

Schnyder Woods and Orthogonal Surfaces

Stefan Felsner

Florian Zickfeld

Technische Universität Berlin, Fachbereich Mathematik

Straße des 17. Juni 136, 10623 Berlin, Germany

{felsner,zickfeld}@math.tu-berlin.de

Abstract

In this paper we study connections between planar graphs, Schnyder woods, and orthogonal surfaces. Schnyder woods and the face counting approach have important applications in graph drawing and the dimension theory of orders. Orthogonal surfaces explain connections between these seemingly unrelated notions. We use these connections for an intuitive proof of the Brightwell-Trotter Theorem which says, that the face lattice of a 3-polytope minus one face has order dimension three. Our proof yields a linear time algorithm for the construction of the three linear orders that realize the face lattice.

Coplanar orthogonal surfaces are in correspondence with a large class of convex straight line drawings of 3-connected planar graphs. We show that Schnyder's face counting approach with weighted faces can be used to construct all coplanar orthogonal surfaces and hence the corresponding drawings. Appropriate weights are computable in linear time.

1 Introduction

In two fundamental papers [17, 18] Schnyder developed a theory of Schnyder labelings and Schnyder woods for planar triangulations. The second paper deals with grid drawings of planar graphs and contains the first of numerous applications of Schnyder woods in the area of graph drawing. For example, the results in [3, 15, 2] use Schnyder woods, and more references can be found in [8].

In [17], 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 $G = (V, E)$, associate an order P_G of height two 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 .

The *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*. 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 [19, 20, 4, 9].

Theorem 1 (Schnyder's Theorem) *A graph is planar if and only if the dimension of its incidence order is at most three.*

In the same paper Schnyder also shows that the incidence poset of vertices, edges and faces of a planar triangulation has dimension four, but the dimension drops to three upon removal of a face. Brightwell and Trotter [5] 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.

Theorem 2 (Brightwell-Trotter Theorem) *The incidence order of the vertices, edges and faces of a 3-connected planar graph G has dimension four. 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 three.*

Note that, by Steinitz's Theorem, the incidence poset of vertices, edges and faces of a 3-connected planar graph is just the face lattice of the corresponding 3-polytope with $\mathbf{0}$ and $\mathbf{1}$ removed.

The original proof of Theorem 2 in [4] was long and technical. Felsner [9] gave a simpler proof. In [10] Felsner showed that every Schnyder wood of a 3-connected planar graph is supported by a rigid orthogonal surface (Theorem 7). An orthogonal surface is called rigid if it supports a unique graph, see Figure 10 (b). By a result of Miller [16], Felsner's result implies Theorem 2. In Section 3 we present an intuitive proof of Theorem 7, that 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. 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 12. The nontrivial point is to show that these adjustments can be combined in such a way that the whole surface becomes rigid.

The rest of the paper is organized as follows. In Section 2 we give definitions and a brief introduction into the structural properties of Schnyder woods and orthogonal surfaces which are required for the discussion in the later parts of this paper. For a more detailed introduction to the topic we refer the reader to [11].

As mentioned above, Section 3 deals with rigid orthogonal surfaces. An orthogonal surface is coplanar if all its generating minima, i.e. the points on the surface which are minimal with respect to the dominance order in \mathbb{R}^3 , lie in the same plane. Section 4.1 is concerned with these coplanar surfaces. The interest in this class originates from their close connection to planar straight line drawings. Connecting the minima of a coplanar surface by straight line segments yields a plane and convex straight line drawing of the graph. Similar approaches for non-coplanar surfaces fail as the drawings need not be crossing-free. We show 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 which is simultaneously rigid and coplanar is the topic of Section 4.2. In Section 5 we discuss further representations of orthogonal surfaces and mention some open problems.

2 Basics on Schnyder Woods and Orthogonal Surfaces

All the proofs omitted in this section can be found in [11, 10, 9]. A *planar map* M is a simple planar graph G together with a fixed planar embedding of G in the plane. Let a_1, a_2, a_3 be three vertices occurring in clockwise order on the outer face of M . A suspension M^σ is obtained by attaching a half-edge that reaches into the outer face to each of these *special vertices*.

