $u_s \neq u_t$. Then, we expand V_{k-1}^+, w_s, w_t and obtain $\Delta(2k+3, k+1, 1) = 0$, since v_k is a dead leaf, see Figure 4.10 (b).

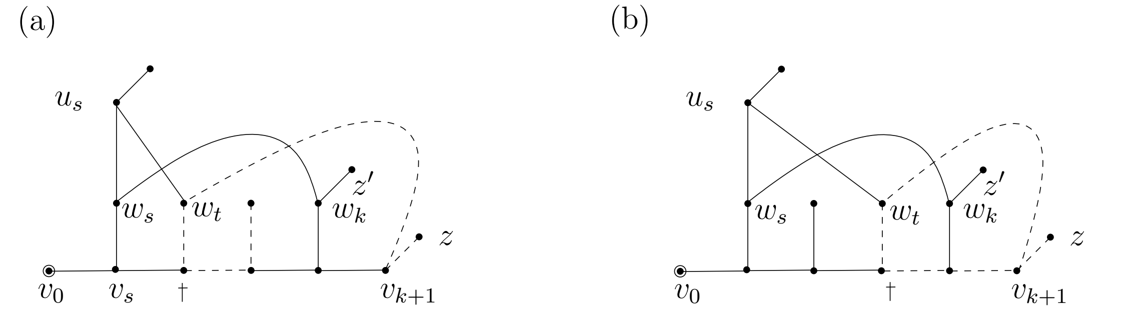

Figure 4.11.

Thus, we may assume that $u_s = u_t$. If $d(u_s) = 2$, then expanding V_k, w_s yields $\Delta(2k+1, k, 1) = 0$, since w_t is a dead leaf. If u_s has a neighbor that is not in $V(F) \cup V_{k+1} \cup W_k \cup \{z'\}$, then we expand

$$V_{t-1}^+, w_s, u_s, w_k, v_k, v_{k-1}, v_{k-2} \dots v_{t+2}$$

and obtain $\Delta(2k+3, k+1, 1)$, since v_t becomes a dead leaf, see Figure 4.11 (a). If $t = k-1$ we expand V_{t-1}^+, w_s, u_s, w_k to obtain the same result, see Figure 4.11 (b).

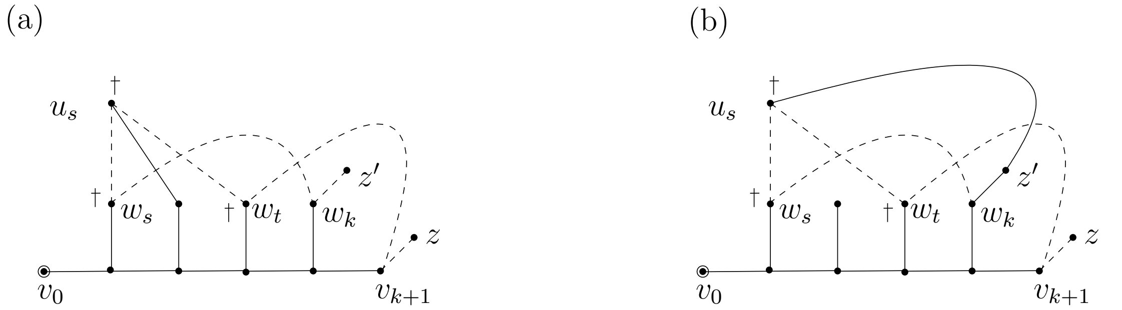

Figure 4.12.

Finally, we may assume that u_s has a third neighbor in $W_k \cup \{z'\}$. If this neighbor is some w_j then expanding V_k^+, w_j yields $\Delta(2k+2, k, 3) = 1$ since u_s, w_s, w_t all become dead leaves, see Figure 4.12 (a). If $u_s \sim z'$, we expand V_k^+, w_k, z' and obtain $\Delta(2k+3, k, 3) = 0$ since u_s, w_s, w_t all become dead leaves, see Figure 4.12 (b). $\square$

4.3 Leafy Trees in Graphs without Cubic Diamonds

In this section we introduce another ingredient, that is essential for the proof of Theorem 4.1, the main result of this chapter. This ingredient is the technique of graph reductions which help us to cope with the case distinctions in the constructions of F'. Theorem 4.9 can be obtained without much effort as a corollary of Theorem 4.1, as we explain in Section 4.4.4. In this section we give an independent, relatively short and simple proof of Theorem 4.9 to demonstrate the use of graph reductions.

cubic diamond **Theorem 4.9.** *Let G be a connected graph without cubic diamonds on at least two vertices. Then, G has a spanning tree T with*

$$\ell(T) \geq 2n_{\geq 3}(G)/7 + \begin{cases} 12/7 & \text{if } G \text{ is cubic}, \\ 2 & \text{otherwise.} \end{cases}$$

Bonsma's main result about spanning trees with many leaves in [15] is the following. Let D be the number of cubic diamonds, that a graph G of minimum degree 3 contains. Then G has a spanning tree T with

$$\ell(T) \geq (2n - D + 12)/7.$$

As for his proof of Theorem 4.8, Bonsma's uses an approach via connected dominating sets. It is easy to generalize Theorem 4.9 such that it also incorporates the variable D that counts the number of cubic diamonds. We simply substitute each of the $D \geq 1$ cubic diamonds of G by a vertex of degree 2 which yields a graph G' with $n_{\geq 3}(G') = n_{\geq 3}(G) - 4D$. Now we take advantage of the fact that our proof technique does not need assumptions about the minimum degree.

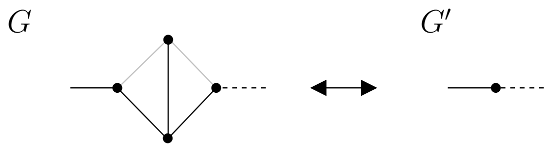

Figure 4.13. Substituting cubic diamonds by degree 2 vertices.

We apply Theorem 4.9 to conclude that G' has a spanning tree T' with $\ell(T') \geq 2n_{\geq 3}(G')/7 + 2$. The tree T' can be extended to a spanning tree of G with $\ell(T) = \ell(T') + D$. This is illustrated in Figure 4.13 where the dashed edge is in T if and only if it is in T'. Thus, we have that

$$\begin{aligned} \ell(T) &\geq 2n_{\geq 3}(G')/7 + 2 + D \\ &= (2n_{\geq 3}(G) - 8D)/7 + 2 + D = (2n_{\geq 3}(G) - D)/7 + 2. \end{aligned}$$

Theorem 4.9 is tight for the cube Q_3. That the linear factor cannot be improved is shown by the graphs obtained from a cycle by replacing every vertex with an N_2 2-necklace. The bound $n_{\neq 2}/4 + 3/2$ from Theorem 4.7 is tight for the graphs obtained from ternary trees by replacing every inner vertex by a cubic triangle, see Figure 4.6 (b). These graphs do not contain cubic diamonds. Thus it is not possible to strengthen Theorem 4.9 by replacing $n_{\geq 3}(G)$ by $n_{\neq 2}(G)$.

Besides demonstrating the use of graph reductions, the proof of Theorem 4.9 that we present has the advantage of being considerably shorter than that in [15]. As in the last section, we introduce some notation and conventions which are also needed for the proof of Theorem 4.1, and we exemplify the use of graph reductions in the proof of Theorem 4.9.

Just like Theorem 4.9, the results in the rest of this chapter are bounds in terms of $n_I(G)$ for $I = \{1, 2\}$. It is therefore convenient to introduce a name for vertices which have degree at most 2 and we call them *goobers*. We adopt this notion *goober* from [54], although there it is defined in a slightly different way. This difference is irrelevant for the proof of Theorem 4.9, and we discuss this issue in more detail in Section 4.4.2, when we introduce Theorem 4.16.

Ignoring that the reductions may disconnect the graph, the main idea behind them is the following. A graph G is reduced to a graph G' with $n_{\geq 3}(G) - n_{\geq 3}(G') = k$, such that every spanning tree of G' can be turned into a spanning tree of G with at least $k/3$ additional leaves. This preserves the desired leaf ratio, see Lemma 4.10 for more details. The main advantage of the reductions is that we can exclude certain cases in the construction of the extension F' for some $F \subset G$.

The proof of Theorem 4.9 uses the three graph reductions that are shown in Figure 4.14. The numbers above the arrows indicate the decrease in $n_{\geq 3}$, and the numbers below the arrows indicate the number of leaves that can be gained in a spanning tree when reversing the reduction. How these leaves can be gained is described in the proof of Lemma 4.10. The white vertices are goobers, black vertices have degree 3. Dashed edges are present in the reduced graph if and only if they are present in the original graph.

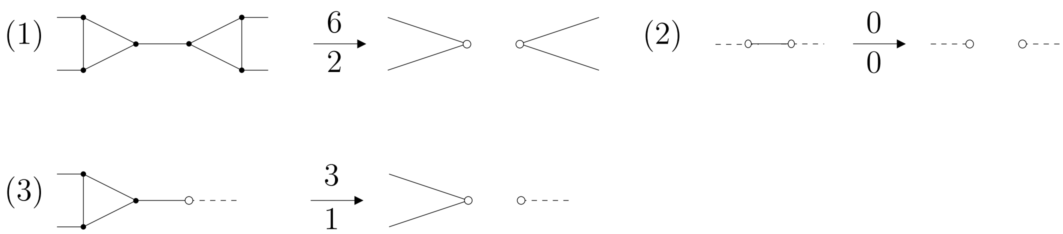

Figure 4.14. Reductions for the proof of Theorem 4.9.

Each of the reductions may be applied if it does not introduce a cubic diamond and no multiple edges incident to a vertex of degree at least 3. A reduction is *admissible* for a graph G if it can be applied without violating this condition. *admissible* *reducible/irreducible*

A graph G is *reducible* if one of the three reduction rules is admissible, and *irreducible*otherwise. We will prove Theorem 4.9 for irreducible graphs using the method introduced in Section 4.2 and then the following lemma is used to deduce the statement of Theorem 4.9.

maximal forest

A forest F of G is a *maximal forest* if G has no forest F' that is a strict supergraph of F. Hence F is a maximal forest of G if and only if it consists of a spanning tree for every component of G. Components with only one vertex will be called *trivial components* in the sequel.

trivial component

Lemma 4.10. *Let G' be the result of applying one of the Reductions (1)–(3) to a connected graph G and let k be the number of non-trivial components of G'. If G' has a maximal forest F with at least $n_{\geq 3}(G')/3 + 2k$ leaves, then G has a spanning tree T with at least $n_{\geq 3}(G)/3 + 2$ leaves.*

Proof. Clearly, G' has one or two connected components and $k \in \{1, 2\}$. A reduction that creates i new goobers, decreases $n_{\geq 3}(G)$ by $3i$, see Figure 4.14. Figure 4.15 shows for $k \geq 1$ how a spanning forest of G with k non-trivial components and i new leaves can be gained for Reductions (1)–(3). Dotted edges are in the forest on the right, if and only if they are in the forest on the left.

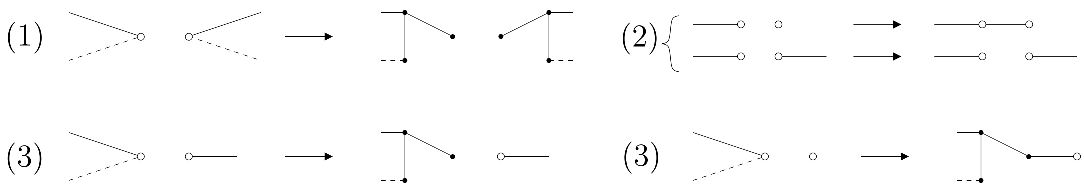

Figure 4.15. Reversing the reductions for the proof of Theorem 4.9.

For $k = 1$ we obtain the following by using the constructions from Figure 4.15.

$$\ell(T) \geq \ell(F) + 2 + i \geq n_{\geq 3}(G)/3 - i + i + 2 = n_{\geq 3}(G)/3 + 2.$$

If $k = 2$, then the two trees which form the maximal forest F must be connected at the cost of two leaves and we obtain

$$\ell(T) \geq \ell(F) + 4 + i - 2 \geq n_{\geq 3}(G)/3 + 2.$$

If $k = 0$, then the applied reduction was Reduction (2), and $G = K_2$. Thus, G has a spanning tree with two leaves which suffices. $\square$

We will show next, how Theorem 4.9 can be deduced from Lemma 4.10 in conjunction with Lemma 4.11. We then also give the proof of Lemma 4.11.

Lemma 4.11. *Let G be a connected, irreducible graph without cubic diamonds on at least two vertices. Then, G has a spanning tree T with*

$$\ell(T) \geq 2n_{\geq 3}(G)/7 + \begin{cases} 12/7 & \text{if } G \text{ is cubic} \\ 2 & \text{otherwise.} \end{cases}$$

Proof of Theorem 4.9. If G is irreducible, then the claim follows directly from Lemma 4.11, otherwise, one of the Reductions (1)–(3) is applicable. If G is reducible, then no component of the reduced graph is cubic. We may therefore assume by induction over the number of edges that the reduced graph has a maximal forest with $2n_{\geq 3}/7+2k$ leaves. Lemma 4.10 then implies that G has a spanning tree T with $\ell(T) \geq 2n_{\geq 3}(G)/7 + 2$. □

Proof of Lemma 4.11. The theorem uses $a = 7/2$, $c = 12/7$ for cubic graphs, $c = 2$ for non-cubic graphs, and $I = \{1, 2\}$.

We choose $a_1 = 3$ and $a_2 = 1/2$, thus the definition of the *leaf potential* that we use here is *leaf potential*

$$\mathcal{P}(F, G) = 3\ell(F) + \ell_d(F)/2 - n_{\neq 2,G}(F) - 7cc(F),$$

and we define $\Delta(x, y, z) = 3y + z/2 - x$.

If G is cubic, then the initial subgraph F with $\mathcal{P}(F, G) \geq -1$ is easily obtained. Let v a vertex of G. Since G has no cubic diamonds, v has a neighbor w, such that the edge $\{v, w\}$ is in no triangle. Thus, expanding v, w yields $\Delta(6, 4, 0) = 6$. In the case that G is not cubic Claim 1 shows that there is an initial graph which satisfies $\mathcal{P}(F, G) \geq 0$.

If F has a non-leaf vertex v that is adjacent to a vertex from $\overline{V(F)}$, then expanding v implies $\Delta(1, 1, 0) = 2$. Thus, we may assume that all vertices on the boundary of F are leaves of F. Figure 4.16 shows four simple operations, that extend a non-spanning subgraph F. The following conventions apply for all figures in this proof. White vertices have degree at most 2, vertices marked with a cross are dead leaves. For all other vertices the shown degrees are to be understood as lower bounds, unless otherwise stated. The encircled vertices already belong to F, the solid edges show how the tree is extended, while dashed edges are not in the tree.

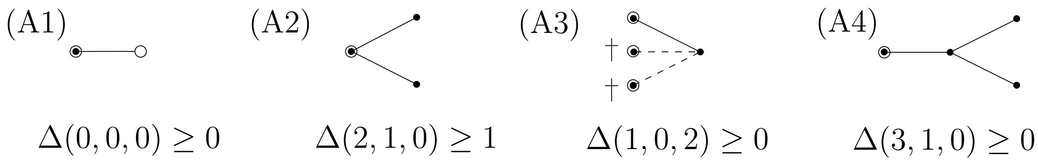

Figure 4.16. Extensions for the proof of Theorem 4.9.

We assume from now on, that none of the operations (A1) – (A4) is applicable. This implies, that every $v \in F$ has at most one neighbor in F^C, i.e. $d_{F^C}(v) \leq 1$ and that every $v \in \overline{V(F)}$, $v \sim F$ has two neighbors in F and one neighbor in $\overline{V(F)}$. Thus, all $v \in \overline{V(F)}$, $v \sim F$ have $d(v) = 3$. We split the rest of the proof into two claims.

Claim 1. Let $F \subset G$ be a, possibly empty, subgraph of G, such that not all vertices in $\overline{V(F)}$ have degree 3. Then, F is extendible.

First, suppose that there is a vertex $v \in \overline{V(F)}$ with $d(v) \geq 4$. Then expanding v yields at least $\Delta(5,4,0) = 7$. Thus, we may assume that $d(v) \leq 3$ for all $v \in F^C$.

Thus, there is v with $d(v) \leq 2$. Let w be a neighbor of v which must have $d(w) = 3$ since otherwise Reduction (2) would be applicable. If $w \sim x \in F$ we obtain $\Delta(1,0,2)$ if $d(v) = 1$. If $d(v) = 2$, then we may proceed with the other neighbor w' of v, unless $w' \sim F$. If $w, w' \sim F$ we obtain $\Delta(2,0,4) = 0$ by expanding x, w, v.

Thus, we may assume that $w \not\sim F$, and we denote the two neighbors of w by x and y. If x or y is a goober, then we obtain $\Delta(2,3,0) = 7$ by expanding w. Since $d_G(v) \leq 2$, we may assume that x has a neighbor z other than v, w, y. If $x \sim y$, then z is adjacent to y as well and not a goober, since Reduction (3) is not admissible. Furthermore, z must have degree at least 4, since there are no cubic diamonds in G. Since we assumed that the vertices in $\overline{V(F)}$ have degree at most 3, we have that $z \in F$. But z has two neighbors $x, y \in \overline{V(F)}$ which contradicts the assumption that (A2) is not applicable. Thus, $x \not\sim y$ and therefore x has $|N(x) \setminus N[w]| = 2$. If $x \not\sim F$, then expanding w, x yields $\Delta(5,4,0) = 7$. If $x \sim z \in F$, then expanding z, x, w yields $\Delta(3,1,1) = 0.5$. $\triangle$

Claim 2. Let $F \subset G$ be a non-empty subgraph of G, such that all vertices of $\overline{V(F)}$ have degree 3. Then, F is extendible.

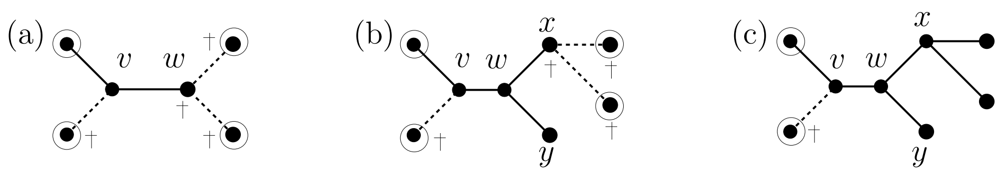

Figure 4.17.

Let $v \in \overline{V(F)}$ have neighbors $s, t \in F$ and $w \in \overline{V(F)}$. If $w \sim F$, then expanding s, v yields $\Delta(2,0,4) = 0$, see Figure 4.17 (a). Thus, w has two neighbors $x, y \in \overline{V(F)}$. If $x \sim F$, then expanding s, v, w yields $\Delta(4,1,4) = 1$, see Figure 4.17 (b). If $|N(x) \setminus \{w, y\}| = 2$, then expanding s, v, w, x yields $\Delta(6,2,1) \geq 0.5$, see Figure 4.17 (c).

Thus, $x \sim y$ and x has a neighbor $z \neq w, y$. If $z \sim F$ then expanding s, v, w, x yields $\Delta(5, 1, 4) = 0$, see Figure 4.18 (a). If $y \sim z$, then, since G has no cubic diamonds $d(z) \geq 4$ which contradicts that $x \not\sim F$. Thus, z has two neighbors $a, b \notin F \cup \{x\}$. If $a \sim y$, then expanding s, v, w, x, z yields $\Delta(7, 2, 2) = 0$, see Figure 4.18 (b).

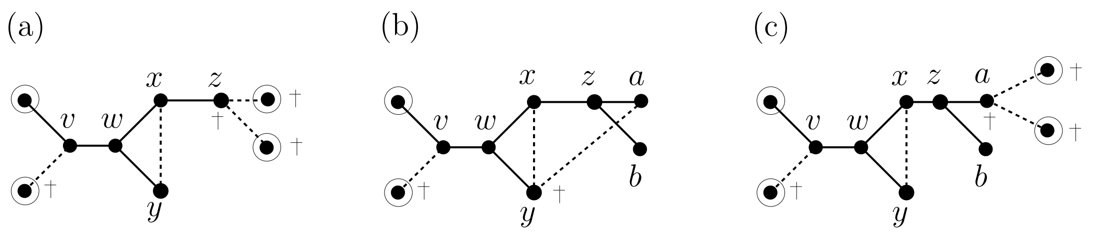

Figure 4.18.