Let M^σ be a suspended 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 (alternatively: red, green, blue) satisfying three rules¹:

- (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) *Rule of vertices*: The labels of the angles at each vertex form, in clockwise order, nonempty intervals of 1's, 2's, and 3's.
- (A3) *Rule of faces*: The labels of the angles at each face form, in clockwise order, a nonempty interval of 1's, a nonempty interval of 2's and a nonempty interval of 3's. At the outer face the same is true in counterclockwise order.




Figure 1: Rule of vertices and rule of faces

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 satisfying the following rules.

- (W1) Every edge e is oriented in one direction or in two opposite directions. The directions of edges are colored such that if e is bidirected the two directions have distinct colors.
- (W2) The half-edge at a_i is directed outwards and colored i .
- (W3) Every vertex v has out-degree one in each color. The edges e_1, e_2, e_3 leaving v in colors 1, 2, 3 occur in clockwise order. Each edge entering v in color i enters v in the clockwise sector from e_{i+1} to e_{i-1} .




Figure 2: Edge orientations and edge colors at a vertex.

- (W4) There is no interior face the boundary of which is a directed monochromatic cycle.

The next theorem shows that Schnyder labelings and Schnyder woods are, essentially, the same.

¹We assume a cyclic structure on the labels so that $i + 1$ and $i - 1$ are always defined.

Theorem 3 *Let M^σ be a suspended 3-connected planar map. The correspondence indicated in Figure 3 is a bijection between the Schnyder labelings (axioms A1, A2, A3) and Schnyder woods (axioms W1, W2, W3, W4) of M^σ .*




Figure 3: The correspondence between angle labels at an edge and the coloring and 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 a suspension explicitly.

Let M 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 one in T_i . Therefore, every 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.

Lemma 1 *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 . The digraph $D_i = T_i \cup T_{i-1}^{-1} \cup T_{i+1}^{-1}$ is acyclic for $i = 1, 2, 3$.*

A proof can be found in [9] or [11].

By the rule of vertices (W3) every vertex has out-degree one 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:

Corollary 1 *T_i is a directed 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 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 M into 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 4.

Lemma 2 *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 3 *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.*

These lemmas are crucial for the applications of the *face-count vector* (v_1, v_2, v_3) of a vertex v with respect to a Schnyder wood which is defined by

$$v_i = \text{the number of faces of } M \text{ contained in region } R_i(v).$$




Figure 4: A Schnyder wood and the regions of the vertex v . The small numbers correspond to edge colors.

Later we will use this vector to construct orthogonal surfaces supporting a given Schnyder wood. The classic application 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

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

and an edge $\{u, v\}$ is mapped by μ to the line segment connecting $\mu(u)$ and $\mu(v)$.

Theorem 4 *The drawing $\mu(M)$ of a 3-connected plane map is convex and plane, i.e., the boundary of every face is a convex polygon.*

We will need another tool from the theory of Schnyder woods, the *edge split*. We start with a lemma from [3] about the generic appearance of a face in a Schnyder wood.

Lemma 4 *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 5, 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

The three edges from the first, third, and fifth set are the special edges of the face.




Figure 5: The generic appearance of a face as described by Lemma 4 and two concrete instances. The small numbers correspond to edge colors.

Sketch of the proof. Recall the rule of faces (A3) for the Schnyder labeling. Apply the rule (Figure 3) for converting a Schnyder labeling into a Schnyder wood. \square

Given a Schnyder wood S let e be a bidirected edge such that one of its directions is colored j and F be the incident face to which 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 6 illustrates the operation.




Figure 6: The two possible types of splits of a non-special bidirected red-green edge in F . The small numbers correspond to angle labels.

Lemma 5 *Let S be a Schnyder wood and e a bidirected edge of S . Then, splitting e yields a Schnyder wood on the resulting graph.*

Proof. Figure 6 shows the angle labelings. It is obvious that the labels at the angles of u , v , w , F_1 and F_2 obey the rule of vertices (A2) and of faces (A3), respectively. Note that (w, v) and (w', u) may also be special edges of F . \square

2.1 Orthogonal surfaces

Consider \mathbb{R}^3 equipped with the dominance order. We write $u \vee v$ and $u \wedge v$ to denote the *join* (component-wise maximum) and *meet* (component-wise minimum) of $u, v \in \mathbb{R}^3$. Let $\mathcal{V} \subset \mathbb{R}^3$ be an antichain, i.e., a set of pairwise incomparable elements. The *filter* generated by \mathcal{V} in \mathbb{R}^3 is the set

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




Figure 7: Part (a): orthogonal surface $\mathfrak{S}_{\mathcal{V}}$ 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 $v_9 = (1, 2, 4)$. Part (b): surface generated by $u_1 = (5, 0, 0)$, $u_2 = (0, 5, 0)$, $u_3 = (0, 0, 5)$, $u_4 = (4, 3, 2)$, $u_5 = (4, 4, 1)$.

The boundary $\mathfrak{S}_{\mathcal{V}}$ of $\langle \mathcal{V} \rangle$ is the *orthogonal surface* generated by \mathcal{V} , Figures 7 (a) and (b) show examples.

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* in $\mathfrak{S}_{\mathcal{V}}$. The *orthogonal arc* of $v \in \mathcal{V}$ in direction of the standard basis vector e_i is the piece of the ray $v + \lambda e_i$, $\lambda \geq 0$, which follows a crease of $\mathfrak{S}_{\mathcal{V}}$. Clearly every vector $v \in \mathcal{V}$ has exactly three orthogonal arcs, one parallel to each coordinate axis. Some orthogonal arcs are unbounded while others are bounded. Observe that $u \vee v$ shares two coordinates with at least one (and perhaps both) of u and v , so every elbow geodesic contains at least one bounded orthogonal arc.

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) *Vertex axiom*. There is a bijection between the vertices of M and \mathcal{V} .
- (G2) *Elbow geodesic axiom*. Every edge of M is an elbow geodesic in $\mathfrak{S}_{\mathcal{V}}$, and every bounded orthogonal arc in $\mathfrak{S}_{\mathcal{V}}$ is part of an edge of M .
- (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 contains exactly three unbounded orthogonal arcs. The example from Figure 7 (a) is not axial, however, removing the point v_5 from the set \mathcal{V} leads to an axial surface, see Figure 8. These definitions have been proposed by Miller [16] who, essentially, also observed the following theorem.

Theorem 5 *Let \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 , which 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 .*

Sketch of the proof. 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: 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 8 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 embed every vertex v at its face count vector $(v_1, v_2, v_3) \in \mathbb{N}^3 \subset \mathbb{R}^3$, i.e., $\mathcal{V} = \{(v_1, v_2, v_3) : v \text{ is a vertex of } M\}$. It can be verified that the canonical map $M \hookrightarrow \mathfrak{S}_{\mathcal{V}}$ is a geodesic embedding. The orthogonal surface in the left part of Figure 8 can be constructed by this rule from the Schnyder wood on the right. The complete proof of the theorem can be found in [11]. \square




Figure 8: A geodesic embedding and the induced Schnyder labeling and Schnyder wood.

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*. Flats are the connected regions of constant gray-value in our drawings of orthogonal surfaces. Formally, let H be the plane $x_i = h$ and $\tilde{F}_1, \dots, \tilde{F}_\ell$, the connected components of the interior of $H \cap \mathfrak{S}_{\mathcal{V}}$. The topological closures F_1, \dots, F_ℓ of these components are i -flats of height h , see Figure 9. The i -flat of $v \in V$ is denoted by $F_i(v)$. A more technical definition of flats is given in [12], the definition there is used for investigations of orthogonal surfaces in dimension 4 and higher.

In the Schnyder labeling the angle φ is labeled i , if the x_i -axis is orthogonal to the flat containing it (see the labels in Figure 8). It is easy to verify properties (A1), (A2) and (A3) for this angle labeling.




Figure 9: Two i -flats with the same i -coordinate.

Given Theorem 5 it is natural to ask questions about *existence* and *uniqueness* of geodesic embeddings.

Question 1 Does every axial orthogonal surface $\mathfrak{S}_{\mathcal{V}}$ support a Schnyder wood?

Answer: No, a surface with three orthogonal arcs meeting in a single point does not, see Figure 10.a. Call a surface *degenerate* if such a pattern occurs. It can be shown that every non-degenerate axial orthogonal surface \mathfrak{S}_V supports a Schnyder wood.