If $a \sim F$, then we obtain $\Delta(7, 2, 4) = 1$ by expanding s, v, w, x, z, see Figure 4.18 (c), and if a has two neighbors which are not in $\{b, y, z\}$, then expanding s, v, w, x, z, a yields $\Delta(9, 3, 1) = 0.5$, see Figure 4.19 (a). Thus, $a \sim b$ and since Reduction (1) is not applicable, see Figure 4.19 (b), a, b must have a common neighbor $c \neq z$. Since G has no cubic diamonds, $d(c) \neq 3$ and therefore $c \in F$. But this is a contradiction to the assumption $a \not\sim F$, and therefore the proof of the claim is complete. $\triangle$

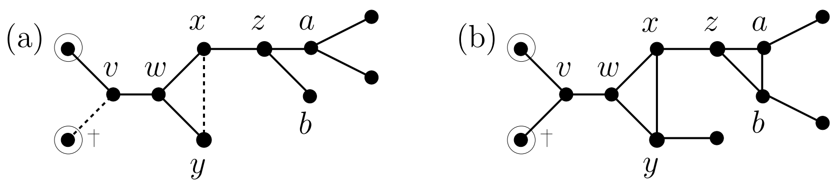

Figure 4.19.

Claims 1 and 2 in conjunction with the construction of the initial tree for cubic graphs yield a spanning subgraph F with sufficient leaf potential. The lemma now follows as described in Section 4.2 $\square$

4.4 Leafy Trees in Graphs without Necklaces and Blossoms

This section is devoted to the proof of the main theorem of this chapter which we repeat here for convenience. We will then sketch the proof and give an overview of the different ingredients that it uses.

Theorem 4.1. *Let G be a simple, connected graph on at least two vertices which contains neither 2-necklaces nor 2-blossoms. Then, G has a spanning tree T with*

$$\ell(T) \geq n_{\geq 3}(G)/3 + \begin{cases} 4/3 & \text{if } G = Q_3, \\ 5/3 & \text{if } G = G_7 \text{ or } G \neq Q_3 \text{ is cubic}, \\ 2 & \text{otherwise.} \end{cases}$$

We introduce a number of reduction rules in Section 4.4.1. These reduction rules are applied to the graph G until an irreducible graph G' is obtained. The rules maintain an invariant which guarantees that if Theorem 4.1 holds for every component of G', it also holds for G. When we show how to construct an extension F' for a subgraph $F \subset G'$, we can therefore restrict our attention to irreducible graphs.

In Section 4.4.2 we formulate the lemmas that take care of the construction of F'. We argue that the central part of the proof of Theorem 4.16 in [54] in fact can be used in our setting as well when the graph F^C has maximum degree 3. We then formulate a lemma that enables us to obtain a forest F that covers all vertices of degree at least 4 and also has enough leaves. This lemma is the core of our proof and also its most technical part. In Section 4.4.3 we combine the tools from Sections 4.4.1 and 4.4.2 to prove Theorem 4.1.

In Section 4.4.4 we discuss two theorems which follow from Theorem 4.1. One of them is a generalization of Theorem 4.1 that does not require a graph without 2-necklaces and 2-blossoms. Instead it incorporates the number of 2-necklaces and 2-blossoms that a graph has into the bound. The second theorem that we prove in Section 4.4.4 gives a bound for the number of leaves that can be obtained in graphs without 2-necklaces when 2-blossom subgraphs are allowed. The bound in this theorem is tight for the flower trees introduced in Section 4.1.

In Section 4.4.5 we prove two lemmas that are needed for the proof of Theorem 4.1 in Section 4.4.3. We delay these proofs until the end of Section 4.4 since they are rather long and technical.

Before we start the preparations for the proof of Theorem 4.1 we now discuss the tightness of the bounds that it states. The bound $n/3 + 4/3$ for cubic graphs is shown to be tight in [54] and it is mentioned that there exists only one graph for which the additive term $4/3$ cannot be increased. This graph is the 3-dimensional cube Q_3, see Figure 4.20 (a). Furthermore two examples of cubic graphs are given for which only an additive term of $5/3$ can be achieved, Figure 4.20 (b) shows one of them, the other one contains a triangle. From this latter example an infinite family of graphs can be constructed for which only $5/3$ can be achieved with the help of Reduction (7) from Figure 4.21. Furthermore, in [54] infinitely many examples are given with no more than $n/3 + 2$ leaves.

The bounds from Theorem 4.1 for non-cubic graphs are best possible as well. Infinitely many graphs with arbitrarily many vertices of higher and lower degree can be constructed that do not admit more than $n_{\geq 3}/3+2$ leaves. Figure 4.20 (c) shows such an example with many degree 2 and degree 4 vertices. This example is closely related to one of the examples from [54].

The reason that the additive term in Theorem 4.1 cannot be increased to 2 for all graphs with maximum degree at least 4 is again only one example. Figure 4.20 (d) shows a graph on $n = 7$ vertices that only admits $4 = n_{\geq 3}/3 + 5/3$ leaves. This graph will be called G_7 in the remainder. This graph is in fact a blossom plus two edges; deleting any edge between two degree 4 vertices yields a 2-blossom.

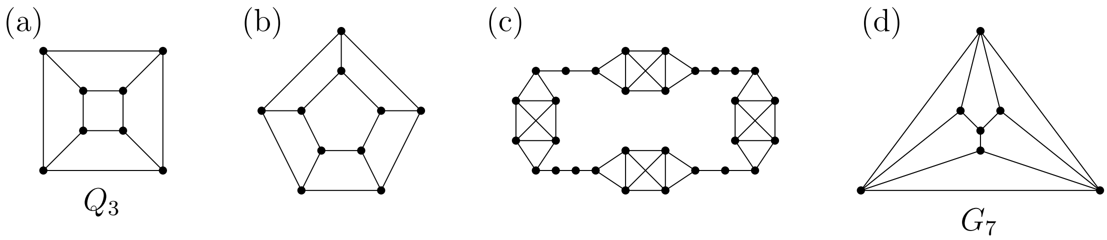

Figure 4.20. Four graphs for which Theorem 4.1 is tight.

4.4.1 Reducible Structures

We have used three graph reductions for the proof of Theorem 4.9 and in this section we introduce further reductions that we need for the proof of Theorem 4.1.

The proof of Theorem 4.1 follows the method that we introduced and exemplified in Section 4.2. That is, it relies on locally extending a forest until it becomes spanning while guaranteeing a certain number of leaves for every intermediate forest. The reductions help to delay the treatment of some substructures which cannot be readily handled during the extension process and they also simplify the case study in the proof of Lemma 4.19.

We first repeat the seven reduction rules defined in [54], and then introduce five new rules which are designed to handle structures containing degree 4 vertices. While the first seven rules are defined in [54] for graphs with maximum degree 3 we define them for arbitrary graphs, but the vertices on which they act must have the same degrees as in the original definition. Reductions (1) – (3) have already been used in the proof of Theorem 4.9 in this way.

The seven reduction rules from [54] consist of graph operations on certain structures, and conditions on when they may be applied. Figure 4.21 shows the operations. The black vertices all have degree 3, and goobers are shown as white vertices. Dashed edges are present in the resulting graph if and only if they exist in the original graph. The numbers above the arrows indicate the decrease in $n_{\geq 3}$, and the numbers below the arrows indicate the number of leaves that can be gained in a spanning tree when reversing the reduction. The following restrictions are imposed on the application of these rules, see Section 3 of [54].

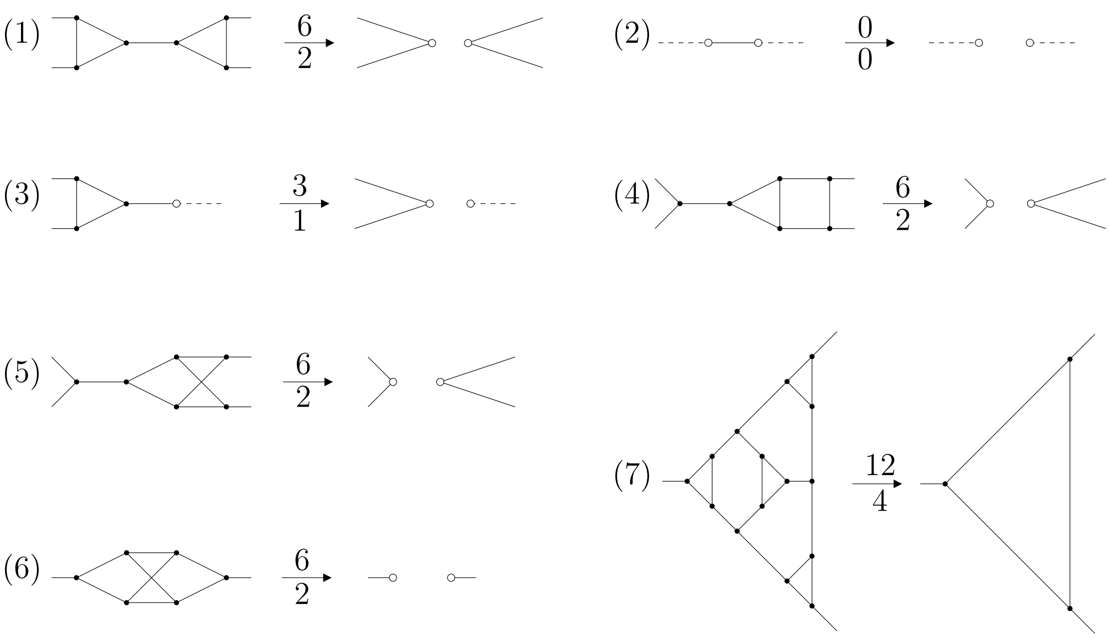

Figure 4.21. The seven low-degree reduction rules.

- Reductions (1), (3), (4), and (5) may not be applied if the two outgoing edges from the left side, or the two outgoing edges from the right side, share a non-goober end vertex. (An outgoing edge from the left and an outgoing edge from the right may share a non-goober end vertex.)
- Reduction (7) may not be applied if any pair of outgoing edges shares a non-goober end vertex.

In other words, a rule may not be applied if it would introduce multi-edges incident with non-goobers, or if it would introduce a diamond. These seven reduction rules will be called the *low-degree reduction rules*.

low-degree reduction rules

We define an invariant that exhibits the properties which should be maintained while applying graph reductions. These properties ensure that induction can be applied in the proof of Theorem 4.1. We denote the graph consisting of two vertices connected by two parallel edges as K_2+e. A graph H is said to *satisfy the invariant* if it fulfills the following conditions.

$K_2 + e$

satisfy the invariant

(i) The graph H is connected, or every component of H contains a goober, and

(ii) every component of H is either simple or it is a $K_2 + e$, and

(iii) H contains neither 2-necklaces nor 2-blossoms.

Lemma 4.12. *Let G' be obtained from G by the application of a low-degree reduction rule. If G satisfies the invariant then so does G'. Furthermore, if G contains a goober, then so does G'.*

Proof. Property (i) obviously holds for all seven low-degree reduction rules by definition. Note that the Rules (1)–(6) only introduce goobers as new vertices and the only new edges are incident to these goobers. Furthermore all other vertex degrees remain unchanged. In conjunction with the applicability conditions, this implies (ii) and (iii) for these six rules. We can also deduce that $\delta(G') \leq 2$ whenever $\delta(G) \leq 2$. It remains to consider Rule (7).

Rule (7) obviously cannot create parallel edges, that is (ii) holds. Furthermore, Rule (7) cannot introduce a 2-blossom since a 2-blossom cannot share a vertex with a triangle induced by three vertices of degree 3. This is not true for 2-necklaces, but if Rule (7) introduces a 2-necklace, two of the outgoing edges share an end vertex, contradicting the condition for applying Rule (7). Finally, Rule (7) does not remove goobers by definition, so the proof of the lemma is complete. ☐

We now introduce five new reduction rules which we call the *high-degree reduction rules*. Each rule again consists of a graph operation and conditions on the applicability.

high-degree reduction rules

Figure 4.22 shows the graph operations for the five rules. The encircled vertices are the terminals and may have further incidences, unlike the other vertices. None of the vertices in the figures may coincide, but there are no restrictions on outgoing edges sharing end vertices, unless this yields parallel edges incident to non-goobers. The numbers above the arrows indicate the decrease in $n_{\geq 3}$, and the numbers below the arrows indicate the number of leaves that can be gained in a spanning tree when reversing the reduction. Since (R4) must disconnect a component, this notion is not relevant for (R4); this rule will be treated separately in the sequel.

Observe that in particular, (R5) seems counterproductive when the goal is to find spanning trees with many leaves, but it is useful to keep the case analysis in the proof of Lemma 4.19 simple.

The following restrictions are imposed on the applicability of the operations from Figure 4.22 to a graph G. First, none of the reduction rules may be applied if it introduces a new 2-necklace or 2-blossom. In addition, the following rule-specific restrictions are imposed. A *bridge* is an edge whose deletion increases the number of components.

bridge

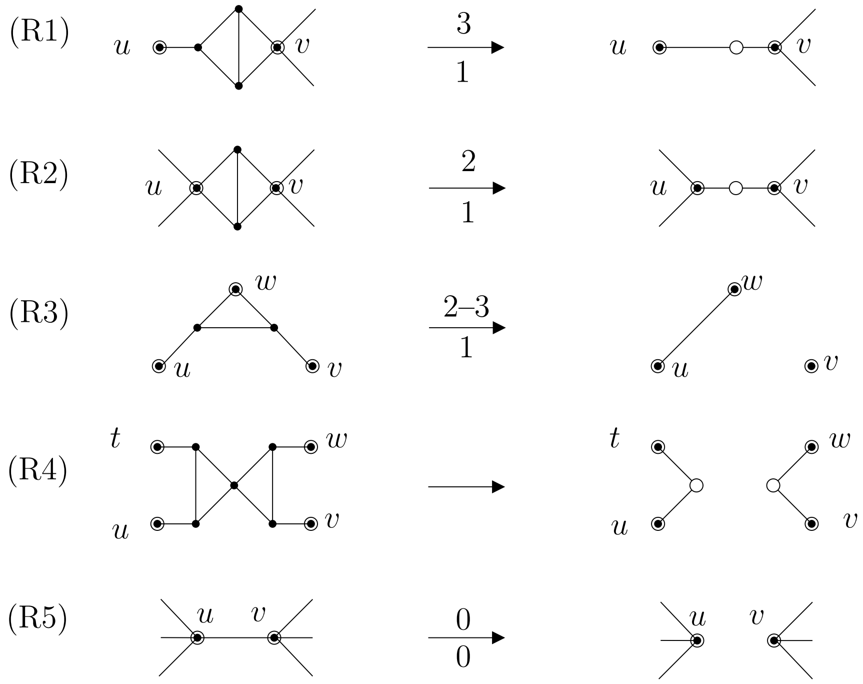

Figure 4.22. The five high-degree reduction rules.

(R1) $d_G(v) \geq 4$.

(R2) $d_G(u) \geq 4$ and $d_G(v) \geq 4$.

(R3) $cc(G') = cc(G)$, the edge uw is not in G, and in addition $d_{G'}(v) \geq 3$, or $d_{G'}(w) \geq 3$, or both.

(R4) $cc(G') > cc(G)$, that is G' is *not* connected.

(R5) $d_G(u) \geq 4$, $d_G(v) \geq 4$, uv may not be a bridge, and $G - uv$ is not cubic.

We generalize the definition from Section 4.3 and call each of the twelve reduction rules *admissible* if it can be applied without violating one of the imposed conditions. In particular the condition that no 2-necklaces or 2-blossoms are introduced will be important later and it also implies that the following lemma holds.

admissible

Lemma 4.13. *Let G' be obtained from G by the application of a high-degree reduction rule. If G satisfies the invariant then so does G'. Furthermore, if G contains a goober, then so does G'.*

We extend the definition of reducibility from Section 4.3. A graph G is *reducible* if one of the low-degree or high-degree reduction rules can be applied, and *irreducible* otherwise. Griggs et al. [54] call a graph irreducible if none of the

reducible/irreducible

low-degree reduction rules can be applied. Clearly, a graph that is irreducible according to our definition is also irreducible according to their definition. We use this in Section 4.4.2, when Lemma 4.18 is introduced.

Note that irreducible graphs satisfying the invariant are simple because of Rule (2). The following property of irreducible graphs substantially simplifies subsequent proofs. Here G_7 denotes the graph from Figure 4.20 (d). We delay the proof of Lemma 4.14 to Section 4.4.5, since it is rather long and technical.

Lemma 4.14 (Edge Deletion)**.** *Let G be an irreducible graph not equal to G_7 with adjacent vertices u and v. If $d(u) = d(v) = 4$, then uv is a bridge, or $G - uv$ is cubic, or one of u, v becomes an inner vertex of a cubic diamond upon deletion of the edge uv.*

We now show that we can reverse all the reduction rules while maintaining spanning trees with a sufficient number of leaves for every component. For the low-degree reduction rules this lemma was implicitly proved in [54] and we have seen how the reconstructions work for Reductions (1)–(3) in the proof of Lemma 4.10. We refer to [54] for the detailed tree reconstructions, but we do repeat the main idea behind the proof here.

Lemma 4.15 (Reconstruction Lemma)**.** *Let G' be the result of applying a reduction rule to a connected graph G, and let k be the number of non-trivial components of G', and $\alpha \geq 0 \in \mathbb{R}$. If G' has a maximal forest with at least $n_{\geq 3}(G')/3 + 2k - \alpha$ leaves, then G has a spanning tree with at least $n_{\geq 3}(G)/3 + 2 - \alpha$ leaves.*

Proof. We prove the statement only for $\alpha = 0$, the reasoning is the same for other values of α.

Case 1. The applied rule was a low-degree reduction rule. Note that $cc(G')$ is either 1 or 2. If G' is connected, that is $cc(G') = 1$, then its maximal forest is a spanning tree and can be turned into a spanning tree of G with $(n_{\geq 3}(G) - n_{\geq 3}(G'))/3$ more leaves. To prove this, it is shown in Section 3 of [54] for every rule how to adapt the tree of G' for G (*tree reconstructions*). So now we assume that $cc(G') = 2$. If $k = 2$ then applying the same tree reconstructions yields a spanning forest of G consisting of two trees, with again $(n_{\geq 3}(G) - n_{\geq 3}(G'))/3$ more leaves in total. These two trees of G can be connected to one spanning tree T by adding one edge which destroys at most two leaves, and we obtain

$$\ell(T) \geq n_{\geq 3}(G')/3 + 2k + (n_{\geq 3}(G) - n_{\geq 3}(G'))/3 - 2 = n_{\geq 3}(G)/3 + 2.$$

If exactly one of the two components is trivial ($k = 1$) then the applied rule must be Rule (2) or (3). In this case, it can be checked that after the tree reconstruction for the non-trivial component, the trivial component can be attached to the tree without decreasing the number of leaves. One leaf is lost but the isolated vertex becomes a leaf, and this yields

$$\ell(T) \geq n_{\geq 3}(G')/3 + 2 + (n_{\geq 3}(G) - n_{\geq 3}(G'))/3 = n_{\geq 3}(G)/3 + 2.$$

If both components of G' are trivial ($k = 0$), then Rule (2) was applied, and $G = K_2$, for which the statement holds. This proves the lemma when a low-degree reduction rule is applied.

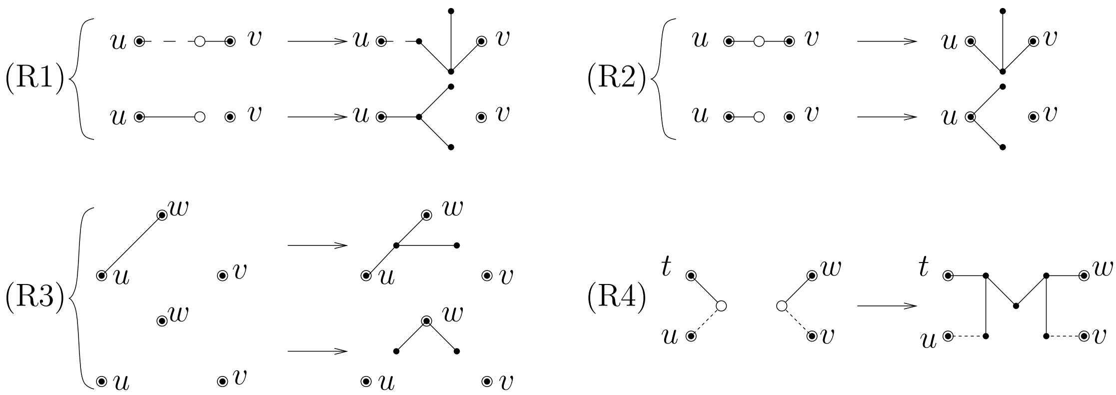

Figure 4.23. Tree constructions for reversing the high-degree reduction rules.