Convention: From now on this paper only deals with *non-degenerate* and *axial* orthogonal surfaces. For brevity we will usually omit these predicates.




Figure 10: **a.** A degenerate pattern. **b.** A non-rigid edge (u, v) , the bend-point $u \wedge v$ dominates w .

Question 2 Is the Schnyder wood supported by the surface \mathfrak{S}_V unique?

Answer: In general not, e.g. in the situation shown in Figure 10.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 existence of such a choice for an edge is caused by a non-rigidity in the sense of the following definition:

An elbow geodesic connecting vertices u and v is *rigid*, if u and v are the only vertices in \mathcal{V} dominated by $u \vee v$. An orthogonal surface \mathfrak{S}_V is a *rigid surface* if all its elbow geodesics are rigid.

The next section is devoted to the proof of a stronger version of Theorem 5, in the sense that for every Schnyder wood there is a rigid surface supporting it.

3 Rigid Orthogonal Surfaces via Flat Shifting

We set Theorem 7 into context before we give a new proof. Miller [16] observed that a rigid orthogonal surface supports exactly one Schnyder wood and proved:

Theorem 6 *Every suspended 3-connected planar map M^σ has a geodesic embedding $M^\sigma \hookrightarrow \mathfrak{S}$ on some rigid orthogonal surface \mathfrak{S} .*

Together with the following proposition from [16] (see also [10]) this implies the Brightwell-Trotter Theorem (Theorem 2).

Proposition 1 *Let \mathfrak{S}_V be a rigid orthogonal surface. Let $M^\sigma \hookrightarrow \mathfrak{S}_V$ be a geodesic embedding and F a bounded region of M . If α_F is the join of the vertices of F , then $w \in F \Leftrightarrow w \leq \alpha_F$.*

Note that α_F as defined above lies on \mathfrak{S}_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}_V if and only if they all lie on a common face of M^σ . It is crucial here, that \mathfrak{S}_V is a rigid surface. If W contains a vertex from each of the three sides of the face F , as shown in Figure 5, then the join is a maximum of \mathfrak{S}_V .

We give a new proof of the following result by Felsner [10], who answered a question by Miller with this extension of Theorem 6.

Theorem 7 *If S is a Schnyder wood of a map M^σ , then there is a rigid axial orthogonal surface \mathfrak{S} and a geodesic embedding $M^\sigma \hookrightarrow \mathfrak{S}$. In particular S is the unique Schnyder wood supported by \mathfrak{S} .*

We now present a new proof of this theorem, Lemma 6 and Lemma 7 are part of this proof. 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. 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 11 shows an example.

- (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 $w \in V$ is such that $F_j(w) = F_j(u)$ and $w_i > u_i$, then $F_i(v) < F_i(w)$ in Γ_i .

The pairs in Γ_i are classified as a -relations (*arc*), p -relations (*preserve*) and r -relations (*repel*). Lemma 6 is the heart of the proof of Theorem 7 as it justifies why the flat shifts (i.e. r -relations) can be combined to obtain a rigid surface.




Figure 11: On the left a non-rigid surface with Schnyder wood S . On the right the corresponding relation Γ_1 . Here, a -relations are represented by red arrows, p -relations by cyan arrows, and the only r -relation by a golden arrow.

Lemma 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 $i = 1$.

We identify the a - and p -relations with edges of the Schnyder wood S . The set of vertices lying on a common 1-flat is strongly connected in S via bidirected green-blue edges. 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 page 6). Let $e = (v, u) \in S$ be a unidirected blue edge and $F(u) < F(v)$ the

corresponding p -relation. Let $F(u_k) > \dots > F(u_1)$ be the set of flats that have an r -relation $F(v) < F(u_j)$ related to e , the order on this set coming from the a -relations. 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 $\{u, u_1\}, \{u_1, u_2\}, \dots, \{u_{k-1}, u_k\}$ towards v . This is legal since the angle of v in the face in question has label 2 (green), see Lemma 5.

Repeat this operation for other r -relations in Γ_1 which come from unidirected blue edges. A symmetric operation is used to introduce edges for all r -relations in Γ_1 which come from unidirected green edges in the Schnyder wood S .