Case 2. The applied rule was a high-degree reduction rule. Note that Rules (R1), (R2), (R3), and (R5) do not increase the number of components, that is $k = 1 = cc(G')$. So for (R5) we do not have to change the spanning tree of G'. For (R1), (R2), and (R3), Figure 4.23 shows how to gain at least one additional leaf in every case. This suffices since each of these rules decreases $n_{\geq 3}$ by at most 3. Here it is essential that (R3) is admissible only if it creates at most one goober. Dashed edges in the figure are present on the right if and only if they are present on the left. Symmetric cases are omitted in the figure. Note that none of the terminals of the operations can lose leaf status, except w in the second reconstruction for (R3). This is compensated by gaining two new leaves here. So in every case enough leaves are gained to maintain the ratio.

Recall that (R4) is only admissible if it disconnects G into two components that are non-trivial, so $k = 2 = cc(G')$. Figure 4.23 shows how to construct a spanning tree for G from the two spanning trees for the components, without decreasing the total number of leaves. Hence the number of leaves of the resulting tree is at least

$$\ell(T) \geq n_{\geq 3}(G')/3 + 2k = n_{\geq 3}(G)/3 - 5/3 + 4 > n_{\geq 3}(G)/3 + 2.$$

This proves the lemma for all reduction rules.

□

4.4.2 Extension Lemmas

Recall that the proof method introduced in Section 4.2 relies on the construction of a graph F' for every non-spanning subgraph F of G, with $F \subset F' \subseteq G$ and $\mathcal{P}(F', G) \geq \mathcal{P}(F, G)$. In this section we summarize the three lemmas that take care of the construction of this extension F' in the proof of Theorem 4.1.

As announced in Section 4.3, we now explain how the definition of goobers in [54] differs from our definition. In [54] Griggs et al. define *goobers* as vertices of degree at most 2 *resulting from a (low-degree) reduction rule.* Considering the reduction rules, it can be seen that this extra condition adds no information (for instance about the possible neighborhoods of goobers). Indeed, no such information is used in the proofs in [54], and thus goobers may simply be defined as we do, i.e. as vertices of degree at most 2. We can therefore restate Theorem 3 from [54] as follows.

goober

Theorem 4.16. *Every irreducible graph G of maximum degree exactly 3 and without cubic diamonds has a spanning tree T with*

$$\ell(T) \geq n_{\geq 3}(G)/3 + \begin{cases} 4/3 & \text{if } G = Q_3, \\ 5/3 & \text{if } G \text{ is cubic}, \\ 2 & \text{otherwise.} \end{cases}$$

We give a short overview of the proof of this statement, as it appears in [54]. The proof method that we have introduced in Section 4.2 was mainly developed in [54] and Theorem 4.16 is proved in [54] with this method. The choice of the parameters is $a = 3$, $I = \{1, 2\}$ and $c \in \{4/3, 5/3, 2\}$, depending on the graph under consideration. In [54] the parameters a_1 and a_2 are chosen as $a_1 = 5/2$ and $a_2 = 1/2$ and we adopt this setting for our proof of Theorem 4.1. Thus, the *leaf potential* is defined as

leaf potential

$$\mathcal{P}(F, G) = 2.5\ell(F) + 0.5\ell_d(F) - n_{\geq 3,G}(F) - 6cc(F),$$

in the rest of Section 4.4 and we also use the short hand

$$\Delta(x, y, z) = 2.5y + 0.5z - x.$$

The main building block of the proof of Theorem 4.16 in [54] is the case study in Section 4 of [54]. The case study proves the following statement that we express using our notation.

Lemma 4.17. *Let G be a graph with maximum degree 3, without diamonds, that is irreducible with respect to the low-degree reduction rules. Let F be a non-empty tree subgraph of G. Then there exists a tree F' with $F \subset F' \subseteq G$ and*

$$2.5(\ell(F') - \ell(F)) + 0.5(\ell_d(F') - \ell_d(F)) - (n_{\geq 3,G}(F') - n_{\geq 3,G}(F)) \geq 0.$$

We now argue that the case study in [54] in fact proves Lemma 4.18 which we need for the proof of Theorem 4.1.

Lemma 4.18 (Extension Lemma)**.** *Let G be a connected irreducible graph, and let $F \subset G$ such that F^C has maximum degree 3 and contains no 2-necklaces. Then F is extendible.*

The most important observation is that the case study that proves Lemma 4.17 does not take advantage of any information about the current tree F. Only information about what we defined as F^C is used. In particular, the fact that F is connected is never used in the proof, and neither are upper bounds on degrees of vertices already included in F. So the maximum degree 3 condition only has to be stated for F^C, and the condition that F is a tree may be removed. Furthermore, an irreducible graph with maximum degree 3 that contains no 2-necklaces does not contain any diamonds as subgraphs. So we may replace the 'without diamonds' condition by the 'no 2-necklace' condition. Our definition of irreducible implies irreducibility with respect to the low-degree reduction rules, so this change neither is a problem. Finally, the resulting graph F' has the same number of components as F, so the expression in Lemma 4.17 simply means that $\mathcal{P}(F', G) \geq \mathcal{P}(F, G)$. This yields Lemma 4.18.

While Lemma 4.18 takes care of the construction of F' in the case that F^C has maximum degree 3, Lemma 4.19 grows trees around vertices of degree at least 4 and yields a graph satisfying the assumptions of Lemma 4.18. Lemma 4.19 is the core of our proof of Theorem 4.1. We delay the proof of Lemma 4.19 to Section 4.4.5, since it is rather long and technical.

Lemma 4.19 (Start Lemma)**.** *Let G be an irreducible graph that is not G_7 and does not have an edge e such that $G - e$ is cubic. Let F be a (possibly empty) subgraph of G, such that F^C contains at least one vertex of degree at least 4, and contains neither 2-necklaces nor 2-blossoms. Then, F is extendible.*

The case that there is an edge e such that $G - e$ is cubic requires additional attention. This is needed in order to preserve the additive term 2 for non-cubic graphs other than G_7 in the induction step. The following lemma guarantees a sufficient leaf potential for the initial subgraph of such an almost cubic graph.

Lemma 4.20. *Let G be an irreducible graph, that contains neither 2-necklaces or 2-blossoms. If G has an edge e such that $G - e$ is cubic, then G has a subgraph F such that $\mathcal{P}(F, G) \geq -0.5$ and F^C has maximum degree 3.*

Proof. Let u, v be the degree 4 vertices incident to e. If u, v have at most one common neighbor, then expanding u, v yields $\Delta(7, 5, 0) = 5.5$. If u, v have three common neighbors, then expanding u yields $\Delta(5, 4, 1) = 5.5$, since v becomes a dead leaf. So we may assume that u, v have two common neighbors x, y.

If z, the fourth neighbor of u, has $|N(z) \setminus \{u, v, x, y\}| = 2$, then expanding u, z yields $\Delta(7, 5, 0) = 5.5$. Otherwise, we may assume that z is adjacent to x. Expanding u then gives $\Delta(5, 4, 1) = 5.5$, since x becomes a dead leaf. $\square$

4.4.3 Proof of the Main Theorem

This section is devoted to combining the tools introduced in Sections 4.4.1 and 4.4.2 in order to prove Theorem 4.1 which we repeat here for convenience.

Theorem 4.1. *Let G be a simple, connected graph on at least two vertices which contains neither 2-necklaces nor 2-blossoms. Then, G has a spanning tree T with*

$$\ell(T) \geq n_{\geq 3}(G)/3 + \begin{cases} 4/3 & \text{if } G = Q_3, \\ 5/3 & \text{if } G = G_7 \text{ or } G \neq Q_3 \text{ is cubic}, \\ 2 & \text{otherwise}. \end{cases}$$

Proof. If G has maximum degree exactly 3, then Theorem 4.1 follows immediately from Theorem 4.16. If G has maximum degree at most 2, G has a spanning tree with two leaves, since we assumed that G is not K_1. For all other graphs we prove the statement by induction over the number of edges. For our induction hypothesis we show that the above statement holds for every irreducible, connected graph which satisfies the invariant.

Induction base. The induction base is the case that G is irreducible. If $G = G_7$, then a spanning tree with $4 = n_{\geq 3}(G)/3 + 5/3$ leaves can be obtained. So we may now assume that G contains at least one vertex of degree at least 4, and is not equal to G_7.

If G has an edge e such that $G-e$ is cubic, then Lemma 4.20 yields a subgraph F with $\mathcal{P}(F, G) \geq -0.5$, such that F^C has maximum degree 3. Then, the Extension Lemma (Lemma 4.18) can be applied iteratively until a spanning subgraph F' is obtained with $\mathcal{P}(F', G) \geq -0.5$. In a spanning graph all leaves are dead, therefore we must have $\mathcal{P}(F', G) \geq 0$, since $\mathcal{P}_G(F', G)$ is integral.

In the remaining case, G fulfills the assumptions of Lemma 4.19 and we start the construction process with an empty subgraph F of G which has $\mathcal{P}(F, G) = 0$. The Start Lemma (Lemma 4.19) shows that, as long as there is at least one vertex of degree at least 4 not in F, we can extend F while maintaining $\mathcal{P}(F, G) \geq 0$. When all vertices of degree at least 4 are included in F, Lemma 4.18 can be applied iteratively until a spanning subgraph F' is obtained with $\mathcal{P}(F', G) \geq 0$.

We may assume that F' is a forest because cycles can be broken by edge deletions without decreasing the number of leaves.

Since all leaves of a spanning subgraph are dead we deduce

$$\begin{aligned} 0 &\leq \mathcal{P}(F', G) = 3\ell(F') - n_{\geq 3}(G) - 6cc(F') \\ \Rightarrow \quad \ell(F') &\geq n_{\geq 3}(G)/3 + 2cc(F'). \end{aligned}$$

We can now add $cc(F') - 1$ edges to F' to obtain a spanning tree, losing at most $2(cc(F') - 1)$ leaves, so the resulting tree has at least $n_{\geq 3}(G)/3 + 2$ leaves.

Induction step. Now, we assume that G is reducible. Then, some reduction rule is admissible, and the reduced graph G' again satisfies the invariant, by Lemmas 4.12 and 4.13. These lemmas also imply that if G contains a goober, then so does G'.

First suppose the reduction rule yields a disconnected graph G'. Then, every resulting component has a goober. So by induction, every non-trivial component C of G' has a spanning tree with at least $n_{\geq 3}(C)/3 + 2$ leaves. Thus, Lemma 4.15 (the Reconstruction Lemma) implies that G has a spanning tree with at least $n_{\geq 3}(G)/3 + 2$ leaves.

Now we suppose that G' is connected. If G' has a spanning tree with at least $n_{\geq 3}(G')/3 + 2$ leaves, then Lemma 4.15 implies that G has a spanning tree with at least $n_{\geq 3}(G)/3 + 2$ leaves. By induction, such a spanning tree for G' can be guaranteed whenever G' is not cubic and $G' \neq G_7$. So now we may assume that G' is cubic or $G' = G_7$. It follows that the applied reduction rule was (7), (R3) or (R5), since all other reduction rules introduce goobers. If (R3) was applied, then in addition it follows that $n_{\geq 3}(G) = n_{\geq 3}(G') + 2$. We consider three cases for G'.

If G' is cubic but not equal to Q_3, then by induction it has a spanning tree T' with $\ell(T') \geq n_{\geq 3}(G')/3 + 5/3$ leaves. When this inequality is not tight, then since both $\ell(T')$ and $n_{\geq 3}(G')$ are integral, we have $\ell(T') \geq n_{\geq 3}(G')/3 + 2$, so the statement for G follows. So now assume $\ell(T') = n_{\geq 3}(G')/3 + 5/3$. If G is cubic as well, then Lemma 4.15 yields that G has a spanning tree with at least $n_{\geq 3}(G)/3 + 5/3$ leaves, which is good enough. If G is not cubic, then the applied reduction rule must be (R3), since (7) cannot yield a cubic graph G' if G is not cubic, and (R5) by definition may not yield a cubic graph. In that case, G has a spanning tree T with

$$\begin{aligned} \ell(T) = \ell(T') + 1 &= n_{\geq 3}(G')/3 + 5/3 + 1 \\ &= n_{\geq 3}(G)/3 - 2/3 + 5/3 + 1 = n_{\geq 3}(G)/3 + 2. \end{aligned}$$

Now suppose $G' = Q_3$. The applied rule is not (R5) by definition, and not (7) since Q_3 contains no triangles. So G' is obtained by applying (R3) to G, and it follows without loss of generality that G is the graph shown in Figure 4.24. There are different possibilities for the other end vertex of the edge e that is drawn as a half edge here. But regardless of this choice, the solid edges show a spanning tree for G with $6 = n_{\geq 3}(G)/3 + 8/3$ leaves, which suffices.

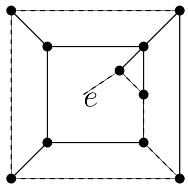

Figure 4.24. An (R3) operation yields a Q_3.

Finally, suppose $G' = G_7$. Then G' contains no triangle induced by three vertices of degree 3, so we may exclude Rule (7). If (R3) was applied, then Lemma 4.15 proves that a spanning tree T of G exists with $\ell(T) \geq n_{\geq 3}(G)/3 + 5/3$. Since in addition $n_{\geq 3}(G) = 9$, rounding gives $\ell(T) \geq \lceil n_{\geq 3}(G)/3+5/3 \rceil = 5 = n_{\geq 3}(G)/3+2$. If (R5) was applied, then Figure 4.25 shows all the possibilities for G, since there are only three ways to add an edge to G_7 such that a simple graph is obtained, when ignoring symmetric cases. For all of these cases, the solid edges in Figure 4.25 show a spanning tree with $5 = n_{\geq 3}(G)/3 + 8/3$ leaves, which suffices. □

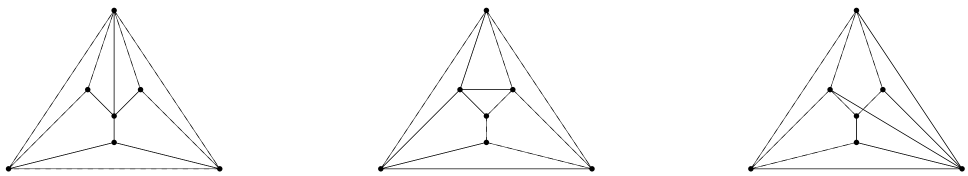

Figure 4.25. An (R5) operation yields a G_7.

4.4.4 Consequences of the Main Theorem

We now show how an even stronger version of Theorem 4.1 can be easily derived, from what we have proven so far. Let H be a 2-blossom or 2-necklace of G. If there exists a non-goober vertex in $V(G) \setminus V(H)$ that is adjacent to both terminals of H, then H is called a *leaf 2-blossom* or a *leaf 2-necklace* respectively. *leaf 2-blossom/ 2-necklace*

Theorem 4.21. *Let G be a simple, connected graph on at least two vertices and let x be the number of non-leaf 2-necklaces in G, and y the number of non-leaf 2-blossoms of G. Then G has a spanning tree T with*

$$\ell(T) \geq (n_{\geq 3}(G) - x - y)/3 + \begin{cases} 4/3 & \text{if } G = Q_3, \\ 5/3 & \text{if } G = G_7 \text{ or } G \neq Q_3 \text{ is cubic}, \\ 2 & \text{otherwise.} \end{cases}$$

Before we prove this theorem, we remark that Theorems 4.6 and 4.9 can be obtained as corollaries. In order to prove Theorem 4.6 note that $4x + 7y \leq n_{\geq 3}(G)$ for every graph G. Theorem 4.9 follows since for graphs without cubic diamonds we have $7x + 7y \leq n_{\geq 3}(G)$.

Proof of Theorem 4.21. If G contains no 2-necklaces or 2-blossoms, then the statement follows directly from Theorem 4.1. Let H be a 2-blossom of G. If $|V(G) \setminus V(H)| \leq 1$, then $n_{\geq 3}(G) = 7$, and a spanning tree T exists with $\ell(T) = 4 = (n_{\geq 3}(G) - 1)/3 + 2$, so the statement holds. Now let H be a 2-necklace of G consisting of k diamonds. If $|V(G) \setminus V(H)| \leq 1$, then a spanning tree T exists with $\ell(T) = k + 2 = (3k + 1 - 1)/3 + 2 = (n_{\geq 3}(G) - 1)/3 + 2$ which proves the statement.

In the remaining case, G contains at least one 2-necklace or 2-blossom H, and $|V(G) \setminus V(H)| \geq 2$. Let H be a 2-necklace or 2-blossom of G with terminals u and v. Since G is connected it follows that, the set $(N(u) \cup N(v)) \setminus V(H)$ contains either two vertices, or contains one non-goober vertex. So in the second case H is a leaf 2-necklace or leaf 2-blossom. Let x' denote the number of leaf 2-necklaces in G, and y' denote the number of leaf 2-blossoms in G.

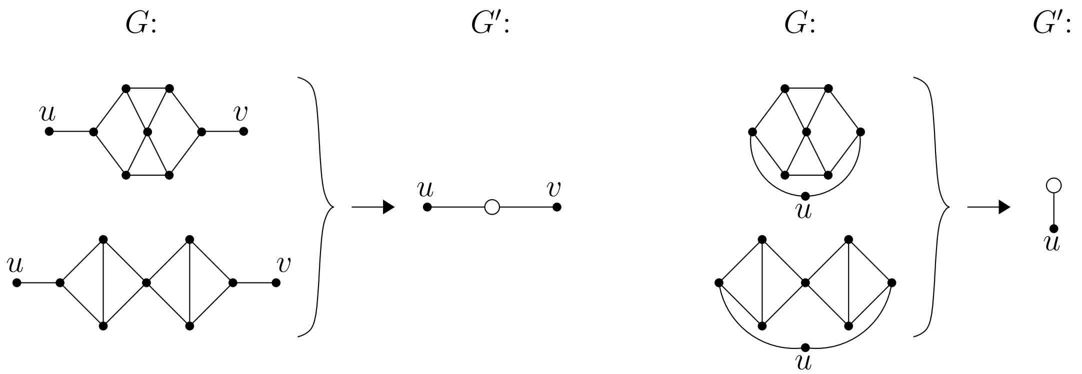

Figure 4.26. Reducing 2-blossoms and 2-necklaces.

We iteratively reduce every 2-necklace or 2-blossom H of G as follows, see Figure 4.26. We contract H into a single vertex w that becomes a goober in the reduced graph. If this yields two parallel edges incident with w, that is when H is a leaf 2-necklace or leaf 2-blossom, then we delete one of the parallel edges. This may turn one other vertex into a goober. Call the resulting graph G', and let d be the number of diamonds that are part of 2-necklaces of G. Note that every diamond necklace N_k that is a 2-necklace of G contributes k to the number of diamonds d. Then we have

$$n_{\geq 3}(G') \geq n_{\geq 3}(G) - 3d - x - 2x' - 7y - 8y'.$$

The graph G' is connected, simple, contains at least two vertices, and at least one goober. Therefore Theorem 4.1 implies that G' has a spanning tree T' with $\ell(T') \geq n_{\geq 3}(G')/3+2$. Figure 4.27 now illustrates how to extend T' iteratively for each of the four cases to obtain a spanning tree T of G.

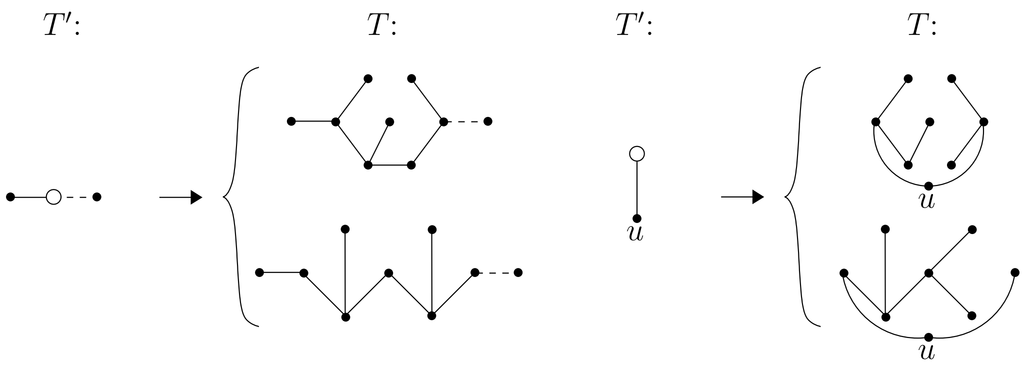

Figure 4.27. Reconstructing trees for 2-blossoms and 2-necklaces.