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

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 $v \neq v'$ from the flat $F(v)$. This yields gaps in the intended cycle C' . In this case 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. □




Figure 12: An orthogonal surface and an associated rigid surface

Let \mathfrak{S} be the orthogonal surface supporting S which is generated by the face counting vectors (c.f. Theorem 5). Let Γ_i^* be the transitive closure of Γ_i which is an order on \mathcal{F}_i by Lemma 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 . In the remaining part of this section we 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 to the rest of the proof of Theorem 7 is the following lemma.

Lemma 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 $P_j(v)$. The complete information for $e = \{x, y\}$ is one of the following:

- (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 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 $u'_i < v'_i$. \square

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

Proof. This follows from Lemma 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 $R_j(v) \subset R_j(u)$. \triangle

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 10.a cannot occur. \triangle

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 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 $[u, v]$ is on the surface $\mathfrak{S}_{\mathcal{V}_\alpha}$.

If edge $e = (u, v)$ is directed in color i from u to v then by Lemma 3 together with Lemma 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} . \triangle

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

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

This completes the proof of Theorem 7. \square

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

Corollary 2 *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. Produce a copy of the vertex set. On this copy build a digraph D_r : 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. 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 as other repel relations associated to the same unidirected blue edge will be 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, 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 7. \square

Theorem 8 *Let P be a 3-polytope with n vertices. Then, a Brightwell-Trotter realizer for P can be computed in $O(n)$ time.*

Proof. Let G be the edge graph of P . As shown by Fusy et al. [13] a Schnyder wood S for G can be computed in $O(n)$ time. With little translational effort this also follows from algorithms for computing orderly spanning trees [6] or canonical orderings [7] which are based on Kant's algorithm [14]. By Corollary 2 a rigid orthogonal surface that induces S can be computed in time $O(n)$. By Proposition 1 such an orthogonal surface yields a Brightwell-Trotter realizer of P . \square

4 Coplanar Surfaces

An orthogonal surface is called *coplanar*, if there exists a constant $c \in \mathbb{R}$ such that every minimum v on the surface fulfills $v_1 + v_2 + v_3 = c$. Schnyder's classic approach of drawing graphs using the face-count 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 5. Geodesic embeddings on coplanar surfaces have the pleasant property that the positions of the vertices in the plane yield a crossing-free and convex straight-line drawing of the underlying graph.

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 7 (b) for an example. The coordinates of the orthogonally projected points are

$$u'_1 = (11, -4, -4)/3, u'_2 = (-4, 11, -4)/3, u'_4 = (4, 1, -2)/3, u'_5 = (4, 4, -5)/3.$$

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 in terms of Schnyder woods is given in Section 4.1. In Section 4.2 we present an example of a Schnyder wood for that no surface inducing it can be rigid and simultaneously coplanar.

4.1 Coplanar Surfaces and Face Weights

We now generalize the classic approach of counting every bounded face with weight one 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 this way.

Theorem 9 *Let \mathfrak{S} be a coplanar orthogonal surface generated by \mathcal{V} supporting a Schnyder wood S on vertex set $V(S) \equiv \mathcal{V}$. Then there is a unique weight function $w : F(S) \rightarrow \mathbb{R}$ 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. A Schnyder wood S and a weight function w define an orthogonal surface $\mathfrak{S}_{S,w}$. This surface, however, need not support the initial Schnyder wood. From the proof of Theorem 5 it follows that a necessary and sufficient condition for an embedding $S \hookrightarrow \mathfrak{S}_{S,w}$ is that

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

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

Proof of Theorem 9. 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 i th coordinate of a_j for $j \neq i$. Subtracting $t = (t_1, t_2, t_3)$ from all generating vectors $v = (v_1, v_2, v_3)$ of the surface \mathfrak{S} , will normalize the figure such 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 the following we assume that \mathfrak{S} is normalized in this sense.

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

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

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

Proof. First suppose that S is the Schnyder wood of a triangulation. For the three special vertices, we only need the single equation $\mathbf{1} \cdot w = c$, with the all-one vector $\mathbf{1}$. This equation

together with the three equations of an inner vertex v is a dependent system: $\mathbb{1} = r_1(v) + r_2(v) + r_3(v)$ 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 $e = vw$ 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

Let e_F be the $(f - 1)$ -dimensional row vector with a single one at the position corresponding to the face F .