The dashed lines represent edges that are in T if and only if they are in T'. For every 2-necklace consisting of k diamonds we gain $k+1$ leaves if it is a leaf 2-necklace, and at least k leaves otherwise. In the case of a leaf 2-necklace, it is essential that the vertex u in Figure 4.27 cannot be a leaf of T'. Similarly, for every 2-blossom we gain three leaves if it is a leaf 2-blossom, and at least two leaves otherwise. So we have

$$\ell(T) \geq \ell(T') + d + x' + 3y' + 2y.$$

From these inequalities it follows that

$$\begin{array}{lllr} \ell(T) \geq & \ell(T') & +d+x'+3y'+2y & \geq \\ & n_{\geq 3}(G')/3+2 & +d+x'+3y'+2y & \geq \\ & (n_{\geq 3}(G)-3d-x-2x'-7y-8y')/3+2 & +d+x'+3y'+2y & \geq \\ & (n_{\geq 3}(G)-x-y)/3+2. & & \end{array}$$

□

The next theorem shows, that every graph without 2-necklaces has a spanning tree with $4n/13+c$ leaves. We have formulated this theorem in Section 4.1 and repeat it here for convenience. The bound $4n_{\geq 3}(G)/13+20/13$ is tight for the cube Q_3 and we have seen in Section 4.1 that the bound $4n_{\geq 3}(G)/13+24/13$ is tight for all *flower trees.* *flower tree*

Theorem 4.5. *Let G be a simple, connected graph on at least two vertices that contains no 2-necklaces. Then, G has a spanning tree T with*

$$\ell(T) \geq 4n_{\geq 3}(G)/13 + \begin{cases} 20/13 & \text{if } G \text{ is cubic,} \\ 24/13 & \text{otherwise.} \end{cases}$$

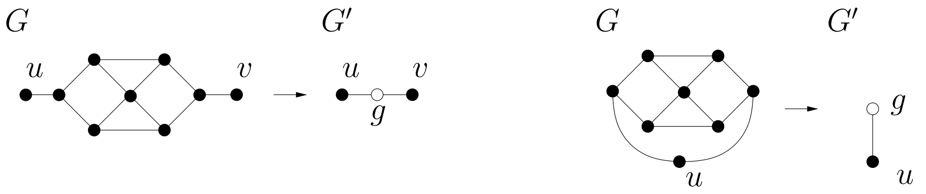

Figure 4.28. Blossom reductions.

Proof. The proof idea is to reduce all 2-blossoms of the graph using the reductions shown in Figure 4.28. Theorem 4.1 is then applied to the resulting graph G', and we obtain a spanning tree T' of G'. The reduction reversals then yield a spanning tree T of G with a sufficient number of leaves.

The bound follows from Theorem 4.1 for graphs without 2-blossoms. Both, Q_3 and G_7 have a spanning tree with four leaves, thus the claim holds for these two graphs. For every other cubic graph or non-cubic graph without 2-blossoms, the claim from Theorem 4.1 is obviously stronger than the claim from Theorem 4.5.

Note that graphs with maximum degree 3 have no 2-blossoms. We may thus assume that G has maximum degree at least 4 and at least one 2-blossom. We reduce all 2-blossoms of G by the reductions shown in Figure 4.28. The reduction shown on the left is for non-leaf 2-blossoms, while the reduction on the right is for leaf 2-blossoms. This yields a graph G' without 2-blossoms and 2-necklaces. We denote the number of non-leaf 2-blossoms by y and the number of leaf 2-blossoms by y'. Then, since the vertex u may become a goober when a leaf 2-blossom is reduced, we have

$$n_{\geq 3}(G') \geq n_{\geq 3}(G) - 7y - 8y'.$$

Since G' has at least one goober Theorem 4.1 shows that G' has a spanning tree T' with

$$3\ell(T') \geq n_{\geq 3}(G') + 6.$$

We build a spanning tree T of G from $T' = T_0$ by iteratively using the extensions shown in Figure 4.29 in an arbitrary order. Dashed edges are in T if and only if they are in T'.

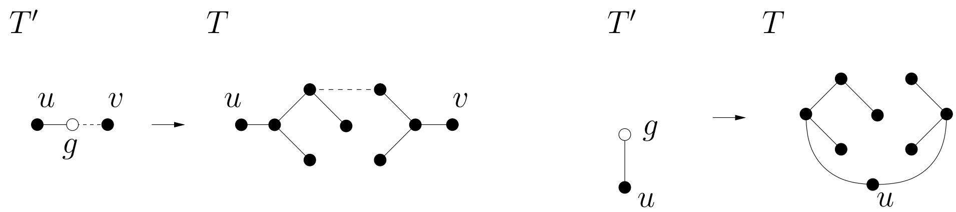

Figure 4.29. Blossom reconstructions.

We denote the intermediate trees by T_i and we have that $T_{y+y'} =: T$ is a spanning tree of G. Let y_1 be the number of extensions reversing the reductions of a non-leaf 2-blossom where u or v is a leaf of T', see Figure 4.29. Then, it holds that

$$\ell(T') \geq 2y_1 + y'.$$

First note, that in T' obviously all the degree 1 goobers created by the reduction of a leaf 2-blossom are leaves. Furthermore, if u or v is a leaf of T' then so is g, see Figure 4.29. The extension of T_i that reverses the reduction of a non-leaf 2-blossom makes u and v inner vertices of T_{i+1}. Thus, if u or v is adjacent to another goober created by the reductions of a non-leaf 2-blossom, then the reversal of that reduction is not counted by y_1. Hence, with every extension counted by y_1 we can associate two leaves of T'.

Let y_2 count the extensions reversing the reduction of a non-leaf 2-blossom that are not counted by y_1, i.e. $y_1 + y_2 = y$. Note that the extensions counted by y_1 create two additional leaves, while all other extensions create three additional leaves. This implies that

$$\begin{aligned}
\frac{13}{4}\ell(T) &= \frac{13}{4}\Big(\ell(T') + 2y_1 + 3(y_2 + y')\Big) \\
&\geq 3\ell(T') + \frac{1}{4}\ell(T') + \frac{13}{2}y_1 + \frac{39}{4}(y_2 + y') \\
&\geq n_{\geq 3}(G') + 6 + \frac{1}{4}(2y_1 + y') + \frac{13}{2}y_1 + \frac{39}{4}(y_2 + y') \\
&= n_{\geq 3}(G') + 6 + 7y_1 + \frac{39}{4}y_2 + 10y' \\
&\geq n_{\geq 3}(G) + 6.
\end{aligned}$$

We conclude that $\ell(T) \geq (4n_{\geq 3}(G) + 24)/13$. Note that this bound is tight if $y_1 = y + y'$ and that this is true for flower trees. □

4.4.5 Dealing with High Degree Vertices

In this section we complete the proof of Theorem 4.1 by presenting the proofs for Lemmas 4.14 and 4.19 which we omitted earlier.

Lemma 4.14 (Edge Deletion)**.** *Let G be an irreducible graph not equal to G_7 with adjacent vertices u and v. If $d(u) = d(v) = 4$, then uv is a bridge, $G - uv$ is cubic, or one of u, v becomes an inner vertex of a cubic diamond upon deletion of the edge uv.*

Proof. Suppose for the sake of contradiction that a non-bridge edge uv exists, between vertices of degree 4, such that $G - uv$ is not cubic and none of u, v becomes an inner vertex of a diamond upon deletion of uv.

Since G is irreducible, no reduction rule is admissible. Clearly, this must mean that a 2-necklace or 2-blossom is introduced when uv is deleted, that is when (R5) is applied to uv. In either case, we will derive a contradiction to the irreducibility of G.

Claim 1. The graph $G - uv$ does not contain a 2-necklace N.

Suppose for the sake of contradiction that $G - uv$ does contain a 2-necklace N. Consider N as a subgraph of G, that is uv is counted towards the degrees of u and v.

We first treat the case that N consists of at least two diamonds. If one of the diamonds in N contains three vertices of degree 3, we can use rule (R1), see Figure 4.30 (a). So now we may assume that one diamond on the end of the necklace contains u as one of the three vertices not shared with the next diamond, and the diamond on the other end of the necklace contains v this way.

If u is a connection vertex of N, then (R2) can be applied, see Figure 4.30 (b). This does not introduce a 2-necklace since the degree 4 vertex v is part of N on the other end. Because v is part of a diamond, this can also not introduce a 2-blossom.

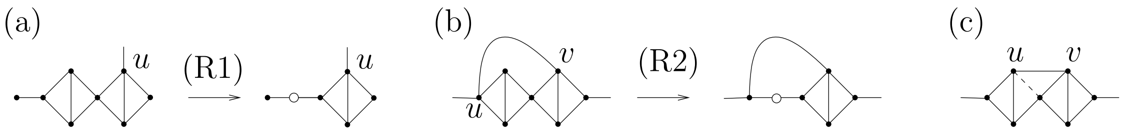

Figure 4.30. Reductions when a long 2-necklace is created.

In the remaining case, both u and v are internal vertices of their respective diamonds. Now it is admissible to apply (R5) to a different edge incident with u, see Figure 4.30 (c), where the dashed edge is the deleted one. This does not introduce a 2-necklace or 2-blossom. Note that u becomes part of a triangle that is induced by degree 3 vertices, for which all outgoing edges have different end vertices. Such a triangle cannot be part of a 2-blossom or 2-necklace. The other end vertex of the deleted edge is still part of a diamond after deletion, and thus is not part of a 2-blossom. It is not part of a 2-necklace since v is in this part of the necklace. This concludes the case where N consists of at least two diamonds.

Now suppose N consists of a single diamond. If u is an inner vertex of this diamond, then v cannot be part of the same diamond since we are dealing with simple graphs. This is then the case we excluded by assumption, see Figure 4.31 (a). So without loss of generality u is one of the connection vertices of the diamond.

Now rule (R1) or (R2) is admissible, depending on whether v is also in the diamond, see Figures 4.31 (b) and (c). This does not introduce a 2-blossom or 2-necklace, since in the case in Figure 4.31 (b), a triangle containing a goober is

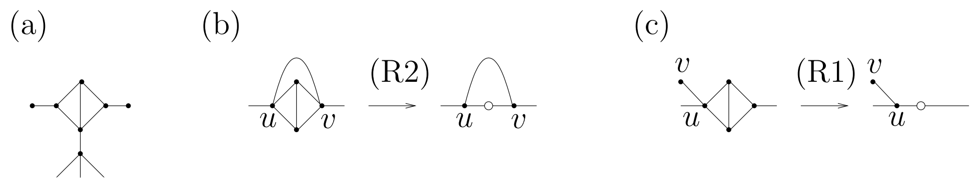

Figure 4.31. Reductions when a cubic diamond is created.

introduced, and in the case in Figure 4.31 (c), v has degree 4 and a goober at distance 2. Note that also no parallel edges are introduced. In the case shown in Figure 4.31 (c) the edges leaving the diamond are distinct, that is deleting uv does not yield a K_4. Otherwise uv would have been a bridge. This shows that it is admissible to apply either (R1) or (R2), which contradicts the irreducibility of G. $\triangle$

Claim 2. The graph $G - uv$ does not contain a 2-blossom B.

Suppose for the sake of contradiction that $G - uv$ does contain a 2-blossom B. For B we use the vertex labels from Figure 4.32 (a). The degree 4 vertex of B is labeled b, its terminals are called c-vertices, and the remaining four vertices are called its a-vertices. Now consider B as a subgraph of G, that is uv is counted towards the vertex degrees. Since $d_G(u) = 4 = d_G(v)$ neither of them is equal to b, since b has degree 4 even after the deletion of uv.

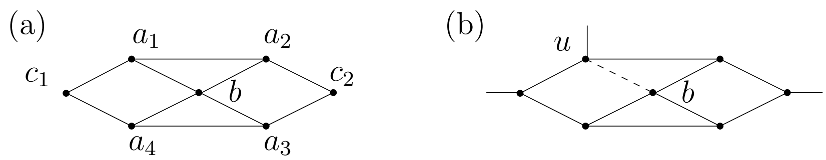

Figure 4.32. The blossom B after deleting uv.

If u is an a-vertex, say without loss of generality $u = a_1$, then it is admissible to delete the edge connecting u to b instead, see Figure 4.32 (b). We argue next that this does not introduce a 2-blossom or 2-necklace. Figure 4.33 shows the possible results of deleting ub in more detail, depending on the position of v.

First suppose $v \neq a_4$, that is we are in one of the situations show in Figures 4.33 (a) – (d). After deleting ub, b becomes part of a triangle that does not share a vertex with another triangle, since we assumed $v \neq a_4$. It follows that b is neither part of a 2-necklace, nor of a 2-blossom. The vertex u may be part of a triangle, these cases are shown in Figures 4.33 (c) and (d). But such a triangle is not part of a diamond, hence u is not part of a 2-necklace. Finally we argue that u is not part of a 2-blossom. Since b is not part of a 2-blossom, its neighbor a_2 is

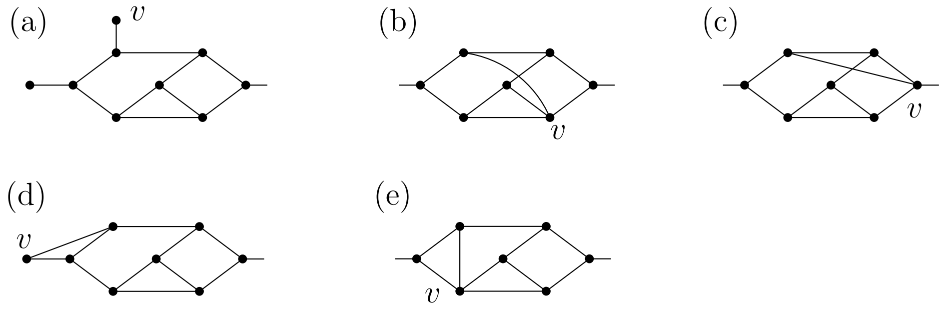

Figure 4.33. Possible results of deleting ub.

not part of a 2-blossom B' unless it is a terminal of B'. In that case it is not part of a triangle, but its neighbor c_2 is, which is impossible. Hence a_2 is not part of a 2-blossom. Thus, if u is part of a 2-blossom B', then it must be a terminal of B', and therefore not part of a triangle, but its neighbor c_1 must be part of a triangle. This is again not possible. This concludes the proof that if $v \neq a_4$, deleting ub is an admissible application of (R5).

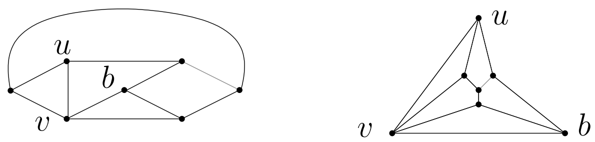

Figure 4.34. Deleting ub yields a blossom if $G = G_7$.

Now we need to consider the case that $u = a_1$ and $v = a_4$, see Figure 4.33 (e). Deleting ub does not introduce a 2-necklace, but there is exactly one way in which it may introduce a 2-blossom that has v as its central degree 4 vertex. Figure 4.34 shows this case, the black edges indicate the new blossom. But now it can be seen that the original graph which includes ub is exactly G_7. This contradicts the assumptions of the lemma. We conclude that if u is an a-vertex and $G \neq G_7$, in every case the edge ub can be deleted by an admissible application of (R5).

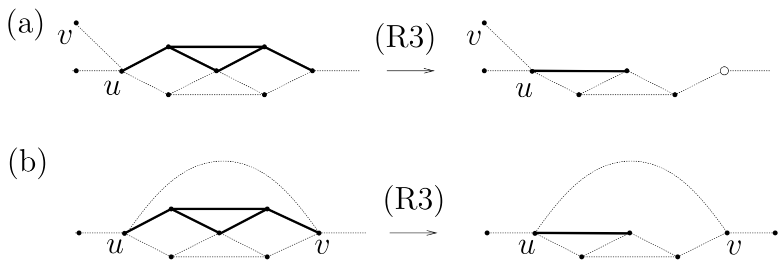

Figure 4.35. More reductions if a 2-blossom is created.

It remains to consider the case that u is a c-vertex. Then, (R3) could be used, see Figures 4.35 (a) and (b). The solid edges indicate the structure reduced by (R3). If in the first case a 2-necklace is introduced, v would be an inner vertex of one of its diamonds, but that is not possible since $d(v) = 4$. In the second case no 2-necklace can be introduced, since v is part of at most one triangle. In neither case a 2-blossom is introduced. $\triangle$

We have thus derived a contradiction to the irreducibility of the graph for all cases where deleting uv would not be an admissible application of (R5) and this proves the lemma. $\square$

Lemma 4.19 shows how to construct an extension F' for F when there is a vertex of degree at least 4 in F^C. Recall that when F and F' have the same number of components, the extension is valid if and only if $\Delta(\Delta n_{\geq 3,G}, \Delta\ell, \Delta\ell_d) \geq 0$, and if a new component is introduced we need $\Delta(\Delta n_{\geq 3,G}, \Delta\ell, \Delta\ell_d) \geq 6$.

Lemma 4.19 (Start Lemma)**.** *Let G be an irreducible graph that is not G_7 and does not have an edge e such that $G - e$ is cubic. Let F be a (possibly empty) subgraph of G, such that F^C contains at least one vertex of degree at least 4, and contains neither 2-necklaces nor 2-blossoms. Then, F is extendible.*

Proof. First suppose F is not the empty graph. If there is a non-leaf vertex v on the boundary of F, then F' can be obtained by expanding v. There is no leaf lost since v was not a leaf, and the newly added vertices are leaves. So the augmentation inequality is satisfied: $\Delta(k, k, 0) \geq 0$. Hence, we may assume in the remainder that only leaves of F have neighbors in $\overline{V(F)}$, or in other words, all vertices on the boundary of F are leaves of F.

The next step is the attempt to augment F using the operations (A1) – (A7), see Figure 4.36. The conventions for the figures in this proof are that encircled vertices belong to $V(F)$ and solid edges show the expansion. White vertices are goobers and other vertex degrees shown are to be understood as lower bounds, except when stated otherwise. Dead leaves are marked with a cross. All of the expansions in Figure 4.36 extend F without creating a new connected component, and satisfy $\Delta(\Delta n_{\geq 3,G}, \Delta\ell, \Delta\ell_d) \geq 0$. Thus, the resulting graph F' is an extension as claimed in the lemma. Together these augmentation rules yield the following claim.

Claim 1. The subgraph F is extendible, if a vertex in $V(F)$ has a goober neighbor in $\overline{V(F)}$ or at least two neighbors in $\overline{V(F)}$, or if there is a vertex $v \in \overline{V(F)}$ with $d_G(v) \geq 4$ at distance at most 2 from F.

If a goober from $\overline{V(F)}$ is adjacent to F then (A1) can be applied. If a vertex in $V(F)$ has at least two neighbors in $\overline{V(F)}$, (A2) can be applied. So from now on we will assume every vertex in $V(F)$ has at most one neighbor in $\overline{V(F)}$, and

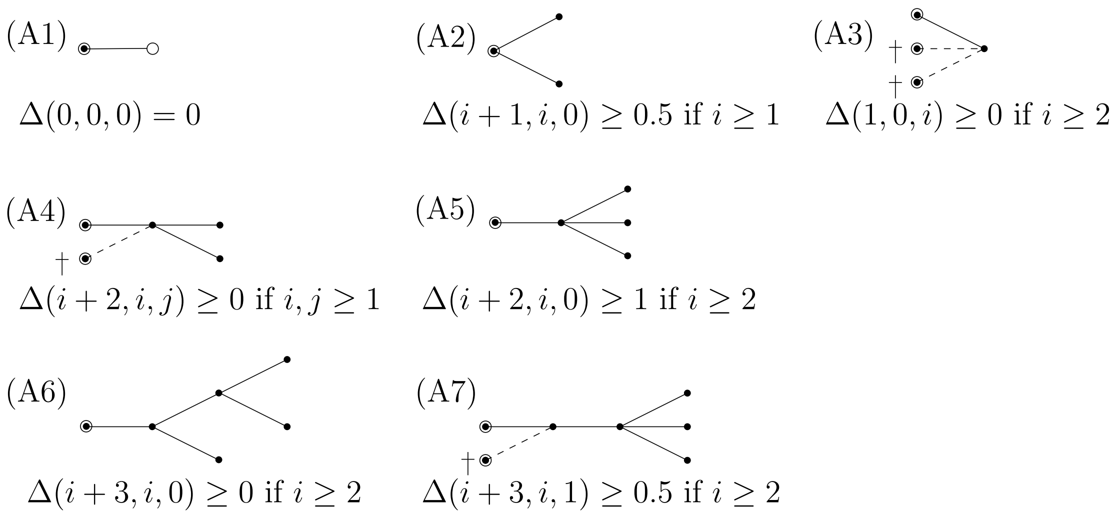

Figure 4.36. Simple augmentations of an existing subgraph.

this neighbor is not a goober. If a vertex $v \in \overline{V(F)}$ with $d_G(v) \geq 4$ is adjacent to a vertex in $V(F)$, then (A3), (A4) or (A5) can be applied. The creation of the dead leaves in (A3) and (A4) follows from the fact that (A2) cannot be applied anymore. If a vertex of degree at least 4 in $\overline{V(F)}$ has distance 2 from a vertex in $V(F)$, then (A6) or (A7) can be applied. $\triangle$

The rest of the proof will handle the more complicated cases when F is the empty graph, or the only vertices of degree higher than 4 in $\overline{V(F)}$ are at a larger distance from $V(F)$. We then introduce a new component of F. This is more complicated because adding a further component comes at a certain cost, more precisely we need that the new component satisfies $\Delta(\Delta n_{\geq 3,G}, \Delta\ell, \Delta\ell_d) \geq 6$.