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

Proof. We distinguish three cases:

Case 1. The boundary of F is a directed cycle C (bidirected edges are allowed on C). From Lemma 4 or more directly from the rule of faces (A3) it follows that the cycle C consists of three directed paths in the three colors. If C is clockwise, the order of the paths is $P_1, P_2,$




Figure 13: Faces with a directed cycle in the boundary.

P_3 and if C is counterclockwise the order of the paths is P_1, P_3, P_2 , see Figure 13. 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 becomes

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

Case 2. We assume that the boundary of F is not a directed cycle. 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 (possibly $w_{i-1} = v_i$). We assume that there are two unidirected special edges of the same color i say $e_1 = (v_1, w_1), e_2 = (w_2, v_2)$.

Subcase $w_1 = v_2$. We first treat the case that e_3 is directed as (w_3, v_3) , this includes the case where e_3 is bidirected. The left of Figure 14 shows the situation with $i = 1$.




Figure 14: Faces without directed cycle, unidirected edges of the same color, and $w_1 = v_2$.

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

$$\mathbb{1} - (r_i(v_1) + r_{i+1}(v_1) + r_{i-1}(v_3)) = e_F.$$

If e_3 is directed as (v_3, w_3) , then, as shown in the right part of Figure 14:

$$\mathbb{1} - (r_{i-1}(w_2) + r_i(w_2) + r_{i+1}(w_3)) = e_F.$$

Subcase $w_1 \neq v_2$. In this case the boundary of F between w_1 and w_2 consists of edges bidirected in colors $i - 1, i + 1$. Let R be the region enclosed by this bidirected path, $P_i(w_1)$, and $P_i(v_2)$, and r the corresponding vector. Note that, as in Figure 15, $R_1(w_1)$, $R_2(w_1)$ and




Figure 15: The region R in the case $i = 1$.

$R_3(v_2)$ partition $B \setminus R$, hence

$$\mathbb{1} - (r_i(w_1) + r_{i+1}(w_1) + r_{i-1}(v_2)) = r.$$

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. \triangle

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 $i - 1, i + 1$, say $e_1 = (v_1, w_1)$, $e_2 = (w_2, v_2)$ and the third special edge is bidirected in colors $i - 1, i + 1$.

Subcase $w_1 = v_2$. The left of Figure 16 shows the situation with $i = 1$.




Figure 16: Faces without directed cycle, and 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,

$$\mathbb{1} - (r_{i-1}(v_2) + r_i(w_2) + r_i(v_1) + r_{i+1}(w_1) - r_i(w_3)) = e_F.$$

Subcase $w_1 \neq v_2$. This is analogous to Case 2, subcase $w_1 \neq v_2$, the right of Figure 16 shows the situation. \triangle

The dimension of the span of the face vectors e_F is $f - 1$. Claim 2 implies that the span of the face vectors is contained in the span of the region vectors $r_i(v)$. Together with Claim 1 this implies that the span of the region vectors is of dimension $f - 1$. Let r_1, \dots, r_{f-1} be a selection of rows which is a basis of this space and let R be the square matrix with rows r_i . Note that in the proof of Claim 1 one such basis is explicitly given. By Claim 2 there is a matrix A with $A \cdot R = I_{f-1}$, that is R is invertible and the linear system $(*)$ has a unique solution. \square

Next we will show how to obtain an efficient algorithm that computes the representation as in Theorem 9 for a given orthogonal surface \mathfrak{S} .

Theorem 10 *Let a non-degenerate, axial, coplanar orthogonal surface \mathfrak{S} be given, which is 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.*

Proof. We will first describe how 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 17 shows a snapshot of the intersection of P with \mathfrak{S} . For the sweep algorithm we need a data structure which 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.

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 Corollary 2. 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 knows all colored and directed edges of S which are induced by W .

We give a description of the algorithm. A priority queue Q and a dynamic tree P (the sweep front) both ordered lexicographically with respect to (x_1, x_2) are the data structures used.




Figure 17: 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 grey lines and vertices are the part of the surface that was already explored.

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. 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, \dots, s_ℓ be the successors of v in P , where s_ℓ is the first one with smaller or equal x_3 -coordinate than v . Remove $s_1, \dots, s_{\ell-1}$ 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 which 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 is the endvertex of the green outgoing edge of v , this edge is not induced by S 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.

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. Normalize all vertex coordinates by subtracting the translation. The faces are now considered one by one. When considering a face F , first determine of which of the possible twenty types F is. As indicated in the proof of Theorem 9 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:

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

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 correspond to the respective regions' weights and the weight of F can be calculated as in the proof of Theorem 9. For example, in the case shown in Figure 14 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 . As every edge lies in at most two inner faces, the runtime is linear: every edge is considered only a constant number of times and the number of edges is linear in the number of vertices for planar graphs. \square

4.2 Rigidity and Coplanarity

The face counting produces a coplanar surface for a given Schnyder wood and in Section 3 we have seen how to construct a rigid surface for a given Schnyder wood. Coplanarity and rigidity are both useful concepts in the realm of orthogonal surfaces. A natural question to ask is therefore, if for every Schnyder wood there is an orthogonal surface having both properties, i.e. that is rigid and coplanar. In this section we will present an example of a Schnyder wood for which a geodesic embedding can be either rigid or coplanar, but not both.

Proposition 2 *The Schnyder wood shown in Figure 18 cannot be embedded on a rigid and simultaneously coplanar surface.*




Figure 18: Schnyder wood on a rigid, but not coplanar surface

Proof. Assume 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 given instance 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:

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

Concatenating the inequalities from the right column of the table we obtain the contradiction $a_3 - b_3 < a_3 - b_3$. \square

5 Further Representations of Orthogonal Surfaces

In Theorem 9 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 .

It is interesting to identify other representations for orthogonal surfaces. In this section we mention some results and problems in this direction. We begin with a positive result.

Proposition 3 *A coplanar orthogonal surface \mathfrak{S} can be represented by a Schnyder wood S and a vector $(\ell_F)_{F \in \mathcal{F}}$ of lengths for the bounded faces of S . The length ℓ_F is the length of an orthogonal i -edge ending in the point (maximum) of \mathfrak{S} which represents F . If F is rightmost on its 1-flat, then ℓ_F is the length of the 1-edge, otherwise it is the length of the 2-edge.*

Figure 19 shows an example. We leave it to the reader to verify that the coplanarity of the surface and the length of the bold edges uniquely determine the length of all edges of the surface. Note that as in the case of Theorem 9 the proposition can be rephrased as a result about the invertibility of a certain matrix. In this case a typical equation corresponding to a




Figure 19: Combinatorial data together with the length of bold edges represent the surface.

row of the matrix is of the form $v_1 - w_1 = \ell_F$ if F is a rightmost face and $v_2 - w_2 = \ell_F$ for the other faces.

Another value that is naturally associated to every bounded face F is the height of the point of \mathfrak{S} which represents F . For simplicity we define the *height* of a point $p \in \mathfrak{S}$ as $h(p) = p_1 + p_2 + p_3$. Is it true that a coplanar orthogonal surface \mathfrak{S} can be represented by a Schnyder wood S and the vector $(h_F)_{F \in \mathcal{F}}$ of heights?

This question can be generalized to the case of non-coplanar orthogonal surfaces. In this case we would like to represent a surface by the heights of the points representing the vertices and faces. The dimension of the vector with an entry for every bounded face and every vertex actually exceeds the number of values needed to determine the generating minima \mathcal{V} of the normalized surface \mathfrak{S} by one. Hence the problem:

- Problem:** (a) Is it true that an orthogonal surface \mathfrak{S} can be represented by a Schnyder wood S and the vector $(h_F)_{F \in \mathcal{V} \cup \mathcal{F}}$ of heights?
(b) The equations implied by the height constraints must have one linear dependence. What is the combinatorial interpretation of this dependence?

In a preliminary version of the paper we have posed this *heights problem* as open. At that time we could only solve (a) and (b) for the special case where the graph supporting the Schnyder wood is a *stacked triangulation*. That is a triangulation that can be constructed starting from a triangle by repeatedly choosing a triangular face and adding a new vertex to this face connecting it to the three vertices on the boundary of the face.