The rest of the proof is divided into three more claims. The first one handles the easiest cases, and the second one handles all remaining cases except those where every degree 4 vertex is the common vertex of two edge-disjoint triangles. This final case is then taken care of in the third claim. Throughout the proof we assume, sometimes implicitly, that none of the situations that have been handled earlier can occur.

Claim 2. Let $v \in \overline{V(F)}$, $d(v) \geq 4$, and $w \in N(v)$. In the following four situations F is extendible: $d(v) \geq 5$, or $d(v) = 4$ and w is a goober, or $d(v) = d(w) = 4$, or $d(v) = 4$ and $N[w] \subset N[v]$.

First note that no vertex in $N[v]$ or $N[w]$ is part of F by Claim 1. If $d(v) \geq 5$, expanding v yields $\Delta(k+1, k, 0) \geq 6.5$, since $k \geq 5$. For $d(v) = 4$ and w a goober, expanding v gives $\Delta(4, 4, 0) = 6$.

Now suppose $d(v) = d(w) = 4$. If vw is a bridge, expanding v and w yields $\Delta(8, 6, 0) = 7$. Note that we assumed $G \neq G_7$ and that $G - uv$ is not cubic. Therefore, Lemma 4.14 shows that either v or w, say v, becomes the inner vertex of a cubic diamond upon deletion of the edge vw. If w has two neighbors not in $N[v]$, then expanding v, w yields $\Delta(7, 5, 1) = 6$, see Figure 4.37 (a). Otherwise w shares two neighbors with v and expanding v yields $\Delta(5, 4, 3) = 6.5$, see Figure 4.37 (b).

So now we may assume that all neighbors of v have degree 3. If $N[w] \subset N[v]$ then either the unique vertex $u \in N[v] \setminus N[w]$ has two neighbors not in $N[v]$, in which case expanding u, v gives $\Delta(7, 5, 1) = 6$, see Figure 4.37 (c), or there is another vertex $x \in N(v) - w$ with $N[x] \subset N[v]$, and v is expanded to obtain $\Delta(5, 4, 2) = 6$, see Figure 4.37 (d). $\triangle$

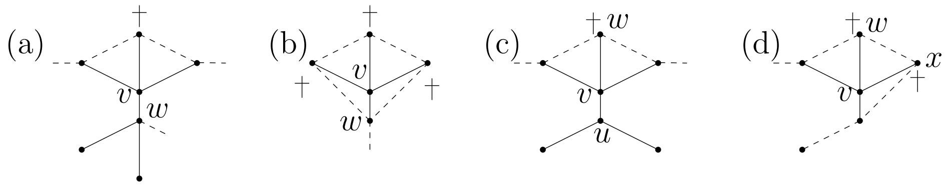

Figure 4.37.

Summarizing, we may now assume that F^C contains no vertices of degree at least 5, and if it contains a vertex v of degree 4, all neighbors of v have degree 3 and have either one or two neighbors not in $N[v]$.

Claim 3. If $\overline{V(F)}$ contains a vertex v with $d(v) = 4$ and a vertex $w \in N(v)$ which has two neighbors $a, b \notin N[v]$, then F is extendible.

We denote the other three neighbors of v by x, y, z. If one of a, b, x, y, z has all of its neighbors in $\{a, b\} \cup N[v]$, we have $\Delta(7, 5, 1) = 6$ by expanding v, w, see Figure 4.38 (a).

If a or b is a goober we obtain $\Delta(6, 5, 0) \geq 6.5$, see Figure 4.38 (b). If a or b is adjacent to a vertex $c \in V(F)$, then expanding v, w will make c a dead leaf and yields $\Delta(7, 5, 1) \geq 6$, see Figure 4.38 (c). If one of a, b, x, y, z has at least two neighbors not in $N[v] \cup \{a, b\}$, we obtain $\Delta(9, 6, 0) = 6$ by expanding v, w and this vertex, see Figure 4.38 (d).

Hence we may assume that a, b, x, y, z each have exactly one neighbor outside $N[v] \cup \{a, b\}$, and that this neighbor is not in F. Since they all have degree at least 3, the vertices a, b, x, y, z must induce three edges. This implies that one of a, b has degree 4 since we already know that x, y, z have degree 3. We may assume without loss of generality that $d(a) = 4$ and $d(b) = 3$. We distinguish two

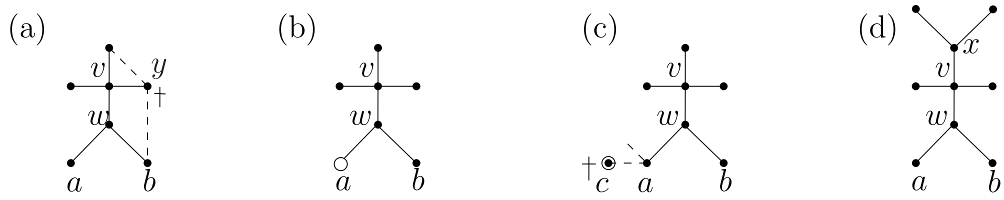

Figure 4.38.

cases depending on whether a is adjacent to b or not. We denote the neighbor of x outside of $N[v] \cup \{a, b\}$ by x', and similarly a', b', y', z' are defined.

Case 1. a is adjacent to x and y while b is adjacent to z.

Consider expanding v, x, z. All vertices in $\{a, b, w, y, x', z'\}$ are adjacent to at least one of v, x, z, thus we have $\Delta(9, 6, 1) = 6.5$ unless $x' = z'$, see Figure 4.39 (a). By an analogous argument with y in the place of x we may now assume that $x' = z' = y'$. Then, expanding v, x yields $\Delta(7, 5, 1) = 6$, since y becomes a dead leaf, see Figure 4.39 (b).

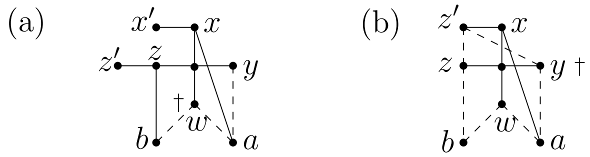

Figure 4.39.

Case 2. a is adjacent to b and x while y is adjacent to z.

If $x' \neq a'$, expanding v, x, a yields $\Delta(9, 6, 1) = 6.5$, see Figure 4.40 (a), so we may assume that $x' = a' =: c$, and this creates a situation symmetric in b and c. By Claim 2 we have that $b \not\sim c$. Now first suppose $b' \sim c$. Then expanding b', c, x, v yields $\Delta(10, 6, 3) = 6.5$, provided b' has a neighbor d other than y, z, see Figure 4.40 (b). Note that $d \in V(F)$ is not possible since augmentation (A6) could have been applied instead.

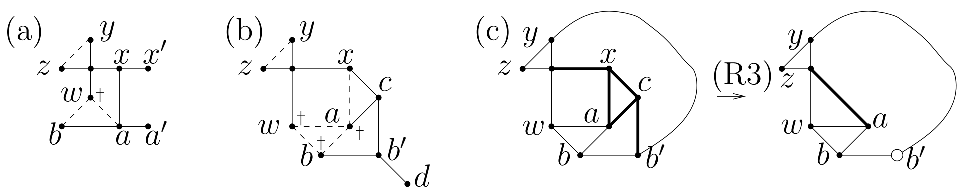

Figure 4.40.

If $N(b') = \{b, c, y\}$, then (R3) is admissible, see Figure 4.40 (c). Since b' becomes a goober this cannot introduce a 2-necklace. Hence it must be that $b' \sim y, z$ and the graph has $\Delta(9, 5, 5) = 6$, see Figure 4.41 (a). This concludes the cases with $b' \sim c$.

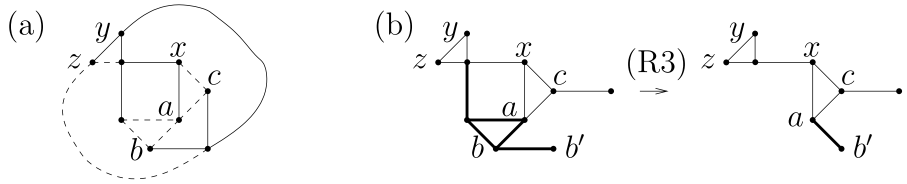

Figure 4.41.

The case $b' \not\sim c$ can be excluded because then (R3) would be admissible, see Figure 4.41 (b). Note that this cannot create a 2-necklace involving y, z since then (R2) would have been admissible. This concludes the proof of Claim 3. $\triangle$

Summarizing Claims 1–3, we may now assume that all neighbors of a degree 4 vertex $v \in \overline{V(F)}$ have degree 3, and have exactly one neighbor not in $N[v]$. In other words, v is the common vertex of two edge-disjoint triangles, see Figure 4.42.

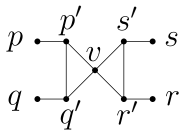

Figure 4.42. The bow tie subgraph

Claim 4. If the graph outside F contains a vertex v with $d(v) = 4$ such that all its neighbors have degree 3 and one neighbor outside $N[v]$, then F is extendible.

We denote the neighbors of v by p', q', r', s' and assume that $p' \sim q'$ and $r' \sim s'$. The neighbor of p' outside $N[v]$ is denoted by p and similarly q, r, s are defined, see Figure 4.42. We split the proof of the claim into three cases.

Case 1. We assume that $p = q$. If p has degree 2, then we expand v, p' to obtain $\Delta(5, 4, 2) = 6$. If p has degree 3, then we can apply (R1), see Figure 4.43 (a). So now without loss of generality p has degree 4. Then by Claim 3, p is also part of two edge-disjoint triangles. So if $p = r$ then also $p = s$. In that case we can expand v, p' to obtain $\Delta(6, 4, 4) = 6$, see Figure 4.43 (b). So now $p \neq r$, $p \neq s$. Consider applying (R2) to the diamond consisting of p, p', q', v, see Figure 4.43 (c). If this introduces a 2-necklace, (R1) could have been applied to the diamond on the other end of this necklace. It cannot introduce a 2-blossom since the triangles of a 2-blossom contain a degree 4 vertex.

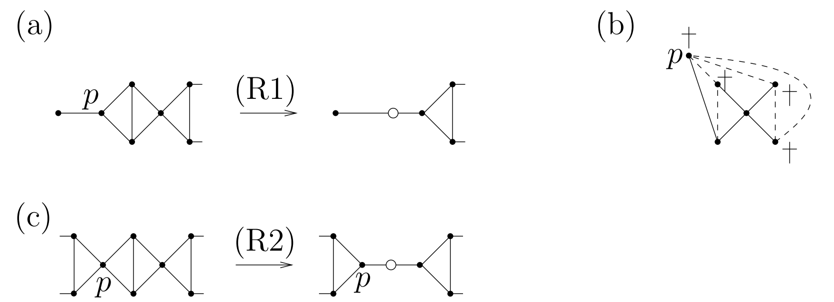

Figure 4.43.

Case 2. We assume that $p = r$. If p has degree 2, then we expand v, p' to obtain $\Delta(5,4,2) = 6$. Now we may assume that $d(p) = 3$ by Claim 3. Also $q \neq s$ since the graph does not contain 2-blossoms and the case that $d(q = s) = 4$ is again excluded by Claim 3. Thus, (R3) is admissible, see Figure 4.44.

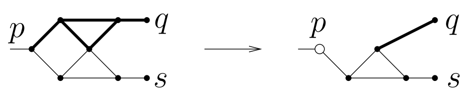

Figure 4.44.

Case 3. We assume that p, q, r, s pairwise different. In this case either (R4) or (R3) is admissible. If v is a cut vertex, (R4) may be used since it increases the number of components. Otherwise, we may assume without loss of generality that p and s are in same connected component of $G - v$ minus the edge sr, and (R3) can be applied without disconnecting the graph. See Figures 4.45 (a) and (b).

$\triangle$

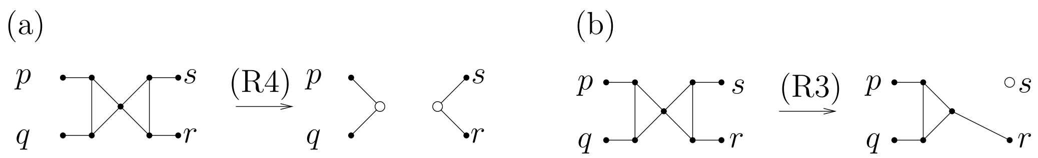

Figure 4.45.

This concludes all possible cases. Whenever the subgraph of G outside of F contains a vertex of degree at least 4, we have shown that G is either reducible, or F is extendible. $\square$

4.5 Conclusions

In Section 4.4 we have proven the main result of this chapter. Graphs without 2-necklaces and 2-blossoms have spanning trees with $n/3+4/3$ leaves. This result generalizes the main results from [54, 15] and can be used to obtain the fastest FPT algorithm for the decision problem MAXLEAF. With respect to Theomem 4.1, the following question remains open.

Problem 4.22. Is there an infinite family of irreducible graphs without 2-necklaces and 2-blossoms that has no spanning tree with more than $n_{\neq 2}/3+5/3$ leaves?

We have also given new proofs for several results of the same type for other graph classes. Our proof method also enabled us give more compact proofs and strengthen some known results.

In Section 4.2 we have improved the result that every graph without triangles with minimum degree at least 3 has a spanning tree with at least $n/3+4/3$ leaves. We prove that every graph without triangles has $n_{\neq 2}/3+2/3$ leaves. It remains open to improve the additive constant in this result.

Problem 4.23. Show that every graph without triangles has a spanning tree with at least $n_{\geq 3}/3+4/3$ leaves.

In discussions with Stefan Felsner and Eric Fusy we have realized that every *stacked triangulation* has a spanning tree with at least $2n/3+1/3$ leaves. This bound is tight for the K_4. An easy proof of this fact uses *Schnyder woods*. A stacked triangulation has a unique Schnyder wood $S=(T_1,T_2,T_3)$ that is the corresponding *3-orientation* is acyclic. Therefore, each face F is incident to an inner vertex v such that both edges incident to F and v are outgoing at v. Hence, there is a color i in which v has no incoming edges, and thus v is a leaf of T_i. Furthermore, the outgoing edges of v in colors $i-1$ and $i+1$ lie on at most one common face, i.e. v is counted in this way as a leaf of T_i at most once. Therefore the three trees have at least $2n-5$ leaves altogether and there must be one tree which has $(2n-5)/3$ leaves among the inner vertices of the stacked triangulation. This tree can be augmented to a spanning tree such that two of the outer vertices also become leaves. All triangulations are 3-connected and stacked triangulations have many 3-cuts. Intuitively, higher connectivity should favor spanning trees with many leaves. Therefore, we propose the following problem.

stacked triangulation

Schnyder wood

3-orientation

Problem 4.24. Does every triangulation with n vertices have a spanning tree with at least $2n/3+1/3$ leaves?

Chapter 5

Small Integer Realizations of Stacked Polytopes

In this chapter we are concerned with polytopes, a central topic of discrete geometry. We focus on *3-polytopes* (that is 3-dimensional polytopes) which have especially kindled the interest of researchers for "obvious" reasons. One of the outstanding results in the theory of 3-polytopes is Steinitz's Theorem [86, 87]. It exhibits a beautiful connection between 3-polytopes and 3-connected planar graphs.

3-polytope

Theorem 5.1 (Steinitz's Theorem)**.** *The edge graphs of* 3*-polytopes are in bijection with the* 3*-connected planar graphs.*

This result completely characterizes all the combinatorial types of 3-polytopes and justifies to use the term of a *combinatorial 3-polytope*, that is a 3-connected planar graph. A *geometric 3-polytope*, that is a convex hull of a point set in $\mathbb{R}^3$, can be seen as a *realization* of a combinatorial polytope. Of course every combinatorial 3-polytope P has infinitely many realizations. The question whether there exist realizations such that all vertices of P are realized with integral coordinates has been answered in the affirmative. Given the existence of such integer realizations, the next thing to ask for is a bound $B(n)$ such that every combinatorial 3-polytope with n vertices has an integer realization with vertex coordinates of absolute value at most $B(n)$. This question is interesting because the bound $B(n)$ allows to bound the ratio of the longest to the shortest edge length, which in turn is useful for efficient visualization of 3-polytopes. To the best of our knowledge the current records for the best bounds are held by Ribo, Rote, and Schulz [73, 74] and we summarize their results in the following theorem.

combinatorial/geometric 3-polytope

realization

Theorem 5.2. *Let P be a combinatorial 3-polytope with n vertices. If P contains a triangle (not necessarily a face), then it can be realized with integral coordinates smaller than 29^n. If P contains no triangle, but at least one quadrangle, then it can be realized with integral coordinates smaller than 47^n. Without further assumptions P can be realized with integral coordinates of absolute value less than 188^n.*

These upper bounds are accompanied by a lower bound of $\Omega(n^{3/2})$. This bound is based on the fact that a strictly convex *grid drawing* of an n-gon needs $\Omega(n^{3/2})$ space, see [1, 90, 2]. A grid drawing of a planar graph is a crossing-free straight

grid drawing

line embedding with integral vertex coordinates and convex faces. In general the faces of a grid drawing do not have to be strictly convex, but of course every grid drawing of a triangulation has strictly convex faces. Note that it is not even known whether the right upper bound for the size of integer realizations of 3-polytopes is polynomial or exponential in n.

In this chapter we focus on realizations of stacked triangulations. We postpone to give a formal definition and content ourselves for the moment with the intuition that realizations of stacked triangulations can be obtained by glueing together simplices. This special structure allows for inductive methods to be used when dealing with stacked triangulations. Another advantage of this class is that every grid drawing of a stacked triangulation in the plane can be lifted to a corresponding *stacked polytope*. A *lifting* of a grid drawing of a 3-connected planar graph G is an assignment of a third coordinate to every vertex position, such that the resulting polytope is a realization of G. Even the smallest non-stacked triangulation, which has six vertices, has non-liftable grid drawings, see [99].

stacked polytope

lifting

There are basically three types of proofs of Steinitz's Theorem. A modern version of Steinitz's own approach can be found in [99]. From this proof it can also be seen that every 3-polytope has an integer realization. A quantitative analysis of the method yields doubly exponential bounds because realizing a polytope with this approach involves going back and forth between a polytope and its polar polytope. The second approach uses the Koebe-Andreev-Thurston Circle Packing Theorem, see [10, 100]. As this may yield irrational coordinates we follow the third approach which uses Tutte's Theorem, see Theorem 5.7 below and [93]. A good exposition of this approach is given by Richter-Gebert in [78] where he also develops a lifting method that we use in this chapter. This lifting method uses grid drawings with associated edge weights such that the so-called equilibrium condition is satisfied at every vertex. To produce small realizations one needs small grid drawings that also allow for small edge weights. This is the approach that we take to produce polynomial integer realizations of *balanced stacked triangulations* which have the maximum possible number of vertices of every height. Linear stacked triangulations have the minimum possible number of vertices of every height, i.e. they have one vertex of every height. For *linear stacked triangulations* we give grid drawings accompanied by an explicit lifting function. This leads to the following results.

balanced/linear stacked triangulation

Theorem 5.3. *Every linear stacked triangulation with n vertices has an integer realization with coordinates of order $O(n^4)$. Balanced stacked triangulations have integer realizations of order $O(n^{2.47})$ and this implies that every stacked triangulation can be realized with coordinates of order 15^n.*

realization

Small grid drawings of 3-connected planar graphs are interesting in their own right and, as opposed to 3-polytopes, good bounds have been obtained. The best known bound for the grid size is $(n-2)\times(n-2)$ and it can be obtained by a variant of the face-counting approach using *Schnyder woods*, see Theorem 1.6 and [83, 38]. *Schnyder wood* The other approach to obtain drawings of size $(n-2)\times(n-2)$ was developed by Chrobak and Kant and extends a partial drawing vertex by vertex, see [24]. A natural idea is to try to lift these efficient drawings. As mentioned above, every grid drawing of a stacked triangulation can be lifted and thus, by the Maxwell-Cremona Theorem (see [28]), admits edge weights satisfying the equilibrium conditions. The task is then to find good upper bounds for admissible edge weights. We have tried to analyze the edge weights for drawings produced by the Schnyder wood method. They did not appear to yield polynomially bounded weights even for balanced stacked triangulations. Due to its iterative nature, the approach of Chrobak and Kant does not seem suitable for being used within this lifting framework.

In the next section we introduce the facts about stacked triangulations that we need in this chapter. In Section 5.2 we discuss polynomial liftings of linear stacked triangulations and in Section 5.3 we treat the case of balanced stacked triangulations. In Section 5.4 we give bounds for integer realizations of brooms, i.e. polytopes that arise from glueing a balanced with a linear stacked polytope.

5.1 Preliminaries

We start by recalling the definition of stacked triangulations which we have already seen in Section 1.4. When working with a stacked triangulation T we always assume that one face is marked, and that this face is the outer face when we use an embedding of T. We define the class of stacked triangulations inductively. We do not need edge orientations in this chapter, and we use the simpler notation uv instead of $\{u,v\}$ to denote an edge between vertices u and v.

- K_3 is a *stacked triangulation*. *stacked triangulation*

- Let $T=(V,E)$ be a stacked triangulation and $\{u,v,w\}$ an unmarked face of T. Then, for a vertex $v'\notin V$, $T'=(V\cup\{v'\}, E\cup\{v'u,v'v,v'w\})$ is a stacked triangulation.

The *height* of the outer vertices is defined to be -1. For an inner vertex v' stacked *height* into a triangle $\{u,v,w\}$ its height is $h(v)=\max\{h(u),h(v),h(w)\}+1$. The height of a face $F=\{u,v,w\}$ is $h(F)=\max\{h(u),h(v),h(w)\}+1$.

With a stacked triangulation T we associate a rooted ternary tree $\mathcal{T}$, see Figure 5.1. Let a crossing-free embedding of T be given. The vertices of $\mathcal{T}$ represent the triangles of T and there is an edge vw in $\mathcal{T}$ if and only if w represents the smallest triangle Δ_w that contains Δ_v in the given embedding. Note that the containment relation only depends on the choice of the outer face, which is fixed by assumption. In order to work with $\mathcal{T}$ it is useful to observe that the leaves of $\mathcal{T}$ are in bijection with the inner faces of T and the inner vertices of $\mathcal{T}$ are in bijection with the inner vertices of T.

- With K_3 we associate a singleton as its tree.
- Let $T = (V, E)$ be a stacked triangulation and $\mathcal{T}$ its tree. Let T' be obtained from T by stacking a vertex $v' \notin V$, into the triangle Δ. The tree $\mathcal{T}'$ associated with T' is obtained from $\mathcal{T}$ by adding three new leaves to the leaf of $\mathcal{T}$ representing Δ.

For every vertex v its height in $\mathcal{T}$ is the number of edges on the shortest path from v to the root. Hence, the height of a vertex or face in T is the same as its height in $\mathcal{T}$. We say that a stacked triangulation has height h if its associated tree has height h.

balanced/linear stacked triangulation

broom

We work with three subclasses of stacked triangulations. A stacked triangulation is a *balanced stacked triangulation* if its associated tree is complete, see Figure 5.1 (a). A stacked triangulation is a *linear stacked triangulation* if its associated tree is a caterpillar, that is a path plus leaves, see Figure 5.1 (b). A stacked triangulation is a *broom* if the associated tree can be obtained from a caterpillar by substituting a leaf of maximum height by a complete ternary tree. For example a broom can be obtained by identifying the outer face of the triangulation in Figure 5.1 (a) with the shaded face of the triangulation in Figure 5.1 (b).

(a) (b)

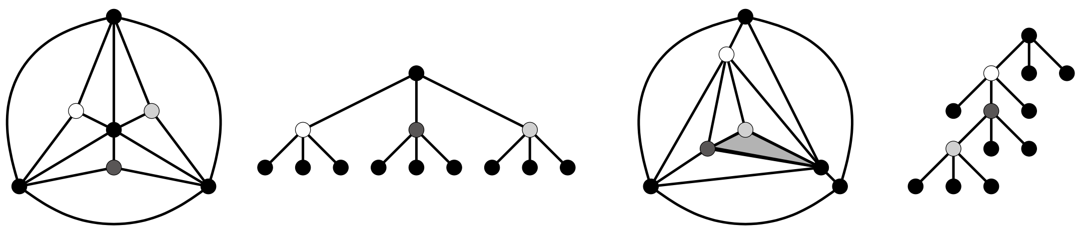

Figure 5.1. A balanced and a linear stacked triangulation with the respective trees.

We collect some easy statistics for these classes of stacked triangulations. A linear stacked triangulation of height h has $h+3$ vertices, $3h+3$ edges, and $2h+1$ bounded faces. A balanced stacked triangulation of height h has $3+\sum_{i=0}^{h-1} 3^i = (3^h+5)/2$ vertices, $3(3^h+1)/2$ edges, and 3^h bounded faces. A balanced stacked triangulation of height h has 3^{h-1} vertices on level $h-1$, and thus about $2/3$ of its vertices have height $h-1$.

5.2 Realization of Linear Stacked Triangulations

We first study the structure of linear stacked triangulations in more detail. For the rest of this section, T shall be a linear stacked triangulation with n vertices and $\mathcal{T}$ its associated tree.

Lemma 5.4. *Let T be a linear stacked triangulation with height $h \geq 1$ and outer face a_1, a_2, a_3. Then, there is a permutation (i, j, k) of $(1, 2, 3)$ such that a_i has degree 3 and $T - a_i$ is a linear stacked triangulation. The outer face of $T - a_i$ is $\{a_j, a_k, v\}$ where v is the third vertex incident to a_i.* *linear stacked triangulation*

Proof. Let v be the unique vertex of T of height 0. Then, because $\mathcal{T}$ is a caterpillar, at most one of the triangles $F_3 = \{a_1, a_2, v\}$, $F_2 = \{a_1, a_3, v\}$, $F_1 = \{a_2, a_3, v\}$ contains further vertices. We may assume that this triangle is F_3, and thus the only neighbors of a_3 are a_1, a_2, and v. Furthermore, v is the root of the tree $\mathcal{T}$ associated with T, and it is adjacent to the two leaves of $\mathcal{T}$ that represent F_1 and F_2. Let $\mathcal{T}'$ be the tree obtained from $\mathcal{T}$ by deleting v, F_1, and F_2. The root of $\mathcal{T}'$ is the unique non-leaf vertex of height 1 of $\mathcal{T}$. The triangulation $T - a_3$ can be obtained by starting with the triangle $\{a_1, a_2, v\}$ and then stacking vertices as encoded in $\mathcal{T}'$. $\square$

We use Lemma 5.4 inductively to remove vertices from T until only a facial triangle remains which we call the *central triangle*. Using the order of removals, we define a *shelling order* s on T. The three vertices of the central triangle receive the numbers $1, 2, 3$. If the vertex v was removed as the ith vertex, we define $s(v) = n - i + 1$, see Figure 5.2. Furthermore, we call an edge a *spine edge* if it is the unique edge incident to a removed vertex that is not incident to the outer face when this vertex is removed. The three outer vertices of T are each incident to one spine edge and the same is true for the three vertices of the central triangle. Every other vertex is incident to two spine edges and the subgraph induced by the spine edges is cycle-free. Thus, this induced subgraph consists of three disjoint paths, the *spines* S_i, $i = 1, 2, 3$ of T which connect the outer triangle with the central triangle. We will denote the triangle that constitutes the outer face before v is removed as $\Delta_{s(v)}$ and the central triangle as Δ_3. This yields a sequence of $n - 2$ triangles. *central triangle* *shelling order* *spine edge* *spine S_i*

We now define the embedding M_T of T that will be lifted. For $v \in S_i$ let $s_i(v)$ be the number of vertices that lie on S_i and come before v in the shelling order s, see Figure 5.2. We define M_T as

$$M_T(v) = \begin{cases} (s_1(v) + 1, 0) & \text{if } v \in S_1 \\ (0, s_2(v) + 1) & \text{if } v \in S_2 \\ (-s_3(v) - 1, -s_3(v) - 1) & \text{if } v \in S_3 \end{cases}$$

and connect adjacent vertices by straight line segments, see Figure 5.2. It is easy to see that this yields a grid embedding. Furthermore, the whole embedding is contained in a square of side length n, and $\Delta_3, \ldots, \Delta_n$ is a sequence of geometric triangles of the form $\{(i, 0), (0, j), (-k, -k)\}$.

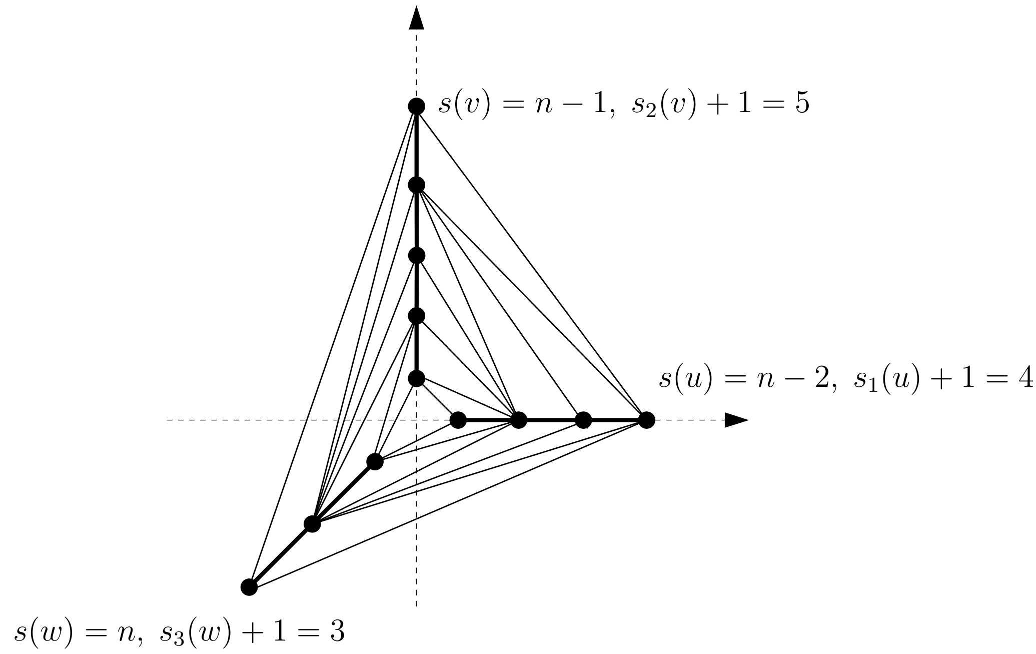

Figure 5.2. A grid embedding M_T of a linear stacked triangulation T. The thick edges are the spine edges of M_T.

We lift M_T by defining a piecewise linear convex function $f_T : \mathbb{R}^2 \to \mathbb{R}_0^+$ that produces creases for all inner edges of T. Note that every non-spine edge lies in some triangle Δ_α. We define f_T as a sum of piecewise linear functions.

$$f_T(x, y) = \sum_{\alpha=3}^{n} f^{\Delta_\alpha}(x, y)$$

Let Δ_α be embedded as $\{(i, 0), (0, j), (-k, -k)\}$. Then, f^{Δ_α} is a piecewise linear convex function that produces creases exactly on the line segments connecting the vertices of Δ_α and the rays

$$\{(t, 0) | t \geq i\},\ \{(0, t) | t \geq j\},\ \{(-t, -t) | t \geq k\}.$$

We define and study the functions f^{Δ_α} in Lemma 5.5 and illustrate them in Figure 5.3.

Lemma 5.5. *Let Δ_α be a triangle with vertices $u = (i,0), v = (0,j)$ and $w = (-k,-k)$. For $(x,y) \in \mathbb{R}^2$ we define $f^{\Delta_\alpha} = f^{ijk} : \mathbb{R}^2 \to \mathbb{R}_0^+$*

$$f^{ijk}(x,y) \quad = \quad \begin{cases} f_1^{ijk}(x,y) & \text{if } x \geq 0, y \geq 0, y \geq -\frac{j}{i}x + j \\ f_2^{ijk}(x,y) & \text{if } x \geq y, y \leq 0, y \leq \frac{k}{k+i}x - \frac{ik}{k+i} \\ f_3^{ijk}(x,y) & \text{if } x \leq 0, y \geq x, y \geq \frac{j+k}{k}x + j \\ f_4^{ijk}(x,y) & \text{otherwise} \end{cases}$$

where

$$\begin{aligned} f_1^{ijk}(x,y) &= k \cdot (i \cdot y + j \cdot x - i \cdot j), \\ f_2^{ijk}(x,y) &= j \cdot (-(i+k) \cdot y + k \cdot x - i \cdot k), \\ f_3^{ijk}(x,y) &= i \cdot (k \cdot y - (j+k) \cdot x - j \cdot k), \\ f_4^{ijk}(x,y) &= 0. \end{aligned}$$

Then, f^{ijk} is a well defined, piecewise linear, continuous, convex function that maps grid points to integer values. Furthermore, f^{ijk} is not differentiable exactly on the set (i.e. f^{ijk} forms creases on),

$$\begin{aligned} A \quad = \quad & \{(x,y) | y = -\tfrac{j}{i}x + j, 0 \leq x \leq i\} \cup \\ & \{(x,y) | y = \tfrac{k}{k+i}x - \tfrac{ik}{k+i}, -k \leq x \leq i\} \cup \\ & \{(x,y) | y = \tfrac{j+k}{k}x + j, -k \leq x \leq 0\} \cup \\ & \{(x,y) | y = 0, x \geq i\} \cup \{(x,y) | y \geq j, x = 0\} \cup \{(x,y) | y = x, x \leq -k\}. \end{aligned}$$

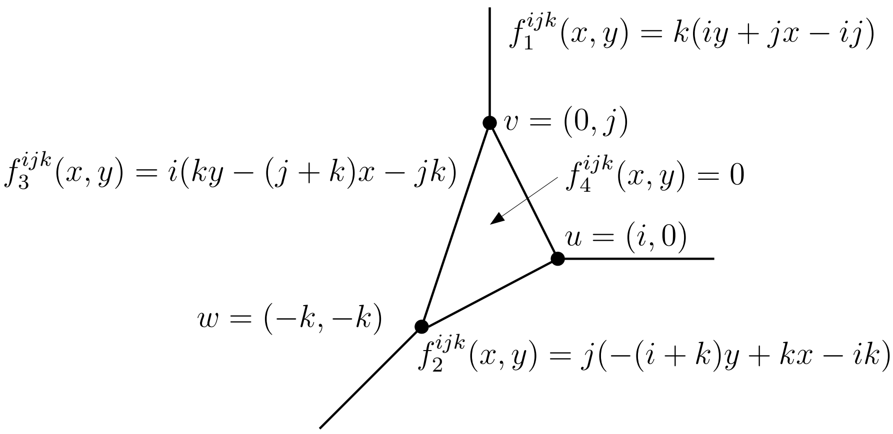

Figure 5.3. The lifting function f^{ijk}.

Proof. We omit the index ijk in the proof and show that

$$f(x,y) = \max\{f_q(x,y) | q = 1,\dots,4\}. \tag{5.1}$$

The calculations for the proof of (5.1) are not hard and we only show that for $(x,y) \in \{(x,y) | x \geq y, y \leq 0, y \leq \frac{k}{k+i}x - \frac{ik}{k+i}\}$ the maximum is attained by f_2.

$$\begin{aligned}
f_2(x,y) - f_1(x,y) &= -ijy - jky - iky + (jkx - jkx) + (ijk - ijk) \geq 0\\
f_2(x,y) - f_3(x,y) &= (x-y)(ij + ik + jk) \geq 0\\
f_2(x,y) - f_4(x,y) &= -j(i+k)y + jkx - ijk \geq -jkx + ijk + jkx - ijk = 0
\end{aligned}$$

This proves that f is a well defined, piecewise linear, continuous, convex, function that maps grid points to integer values. It is easy to see from the definitions of the f_q that the four planes of the form $H_q = (x, y, f_q(x,y))$, are given by the equations below. For example $z = f_2(x,y) = -j(i+k)y + jkx - ijk$ and this verifies the description of H_2 below.

$$\begin{aligned}
-jkx - iky + z &= -ijk && \text{for } q = 1,\\
-jkx + j(i+k)y + z &= -ijk && \text{for } q = 2,\\
i(j+k)x - iky + z &= -ijk && \text{for } q = 3,\\
z &= 0 && \text{for } q = 4.
\end{aligned}$$

Since the normal vectors are not parallel this shows that the planes H_i are not parallel either and it follows that f produces creases exactly on A. □

We now formulate the main result of this section and conclude that linear stacked polytopes can be realized with polynomially bounded integer coordinates.

realization **Theorem 5.6.** *Let T be a linear stacked triangulation. Then, T has a realization with integer coordinates of absolute value bounded by $n \times n \times 3n^4$.*

lifting *Proof.* We embed T with the embedding M_T and define the *lifting* for $(x,y) \in \mathbb{R}^2$ as $f_T(x,y)$. If f and g are continuous, piecewise linear and convex functions not differentiable on A_f and A_g respectively, then $f+g$ is continuous, piecewise linear and convex as well. Furthermore, the set A_{f+g} on which $f+g$ has no derivative is exactly $A_{f+g} = A_f \cup A_g$. Thus f_T is continuous, piecewise linear and convex and produces creases exactly on the line segments representing the non-spine edges of T and the three rays $\{(t,0) | t \geq 1\}$, $\{(0,t) | t \geq 1\}$, $\{(-t,-t) | t \geq 1\}$. Let H be the plane defined by the images of the vertices of Δ_n and H^+ the halfspace of H that contains the origin. Then, the set

$$H^+ \cap \{(x,y,z) \in \mathbb{R}^3 | z \geq f_T(x,y)\}$$

is a realization of T. The bounds for the x- and y-coordinate follow from the embedding M_T. Each of the functions f^{Δ_α} is bounded on the square $[0,n] \times [0,n]$ by $3n^3$ and f_T is a sum of $n-4$ of these functions. This shows the bound for the z-coordinate. □

5.3 Realization of Balanced Stacked Triangulations

We denote the *balanced stacked triangulation* of height h by B_h. We first explain the lifting framework that we use for the realization of balanced stacked triangulations. We then show how balanced stacked polytopes can be realized efficiently.

balanced stacked triangulation B_h

In the introduction of this chapter we have mentioned that good grid drawings of stacked triangulations exist, that every such drawing is liftable and also that this is not sufficient to obtain small integer realizations. We now sketch how the edge weights come into play in Richter-Gebert's lifting method [78] and what further ingredient is needed to obtain good liftings with this approach. Let M be a planar map, that is a crossing-free embedding of a planar graph into $\mathbb{R}^2$. The *cells* of M are the connected regions of $\mathbb{R}^2 \setminus M$. We denote the bounded cells by c_i for $i = 1, \ldots, f-1$ and by c_0 the unbounded outer cell. Given an embedding of a graph with vertex positions $p(v_1), \ldots, p(v_n)$ and edge weights w_{v_i,v_j} we say that a vertex u is in equilibrium if

cell

$$\sum_{v_i : v_i u \in E} w_{v_i,u}(p(v_i) - p(u)) = 0.$$

Theorem 5.7. *[Tutte's Theorem] Let $G = (\{1, \ldots, n\}, E)$ be a 3-connected planar graph that has a cell c_0 with vertices $(k+1, \ldots, n)$ for some $k < n$. Let $p_{k+1}, \ldots, p_n$ be the vertices (in this order) of a convex $(n-k)$-gon and E' the set of edges that are not in c_0. Let $w : E' \mapsto \mathbb{R}^+$ be an assignment of positive weights to the edges from E'.*

1. *There are unique positions $p_1, \ldots, p_k \in \mathbb{R}^2$ for the interior vertices such that all interior vertices are in equilibrium.*

2. *The bounded cells of this embedding of G are then realized as non-overlapping strictly convex polygons.*

We call an embedding as described in the theorem a *Tutte embedding*. Next, we briefly outline how Richter-Gebert [78] defines a lifting function for a given Tutte embedding $\mathcal{P} = (p_1, \ldots, p_n)$ of a graph G with edge-weights $w_{u,v}$ for $uv \in E$. This lifting function is used to define a realization of a polytope $P(G)$ with edge graph G. We assume that $\mathbb{R}^2$ is embedded into $\mathbb{R}^3$ at the plane $z = 1$ and each p_i has *homogenized coordinates* $(x_i, y_i, 1)$.

Tutte embedding

homogenized coordinates

For an oriented edge (b, t) of G there is a unique adjacent cell L to the left of it and a unique adjacent cell R to the right of it. We call the ordered quadruple $(b, t \mid L, R)$ an *oriented patch* of (G, P). If $(b, t|L, R)$ is an oriented patch, then $(t, b \mid R, L)$ is as well. For every interior cell c_i we define a *lifting vector* $q_i \in \mathbb{R}^3$ as follows.

oriented patch $(b, t \mid L, R)$

lifting vector q_i

- $q_1 = (0, 0, 0)$
- $q_L = w_{b,t}(p_b \times p_t) + q_R$ if $(b, t|L, R)$ is an oriented patch of (G, P).

In [78] it is shown that the equilibrium condition of the Tutte embedding implies that the q_i are well-defined.

Let $c_{i_0} = c_1$, and let $c_{i_0}, c_{i_1}, \ldots, c_{i_\ell}$ be a sequence of cells such that $(t_{i_j}, b_{i_j}|c_{i_{j-1}}, c_{i_j})$ is a patch of (G, P) for $j = 1, \ldots, \ell$. We define the lifting function $f_{\mathcal{P}}$ with domain $\mathrm{conv}(p_{n-k}, \ldots, p_n)$ for $p = (x, y, 1) \in c_{i_\ell}$ as

$$f_{\mathcal{P}}(p) = \langle p, q_{i_\ell} \rangle = \langle p, \sum_{j=1}^{\ell} w_{b_{i_j}, t_{i_j}}(p_{b_{i_j}} \times p_{t_{i_j}}) \rangle = \sum_{j=1}^{\ell} w_{b_{i_j}, t_{i_j}} \det(p, p_{b_{i_j}}, p_{t_{i_j}}).$$

In [78] this approach is used to prove upper bounds for the size of integer realizations of 43^n if there is a triangular face and 2^{13n^2} in general. Our focus is on realizing stacked triangulations, which have only triangular cells. The next lemma follows from the definition of $f_{\mathcal{P}}$, and the observation that the determinant $\det(p, p_{b_{i_j}}, p_{t_{i_j}})$ corresponds to twice the area of the triangle spanned by p, $p_{b_{i_j}}$, and $p_{t_{i_j}}$.

Lemma 5.8. *Let $\mathcal{P}$ be a Tutte embedding of a triangulation with non-negative integral vertex coordinates no larger than $c_1 \cdot n^p$ and integral edge weights bounded by $c_2 \cdot n^q$. Then there is a realization of P with integral vertex coordinates bounded by $c_1^2 \cdot c_2 \cdot n^{2p+q+1}$.*

We now present an embedding M_h of B_h that can be lifted with non-negative integral vertex coordinates bounded by $4n/3 \times 4n/3 \times 16n^3/9$. This follows from Lemma 5.8 since M_h can be translated to a Tutte embedding with non-negative integral vertex coordinates no larger than $4n/3$ and all edge weights equal to 1. In Theorem 5.10 we show that a better bound can be obtained.

Lemma 5.9. *Let B_h be the balanced stacked triangulation of height h.*

- *Let M_1 be the straight line embedding of K_4 into $\mathbb{R}^2$ with vertex coordinates $(-1, -1), (1, 0), (0, 1), (0, 0)$.*
- *Let M_{h+1} be the straight line embedding obtained from M_h by multiplying all vertex coordinates by 3 and adding a new vertex into the barycenter of every bounded face.*

Then M_h is a Tutte embedding of B_h with integral vertex coordinates and edge weights all equal to 1 for all $h \in \mathbb{N}$. Furthermore, all vertices have coordinates in $\{-3^{h-1}, \ldots, 3^{h-1}\}^2 \subset \{\lceil -2n/3 \rceil, \ldots, \lfloor 2n/3 \rfloor\}^2$.

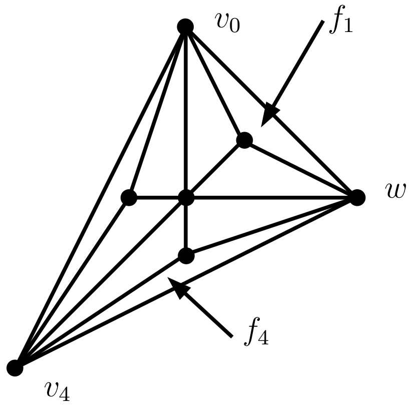

Figure 5.4. The embedding M_2 of the balanced stacked triangulation B_2.

Proof. The claims obviously hold for M_1, so we proceed by induction. For an example see Figure 5.4. Say the claim holds for all $1 \leq i \leq h$. In M_{h+1} all vertices of height $i \leq h-1$ have integer coordinates by the induction hypothesis. A vertex p of height h has coordinates of the form $1/3 \cdot (3p_1 + 3p_2 + 3p_3)$ where the p_i are vertices of height at most $h-1$. As the p_i all are integral, so is p. With regard to the edge weights, we know that M_h is in equilibrium with edge weights 1.

Thus, $3 \cdot M_h$ is in equilibrium as well and we only have to show that the map M'_{h+1} induced by the edges incident to vertices of height h is in equilibrium when all edge weights are chosen to be 1. The vertices of height h lie in the barycenter of their neighbors and are therefore in equilibrium. Let p be a vertex of height $1 \leq i \leq h-1$ and $p_1, \ldots, p_r$ its neighbors in M'_{h+1}, i.e. its neighbors of height h. Furthermore, let $u_1, \ldots u_r$ be the neighbors of p of height at most $h-1$. Then,

$$\begin{aligned}\sum_{i=1}^{r}(p_i - p) &= (1/3(u_1 + u_r + p) - p) + \sum_{i=1}^{r-1}(1/3(u_i + u_{i+1} + p) - p) \\ &= 1/3\left((u_1 - p) + (u_r - p) + \sum_{i=1}^{r-1}(u_i - p) + \sum_{i=2}^{r}(u_i - p)\right) = 0.\end{aligned}$$

The last equality uses the induction hypothesis. We have seen in Section 5.1 that B_h has $(3^h + 5)/2$ vertices which proves the last claim. □

Theorem 5.10. *There is a realization P_h of B_h with non-negative integral vertex coordinates bounded by $4n/3 \times 4n/3 \times O(n^{2.47})$.* *realization*

Proof. We consider the embedding M_h of B_h embedded in $\mathbb{R}^3$ at $z = 1$. Let $w(h) = (3^{h-1}, 0, 1)$, and let its neighbors in counterclockwise order starting with $(0, 3^{h-1}, 1)$ be labeled $v_0(h), \ldots, v_{2^h}(h)$, that is $(-3^{h-1}, -3^{h-1}, 1) = v_{2^h}(h)$, see Figure 5.4.

The bounded faces incident to w are labeled $f_1, \ldots, f_{2^h}$ in counterclockwise order starting with the face incident to $v_0(h)$. If we fix q_1, the lifting vector associated with f_1 to be $(0,0,0)$ then $v_{2^h}(h)$ will have the largest z-coordinate in the resulting *lifting* of B_h and we will calculate this value now. For the following calculations we use the definition of the lifting vectors and the fact that the $v_i(h)$ are convex combinations of $v_0(h)$, $v_{2^h}(h)$ and $w(h)$ for $1 \le i \le 2^h - 1$.

lifting

$$\begin{aligned} q_{2^h} &= \sum_{i=1}^{2^h-1} w(h) \times v_i(h) \\ &= w(h) \times \left(\sum_{i=1}^{2^h-1} v_i(h) \right) = w(h) \times (\alpha_h v_0(h) + \beta_h w(h) + \gamma_h v_{2^h}(h)) \\ &= \alpha_h w(h) \times (v_0(h) + v_{2^h}(h)) \end{aligned}$$

where $\alpha_h, \gamma_h \in \mathbb{R}$ and we used for the last equality that $\alpha_h = \gamma_h$ by symmetry. Let $v_\Sigma(h)$ denote the sum $\left(\sum_{i=1}^{2^h-1} v_i(h)\right)$.

We will now calculate α_h. Since $v_1 = (v_0(1)+v_2(1)+w(1))/3$ we obviously have $\alpha_1 = 1/3$. We explain how α_{h+1} can be obtained from α_h. Note that the value of α_{h+1} is not affected by scaling all z-coordinates of the embedding by a factor 3. Therefore, we may work with an embedding M'_{h+1} that is obtained from M_{h+1} by placing all vertices at height $z = 3$.

Every vertex of height h that is adjacent to $w(h+1)$ in M'_{h+1} is in the barycenter of $w(h+1)$ and two vertices that are adjacent to $w(h)$ in M_h. In M'_{h+1} the points of height at most $h-1$ incident to $w(h+1)$ contribute three times as much to $v_\Sigma(h+1)$ as they contribute to $v_\Sigma(h)$. In addition every such vertex contributes twice to $v_\Sigma(h+1)$ via its two neighbors of height h that are also adjacent to $w(h+1)$. This implies that

$$v_\Sigma(h+1) = 5v_\Sigma(h) + v_0(h) + v_{2^h}(h)$$

because $v_0(h)$ and $v_{2^h}(h)$ are contributed additionally through $v_1(h+1)$ and $v_{2^{h+1}-1}(h+1)$, respectively. Since we expressed $v_\Sigma(h+1)$ with respect to M_h we have to scale by a factor $1/3$, and obtain

$$\alpha_{h+1} = \frac{5}{3} \cdot \alpha_h + \frac{1}{3} = \frac{1}{3} \cdot \sum_{i=0}^{h} \left(\frac{5}{3}\right)^i = \frac{1}{2}\left(\left(\frac{5}{3}\right)^{h+1} - 1\right).$$

The last calculation implies that the z-coordinate of $v_{2^h}(h)$ is $(15^h - 9^h)/6$.

$$\langle v_{2^h}, \alpha_h w \times (v_0 + v_{2^h})\rangle = -3^{h-1} \cdot \frac{1}{2}\left(\left(\frac{5}{3}\right)^{h} - 1\right) \cdot \left(-3^h\right) = (15^h - 9^h)/6$$

The claim follows by expressing h in terms of n. □

Corollary 5.11. *Let $\alpha \in (1, 3]$ and P be a stacked polytope such that its edge graph T has height h and $n \geq \alpha^h$ vertices. Then, there is a realization of P with integral vertex coordinates of absolute value bounded by $n^{\log_\alpha 15}$.*

Proof. In the proof of Theorem 5.10 we show that P_h can be realized with coordinates of absolute value less than 15^h. If F is a face of P but not of P_h we truncate it off P_h by intersecting P_h with a halfspace defined by the points representing the vertices of F. Thus, we obtain a realization of P from that of P_h. Expressing 15^h in terms of the number of vertices of P we obtain

$$15^h = \alpha^{h \log_\alpha 15} \leq n^{\log_\alpha 15}.$$

□

Corollary 5.12. *Every stacked polytope has an integer realization with vertex coordinates of absolute value bounded by 15^n.*

Proof. This follows from the fact that $n \geq h + 3$ for every stacked triangulation. □

5.4 Realization of Brooms

As mentioned in the introduction stacked polytopes can be obtained by glueing together simplices or other stacked polytopes. We show a way to gain some control over the coordinates during such a glueing operation. This should be seen as a further step towards a polynomial bound for stacked polytopes which are not covered by the results from Sections 5.2 and 5.3

Theorem 5.13. *A broom T has a realization with integral vertex coordinates of absolute value bounded by $O(n^{7.93})$* *broom realization*

Proof. Let a balanced stacked polytope B_h and a linear stacked polytope L of height h' be given. We glue these polytopes to obtain a polytope P that is a realization of the broom T.

We use the embeddings and realizations of balanced and linear stacked polytopes that we presented in Sections 5.2 and 5.3. We denote the vertices of the outer face Δ of the embedding of B_h by a_1, a_2, a_3. The vertices of the central triangle Δ' of L are denoted by a_1', a_2', a_3'.

We aim to use an affine transformation in order to map the face Δ' to the face Δ and simultaneously stretch L such that the resulting object is convex. In order to guarantee this, we need to consider the faces adjacent to Δ and Δ' in B_h respectively L. Let b_1 be the vertex other than a_1 which forms a face of B_h

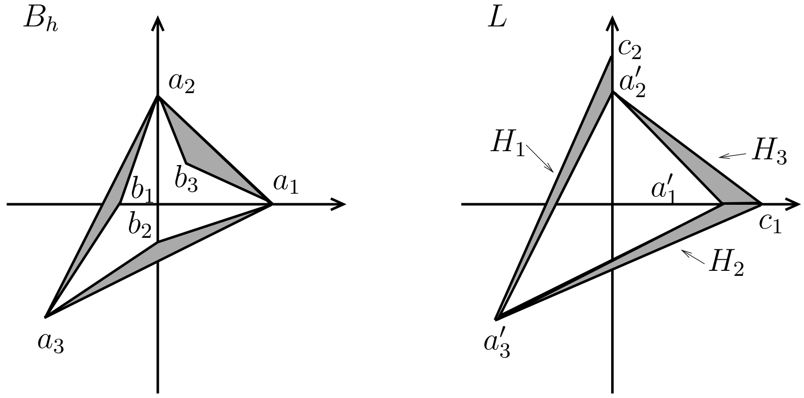

Figure 5.5. Glueing a balanced stacked polytope B_h and a linear stacked polytope L.

with a_2 and a_3, see Figure 5.5. Similarly b_2 and b_3 are defined, and for L we define c_1 and c_2 in this way. We may assume without loss of generality that the first vertex in the shelling order of L is c_1 and the first one on another spine, is c_2, as shown in Figure 5.5. We now describe the details of the realizations of B_h and L that we need.

For the realization of B_h we start with a K_4 embedded with vertex coordinates $(-3,-3,1),(3,0,1),(0,3,1)$ and $(0,0,1)$. We use this embedding to keep the notation simpler. As in the proof of Theorem 5.10 we can calculate the coordinates of vertices of B_h that we need as

$$\begin{aligned}
a_1 &= (3^h,0,0), & b_1 &= (-\tfrac{1}{2}(3^h-3),0,\ell),\\
a_2 &= (0,3^h,0), & b_2 &= (0,-\tfrac{1}{2}(3^h-3),\ell),\\
a_3 &= (-3^h,-3^h,k), & b_3 &= (\tfrac{1}{2}(3^h-3),\tfrac{1}{2}(3^h-3),0),
\end{aligned}$$

with

$$k=\tfrac{3}{2}(15^h-9^h), \quad \text{and} \quad \ell=\tfrac{3}{4}(15^h-9^h)-\tfrac{9}{4}(5^h-3^h).$$

The normal vector of the plane defined by the a_i in B_h is $(k3^h,k3^h,3^{2h+1})$. The coordinates in the lifting of L are

$$\begin{aligned}
a'_1 &= (1,0,0), & c_1 &= (2,0,1),\\
a'_2 &= (0,1,0), & c_2 &= (0,2,\tfrac{i(i+1)}{2}),\\
a'_3 &= (-1,-1,0), & &
\end{aligned}$$

for some $1\le i\le n$. We use the following affine transformation A to glue the two polytopes, by applying it to L, that is $A\cdot a'_i=a_i$. Here, α is the stretching factor that we use to guarantee convexity. We will have to bound α in order prove the claimed bound for the vertex coordinates.

$$
\begin{aligned}
A(x,y,z) &= \begin{pmatrix} 3^h & 0 & \alpha k 3^h \\ 0 & 3^h & \alpha k 3^h \\ -\frac{k}{3} & -\frac{k}{3} & \alpha 3^{2h+1} \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ \frac{k}{3} \end{pmatrix} \\
&= \begin{pmatrix} 3^h x \\ 3^h y \\ \frac{k}{3}(1-x-y) \end{pmatrix} + z\alpha \begin{pmatrix} k3^h \\ k3^h \\ 3^{2h+1} \end{pmatrix}
\end{aligned}
$$

It is easy to check, that A has full rank for $\alpha > 0$ and $k, h \geq 0$. Let H_1 be the plane spanned by $a_2, a_3, A \cdot c_2$, let H_2 be spanned by $a_1, a_3, A \cdot c_1$, and H_3 by $a_1, a_2, A \cdot c_1$, see Figure 5.5. In order to check the convexity conditions we describe the planes H_i in the form $\langle n_i, (x,y,z) \rangle = h_i$ with

$$
\begin{aligned}
n_1 &= \Big(A \cdot c_2 - a_2\Big) \times \Big(a_3 - a_2\Big), \\
n_2 &= \Big(A \cdot c_1 - a_1\Big) \times \Big(a_3 - a_1\Big), \\
n_3 &= \Big(a_2 - a_1\Big) \times \Big(A \cdot c_1 - a_1\Big).
\end{aligned}
$$

To obtain convexity we determine α such that $\langle n_i, b_i \rangle < h_i$ for $i = 1, 2, 3$. We omit the calculations since they are lengthy but straightforward. The calculations imply that it suffices to choose $\alpha > 1$ to satisfy $\langle n_1, b_1 \rangle < h_1$ as well as $\langle n_2, b_2 \rangle < h_2$. The inequality $\langle n_3, b_3 \rangle < h_3$ implies the condition

$$
\alpha > \frac{1}{6}\left(\left(\frac{5}{3}\right)^h - 1\right).
$$

We can thus choose $\alpha = (5/3)^h$ which yields that

$$
A(x,y,z) = \begin{pmatrix} 3^h x + 5^h k z \\ 3^h y + 5^h k z \\ \frac{k}{3}(1-x-y) + 3 \cdot 15^h z \end{pmatrix} = \begin{pmatrix} 3^h x + \frac{3}{2}5^h(15^h - 9^h)z \\ 3^h y + \frac{3}{2}5^h(15^h - 9^h)z \\ \frac{k}{3}(1-x-y) + 3 \cdot 15^h z \end{pmatrix}.
$$

The polytope P has $n = (3^h + 5)/2 + h'$ vertices. Theorem 5.10 implies that the z-coordinates of L are of order $O(n^4)$ and $75^h = O(n^{3.93})$. Therefore we obtain a realization of P of order $O(n^{7.93})$. $\square$

5.5 Conclusions

In this chapter we have shown that linear and balanced stacked triangulations have realizations with integral vertex coordinates of polynomially bounded absolute value. Let $\epsilon > 0$ and consider the set of all stacked triangulations with n

vertices and height h, such that $n \leq (1+\epsilon)^h$. From Corollary 5.11 it follows that if all these stacked triangulations can be lifted with integral vertex coordinates of absolute value bounded by a polynomial in n, then this will imply that all stacked triangulations can be lifted with integral vertex coordinates of absolute value bounded by a polynomial. The main open question that remains is the following.

Problem 5.14. Do all stacked polytopes have realizations with integral vertex coordinates of absolute value bounded by a polynomial?

Conclusions

The four main topics of this thesis are the connections of orthogonal surfaces and Schnyder woods, bounds for the number of planar orientations with prescribed out-degrees, spanning trees with many leaves, and small integer realizations of stacked polytopes. We have given a summary of the main results of each chapter in the Introduction. We now give a collection of interesting open problems from the various chapters. This collection is not complete and more open problems can be found in the concluding sections of the respective chapters.

In Section 1.4 we have introduced the operations edge split and edge merge for Schnyder woods. Besides other applications in [12, 13], these operations have proven to be useful to give a new and simple proof of the Brightwell-Trotter Theorem, see Section 2.1. In the split merge transition graph $\mathcal{S}(n)$ two Schnyder woods are adjacent if they can be obtained from each other by a single split respectively merge operation. In Section 1.4 we have proved a few results about the degrees of the transition graph $\mathcal{S}(n)$. In this context we think that the following question is worth further efforts, see Problem 1.20.

Problem 1. *Can the transition graph $\mathcal{S}(n)$ be used to define a rapidly mixing Markov chain that yields a uniform random sampler for Schnyder woods?*

In Chapter 2 we have studied the connections of orthogonal surfaces and Schnyder woods. In Section 2.3 we have shown how every normalized orthogonal surface $\mathfrak{S}$ can be encoded by a Schnyder wood plus a so-called height value for every minimum and maximum of the surface. The core of the rather complicated proof uses the augmented balance matrix $C'(\mathfrak{S})$. This matrix is invertible and we think that a better understanding of the following problem could help to simplify the proof of Theorem 2.14, see Problem 2.20.

Problem 2. *What is the combinatorial interpretation of the solution of the following linear equation system?*

$$C'(\mathfrak{S}) \cdot y = \vec{e_1}$$

The key for the proof of Theorem 2.14 is a result that Felsner obtained when working on triangle contact representations, see [6] for more on this topic. Progress on Problem 2 could conversely help to answer open questions related to triangle contact representations.

In Chapter 3 we give upper and lower bounds for the maximum number of planar orientations with prescribed out-degrees for different out-degree functions.

In most cases the lower and the upper bound are different, and thus there remain many possibilities for improvement. Among these, the following is particularly interesting, see Problem 3.41.

Problem 3. *Improve the upper bound of 8^n for the number of Schnyder woods of a planar map with n vertices.*

In [29], Páidí Creed shows that counting Eulerian orientations of planar maps is $\#P$-complete. We have shown for some more restricted instances of out-degree functions that counting the number of orientations is $\#P$-complete. For other instances, this question remains open, see Problem 3.44.

Problem 4. *Is it $\#P$-complete to count Eulerian orientations of 4-regular graphs?*

The topic of Chapter 4 are lower bounds for the maximum number of leaves of a spanning tree for a given graph. We give tight lower bounds for some graph classes that are defined by exclusion of certain subgraphs. Planar triangulations are a far more restricted graph class than those that we have considered in Chapter 4. A simple argument using Schnyder woods shows that every stacked triangulation with n vertices has a spanning tree with at least $2n/3 + 1/3$ leaves. This triggered the question if a similar bound can be obtained for all triangulations, see Problem 4.24.

Problem 5. *Does every triangulation on n vertices have a spanning tree with at least $2n/3 + 1/3$ leaves?*

The topic of Chapter 5 are small integer realizations of stacked polytopes. We have shown that some subclasses of stacked polytopes have realizations with polynomially bounded integral coordinates. The question whether this is possible for all stacked polytopes remains open, see Problem 5.14.

Problem 6. *Do all stacked polytopes have realizations with integral vertex coordinates of absolute value bounded by a polynomial?*

This concludes our selection of open problems related to this thesis.

Bibliography

[1] D. M. Acketa and J. D. Žunić, *On the maximal number of edges of convex digital polygons included into a square grid*, Počítače a Umelá Intelegencia, 1 (1982), pp. 549–558. 151

[2] D. M. Acketa and J. D. Žunić, *On the maximal number of edges of convex digital polygons included into an $m \times m$-grid*, J. Comb. Theory Ser. A, 69 (1995), pp. 358–368. 151

[3] D. Adams, *Life, The Universe, and Everything*, The Hitchhiker's Guide to the Galaxy, Pan Books, UK, 1982. iii

[4] D. Adams, *So Long, and Thanks for All the Fish*, The Hitchhiker's Guide to the Galaxy, Pan Books, UK, 1984. iii

[5] N. Alon, *Transversal numbers of uniform hypergraphs*, Graphs and Combinatorics, 6 (1990), pp. 1–4. 104

[6] M. Badent, C. Binucci, E. D. Giacomo, W. Didimo, S. Felsner, F. Giordano, J. Kratochvil, P. Palladino, M. Patrignani, and F. Trotta, *Homothetic triangle contact representations of planar graphs*, in Proc. 19th Canad. Conf. on Comp. Geom., 2007, pp. 233–236. 39, 51, 167

[7] I. Bárány and G. Rote, *Strictly convex drawings of planar graphs*, Documenta Mathematica, 11 (2006), pp. 369–391. 1

[8] R. J. Baxter, *F model on a triangular lattice*, J. Math. Physics, 10 (1969), pp. 1211–1216. 54, 60, 61, 100

[9] R. J. Baxter, *Exactly solved models in statistical mechanics*, Academic Press, 1982. 53, 58

[10] A. I. Bobenko and B. A. Springborn, *Variational principles for circle patterns, and Koebe's theorem*, Transactions Amer. Math. Soc., 356 (2004), pp. 659–689. 152

[11] H. L. Bodlaender, *On linear time minor tests with depth-first search*, J. Algorithms, 14 (1993), pp. 1–23. 101, 105

[12] N. Bonichon, *A bijection between realizers of maximal plane graphs and pairs of non-crossing dyck paths*, Discr. Math., 298 (2005), pp. 104–114. 62, 167

[13] N. Bonichon, S. Felsner, and M. Mosbah, *Convex drawings of 3-connected planar graphs*, Algorithmica, 47 (2007), pp. 399–420. 1, 12, 55, 167

[14] N. Bonichon, B. Le Saëc, and M. Mosbah, *Wagner's theorem on realizers*, in Proc. 29th Int. Col. on Autom., Lang., and Prog., ICALP 02, vol. 2380 of LNCS, 2002, pp. 1043–1053. 14

[15] P. S. Bonsma, *Sparse cuts, matching-cuts and leafy trees in graphs*, PhD thesis, University of Twente, Enschede, the Netherlands, 2006. 101, 102, 103, 104, 105, 113, 118, 119, 149

[16] P. S. Bonsma, T. Brueggemann, and G. J. Woeginger, *A faster FPT algorithm for finding spanning trees with many leaves*, in Proc. 28th Int. Symp. Math. Found. Computer Science, MFCS 03, vol. 2747 of LNCS, 2003, pp. 259–268. 101, 105

[17] P. S. Bonsma and F. Zickfeld, *Spanning trees with many leaves in graphs without diamonds and blossoms*. arXiv: 0707.2760v1, 2007. Accepted for 8th Latin American Theoretical Informatics, 2008. viii, 105

[18] G. R. Brightwell, *Personal communication with S. Felsner*, 2006. 77

[19] G. R. Brightwell and W. T. Trotter, *The order dimension of convex polytopes*, SIAM J. Discrete Math., 6 (1993), pp. 230–245. 1, 24, 25

[20] G. R. Brightwell and W. T. Trotter, *The order dimension of planar maps*, SIAM J. Discrete Math., 10 (1997), pp. 515–528. 1, 24

[21] N. J. Calkin and H. S. Wilf, *The number of independent sets in a grid graph*, SIAM J. Discrete Math., 11 (1998), pp. 54–60. 58, 69, 71

[22] Y. Caro, D. B. West, and R. Yuster, *Connected domination and spanning trees with many leaves*, SIAM J. Discrete Math., 13 (2000), pp. 202–211 (electronic). 104, 113

[23] Y. Chiang, C. Lin, and H. Lu, *Orderly spanning trees with applications*, SIAM J. Comput., 34 (2005), pp. 924–945. 29

[24] M. Chrobak and G. Kant, *Convex grid drawings of 3-connected planar graphs*, Internat. J. Comput. Geom. Appl., 7 (1997), pp. 211–223. 29, 153

[25] R. Chuang, A. Garg, X. He, M. Kao, and H. Lu, *Compact encodings of planar graphs via canonical orderings and multiple parentheses*, in Proc. 25th Int. Col. on Autom., Lang., and Prog., ICALP 98, vol. 1443 of LNCS, 1998, pp. 118–129. 29

[26] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*, MIT Press, Cambridge, MA, second ed., 2001. 36

[27] J. R. Correa, C. Fernandes, M. Matamala, and Y. Wakabayashi, *A 5/3-approximation for finding spanning trees with many leaves in cubic graphs*, in 5th W. Approximation and Online Algorithms, 2007. 101

[28] H. Crapo and W. Whiteley, *Plane self stresses and projected polyhedra I: the basic pattern*, Structural Topology, 20 (1993), pp. 55–78. 153

[29] P. Creed, *Sampling Eulerian orientations of triangular lattice graphs*, 2007. arXiv:cs/0703031v1. 92, 93, 168

[30] P. Dagum and M. Luby, *Approximating the permanent of graphs with large factors*, Theoretical Computer Science, 102 (1992), pp. 283–305. 94, 96

[31] H. de Fraysseix and P. O. de Mendez, *On topological aspects of orientations*, Discr. Math., 229 (2001), pp. 57–72. 9, 53

[32] H. de Fraysseix, P. O. de Mendez, and P. Rosenstiehl, *Bipolar orientations revisited*, Discr. Appl. Math., 56 (1995), pp. 157–179. 73, 74

[33] R. G. Downey and M. R. Fellows, *Parameterized computational feasibility*, in Feasible mathematics, II (1992), vol. 13 of Progr. Comput. Sci. Appl. Logic, Birkhäuser Boston, 1995, pp. 219–244. 105

[34] R. G. Downey and M. R. Fellows, *Parameterized complexity*, Springer-Verlag, New York, 1999. 105

[35] V. Estivill-Castro, M. R. Fellows, M. A. Langston, and F. A. Rosamond, *FPT is P-time extremal structure I*, in ACiD 2005, vol. 4 of Texts in algorithmics, King's College Publications, 2005, pp. 1–41. 101, 105

[36] M. R. Fellows, C. McCartin, F. A. Rosamond, and U. Stege, *Coordinatized kernels and catalytic reductions: an improved FPT algorithm for max leaf spanning tree and other problems*, in In Proc. 20th Conf. Found. Software Technology and Theoretical Computer Science, FSTTCS 2000, vol. 1974 of LNCS, 2000, pp. 240–251. 105

[37] S. Felsner. http://www.math.tu-berlin.de/~felsner/Schnyder.bib. 1

[38] S. Felsner, *Convex drawings of planar graphs and the order dimension of 3-polytopes*, Order, 18 (2001), pp. 19–37. 1, 3, 24, 25, 153

[39] S. Felsner, *Geodesic embeddings and planar graphs*, Order, 20 (2003), pp. 135–150. 1, 23, 25

[40] S. Felsner, *Geometric Graphs and Arrangements*, Vieweg Verlag, 2004. 1, 2, 3, 7

[41] S. Felsner, *Lattice structures from planar graphs*, Elec. J. Comb., (2004). R15. 1, 9, 11, 53, 55

[42] S. FELSNER, E. FUSY, M. NOY, AND D. ORDEN, *Baxter families and more: Bijections and counting*, 2007. in preparation. 69

[43] S. FELSNER, C. HUEMER, S. KAPPES, AND D. ORDEN, *Binary labelings for plane quadrangulations and their relatives*. arXiv:math.CO/0612021, 2007. Submitted. 69

[44] S. FELSNER AND S. KAPPES, *Orthogonal surfaces*. arXiv: math.CO/0602063, 2006. Submitted. 6

[45] S. FELSNER AND F. ZICKFELD, *Schnyder woods and orthogonal surfaces*, in Proc. 14th Int. Symp. Graph Drawing, GD 06, vol. 4372 of LNCS, 2006, pp. 417–429. vii

[46] S. FELSNER AND F. ZICKFELD, *On the number of α-orientations*, in Proc. 33rd Int. W. Graph Theoretic Concepts in Computer Science, WG 07, vol. 4769 of LNCS, 2007, pp. 190–201. vii

[47] S. FELSNER AND F. ZICKFELD, *On the number planar orientations with prescribed degrees*. arXiv: math.CO/0701771v2, 2007. Submitted. vii

[48] S. FELSNER AND F. ZICKFELD, *Schnyder woods and orthogonal surfaces*, Discrete and Computational Geometry, (2007). DOI 10.1007/s00454-007-9027-9. vii

[49] J. FLUM AND M. GROHE, *Parameterized complexity theory*, Springer, Berlin, 2006. 105

[50] E. FUSY, *Transversal structures on triangulations, with application to straight-line drawing.*, in Proc. 13th Int. Symp. Graph Drawing, GD 05, vol. 3843 of LNCS, 2005, pp. 177–188. 55

[51] E. FUSY, *Combinatorics of Plane Maps with Algorithmic Applications*, PhD thesis, École Polytechnique, 2007. 69

[52] E. FUSY, D. POULALHON, AND G. SCHAEFFER, *Dissections and trees, with applications to optimal mesh encoding and to random sampling*, in Proc. 16th ACM-SIAM Sympos. Discrete Algorithms, SODA 05, 2005, pp. 690–699. 29, 55

[53] M. R. GAREY AND D. S. JOHNSON, *Computers and intractability*, Freeman, San Francisco, 1979. 101

[54] J. R. GRIGGS, D. J. KLEITMAN, AND A. SHASTRI, *Spanning trees with many leaves in cubic graphs*, J. Graph Theory, 13 (1989), pp. 669–695. 101, 102, 103, 104, 108, 109, 111, 119, 124, 125, 126, 128, 129, 131, 132, 149

[55] J. R. GRIGGS AND M. WU, *Spanning trees in graphs of minimum degree 4 or 5*, Discr. Math., 104 (1992), pp. 167–183. 104

[56] T. W. Haynes, S. T. Hedetniemi, and P. J. Slater, *Fundamentals of domination in graphs*, vol. 208 of Monographs and Textbooks in Pure and Applied Mathematics, Marcel Dekker Inc., New York, 1998. 113

[57] R. A. Horn and C. R. Johnson, *Matrix analysis*, Cambridge University Press, Cambridge, 1990. Corrected reprint of the 1985 original. 71

[58] M. Jerrum, A. Sinclair, and E. Vigoda, *A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries*, J. ACM, 51 (2004), pp. 671–697. 98

[59] G. Kant, *Drawing planar graphs using the canonical ordering*, Algorithmica, 16 (1996), pp. 4–32. 29

[60] S. Kappes, *Orthogonal Surfaces - A Combinatorial Approach*, PhD thesis, Technische Universität Berlin, 2007. 6

[61] P. W. Kasteleyn, *Graph theory and crystal physics*, in Graph Theory and Theoretical Physics, Academic Press, London, 1967, pp. 43–110. 53

[62] R. W. Kenyon, J. G. Propp, and D. B. Wilson, *Trees and matchings*, Elec. J. Comb., 7 (2000). 65

[63] D. J. Kleitman and D. B. West, *Spanning trees with many leaves*, SIAM J. Discrete Math., 4 (1991), pp. 99–106. 101, 102, 103, 104, 105, 108, 109, 112

[64] K. B. Knauer, *Partial orders on orientations via cycle flips*, master's thesis, Technische Universität Berlin, 2007. 55

[65] E. H. Lieb, *The residual entropy of square ice*, Physical Review, 162 (1967), pp. 162–172. 53, 58, 59, 60, 71, 76, 78, 100

[66] C. Lin, H. Lu, and I. Sun, *Improved compact visibility representation of planar graphs via Schnyder's realizer*, SIAM J. Discrete Math., 18 (2004), pp. 19–29. 1

[67] N. Linial and D. G. Sturtevant, 1987. Unpublished result. 102

[68] C. H. C. Little, *A characterization of convertible* $(0,1)$*-matrices*, J. Combin. Theory Ser. B, 18 (1975), pp. 187–208. 98

[69] L. Lovász and M. D. Plummer, *Matching Theory*, no. 29 in Annals of Discrete Mathematics, Akadémiai Kiadó - North Holland, 1986. 53, 65, 98

[70] S. Melang, *Bipolare Orientierungen planarer Graphen*, master's thesis, Technische Universität Berlin, 2006. 83

[71] M. Mihail and P. Winkler, *On the number of Eulerian orientations of a graph*, Algorithmica, 16 (1996), pp. 402–424. 92, 94

[72] E. N. Miller, *Planar graphs as minimal resolutions of trivariate monomial ideals*, Documenta Mathematica, 7 (2002), pp. 43–90. 7, 25

[73] A. R. Mor, *Realization and Counting Problems for Planar Structures: Trees, Linkages, Polytopes and Polyominoes*, PhD thesis, Freie Universität Berlin, 2005. 53, 151

[74] A. R. Mor, G. Rote, and A. Schulz, *Embedding 3-polytopes on a small grid*, in Proc. of the 23rd Ann. Symp. on Comput. Geom., 2007, pp. 112–118. 151

[75] D. Poulalhon and G. Schaeffer, *Optimal coding and sampling of triangulations*, Algorithmica, 46 (2006), pp. 505–527. 62

[76] F. P. Preparata and S. J. Hong, *Convex hulls of finite sets of points in two and three dimensions*, Comm. ACM, 20 (1977), pp. 87–93. 29

[77] F. P. Preparata and M. I. Shamos, *Computational Geometry*, Texts and Monographs in Computer Science, Springer-Verlag, New York, 1985. An introduction. 29

[78] J. Richter-Gebert, *Realization Spaces of Polytopes*, vol. 1643 of Springer Lecture Notes, Springer, 1996. 152, 159, 160

[79] N. Robertson, P. D. Seymour, and R. Thomas, *Permanents, Pfaffian orientations, and even directed circuits*, Ann. of Math. (2), 150 (1999), pp. 929–975. 98, 99

[80] P. Rosenstiehl, *Embedding in the plane with orientation constraints: the angle graph*, Ann. New York Acad. Sci., (1983), pp. 340–346. 74

[81] G. Rote, *The number of spanning trees in a planar graph*, in Oberwolfach Reports, EMS, 2005, pp. 969–973. `http://page.mi.fu-berlin.de/rote/about_me/publications.html`. 53

[82] W. Schnyder, *Planar graphs and poset dimension*, Order, 5 (1989), pp. 323–343. 1, 24

[83] W. Schnyder, *Embedding planar graphs on the grid*, in Proc. 1st ACM-SIAM Sympos. Discrete Algorithms, SODA 90, 1990, pp. 138–148. 1, 153

[84] N. J. A. Sloane, *The on-line encyclopedia of integer sequences*. `http://www.research.att.com/~njas/sequences`. 79

[85] R. Solis-Oba, 2-*approximation algorithm for finding a spanning tree with maximum number of leaves*, in Proc. 6th Ann. European Symp. Algorithms, ESA 98, vol. 1461 of LNCS, 1998, pp. 441–452. 101

[86] E. STEINITZ, *Polyeder und Raumeinteilungen*, in Encyklopädie der mathematischen Wissenschaften, mit Einschluss ihrer Anwendungen, Dritter Band: Geometrie, III.1.2., Heft 9, Kapitel III A B 12, W. F. Meyer and H. Mohrmann, eds., B. G. Teubner, Leipzig, 1922, pp. 1–139. 24, 151

[87] E. STEINITZ AND H. RADEMACHER, *Vorlesungen über die Theorie der Polyeder*, Springer Verlag, Berlin, 1934. Reprint, Springer Verlag 1976. 24, 151

[88] R. TAMASSIA AND I. G. TOLLIS, *A unified approach to visibility representations of planar graphs*, Discrete and Computational Geometry, 1 (1986), pp. 321–341. 53, 74

[89] H. N. V. TEMPERLEY, *Enumeration of graphs on a large periodic lattice*, in Combinatorics: Proceedings of the British Combinatorial Conference 1973, London Math. Soc. Lecture Note Series ♯13, 1974, pp. 155–159. 66

[90] T. THIELE, *Extremale Probleme für Punktmengen*, master's thesis, Freie Universität Berlin, 1991. 151

[91] W. T. TROTTER, *Combinatorics and Partially Ordered Sets: Dimension Theory*, The Johns Hopkins University Press, 1992. 24

[92] W. T. TROTTER, *Partially ordered sets*, in Handbook of Combinatorics, L. Graham, Grötschel, ed., Elsevier, Amsterdam, 1995, pp. 433–480. 24

[93] W. TUTTE, *How to draw a graph*, Proc. London Math. Soc., 13 (1963), pp. 743–767. 152

[94] W. T. TUTTE, *A short proof of the factor theorem for finite graphs*, Canadian J. Mathematics, 6 (1954), pp. 347–352. 98

[95] W. T. TUTTE, *A census of planar triangulations*, Canadian J. Mathematics, 14 (1962), pp. 21–38. 62

[96] S. P. VADHAN, *The complexity of counting in sparse, regular, and planar graphs*, SIAM J. Comput., 31 (2001), pp. 398–427. 93

[97] K. WAGNER, *Bemerkungen zum Vierfarbenproblem*, in Jahresbericht Deutsche Math.-Vereinigung, vol. 46, 1936, pp. 26–32. 14

[98] D. R. WOODS, *Drawing Planar Graphs*, PhD thesis, Stanford University, 1982. Technical Report STAN-CS-82-943. 53, 74

[99] G. M. ZIEGLER, *Lectures on Polytopes*, Springer, New York, 1995. 152

[100] G. M. ZIEGLER, *Convex polytopes: Extremal constructions and f-vector shapes*. Park City Mathematical Institute (PCMI 2004) Lecture Notes (E. Miller, V. Reiner, B. Sturmfels, eds.). With an Appendix by Th. Schröder and N. Witte; Preprint, TU Berlin, 73 pages; `http://www.arXiv.org/math.MG/0411400`, Nov. 2004. 152

Symbol Index

Index