Meanwhile we have an affirmative answer for part (a) of the problem while part (b) remains open. The proof for part (a) can be derived from results which were obtained in our attempt to prove the following conjecture of Jan Kratochvíl²:

Conjecture: Every 4-connected planar triangulation admits a triangle contact representation with equilateral, axis-aligned triangles.

Figure 20 shows an example of such a representation. Triangle contact representations for stacked triangulations and series-parallel graphs have been obtained in [1]. That paper also contains additional references and background information.

Triangle contact representations are closely related to orthogonal surfaces with the property that the points on the surface corresponding to edges of the graph are coplanar. Orthogonal surfaces with this property are always rigid. We know that not all Schnyder woods are supported by such surfaces. Examples are easily derived from the graph of the octahedron

²Conjecture was stated at the Bertinoro Workshop on Graph Drawing 2007.




Figure 20: A 4-connected triangulation and a triangle contact representation.

(cf. [1] for details). An answer to the following question would contain a proof or refutation of the above conjecture.

Question: Which planar 3-connected graphs are supported by an edge-coplanar orthogonal surface?

Acknowledgements. Stefan Felsner thanks Nicolas Bonichon for inspiring discussions. Florian Zickfeld was partially supported by the DFG Leibniz grant of Günter M. Ziegler and the Studienstiftung des deutschen Volkes.

References

- [1] 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*, 2007.
- [2] I. BÁRÁNY AND G. ROTE, *Strictly convex drawings of planar graphs*, *Documenta Mathematica*, 11 (2006), pp. 369–391.
- [3] N. BONICHON, S. FELSNER, AND M. MOSBAH, *Convex drawings of 3-connected planar graphs*, *Algorithmica*, 47 (2007), pp. 399–420.
- [4] G. BRIGHTWELL AND W. T. TROTTER, *The order dimension of convex polytopes*, *SIAM J. Discrete Math.*, 6 (1993), pp. 230–245.
- [5] G. BRIGHTWELL AND W. T. TROTTER, *The order dimension of planar maps*, *SIAM J. Discrete Math.*, 10 (1997), pp. 515–528.
- [6] Y. CHIANG, C. LIN, AND H. LU, *Orderly spanning trees with applications to graph encoding and graph drawing*, *SIAM J. Comput.*, 34 (2005), pp. 924–945.
- [7] 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.*, 1998, pp. 118–129.

- [8] S. FELSNER. <http://www.math.tu-berlin.de/~felsner/Schnyder.bib>.
- [9] S. FELSNER, *Convex drawings of planar graphs and the order dimension of 3-polytopes*, Order, 18 (2001), pp. 19–37.
- [10] S. FELSNER, *Geodesic embeddings and planar graphs*, Order, 20 (2003), pp. 135–150.
- [11] S. FELSNER, *Geometric Graphs and Arrangements*, Vieweg Verlag, 2004.
- [12] S. FELSNER AND S. KAPPES, *Orthogonal surfaces*. arXiv: math.CO/0602063, 2006. Submitted.
- [13] E. FUSY, D. POULALHON, AND G. SCHAEFFER, *Dissection and trees, with applications to optimal mesh encoding and random sampling*, in Proc. 16. ACM-SIAM Sympos. Discrete Algorithms, 2005, pp. 690–699.
- [14] G. KANT, *Drawing planar graphs using the canonical ordering*, Algorithmica, 16 (1996), pp. 4–32.
- [15] C. LIN, H. LU, AND I.-F. SUN, *Improved compact visibility representation of planar graphs via Schnyder’s realizer*, SIAM J. Discrete Math., 18 (2004), pp. 19–29.
- [16] E. MILLER, *Planar graphs as minimal resolutions of trivariate monomial ideals*, Documenta Math., 7 (2002), pp. 43–90.
- [17] W. SCHNYDER, *Planar graphs and poset dimension*, Order, 5 (1989), pp. 323–343.
- [18] W. SCHNYDER, *Embedding planar graphs on the grid*, in Proc. 1st ACM-SIAM Sympos. Discrete Algorithms, 1990, pp. 138–148.
- [19] W. T. TROTTER, *Combinatorics and Partially Ordered Sets: Dimension Theory*, The Johns Hopkins University Press, 1992.
- [20] W. T. TROTTER, *Partially ordered sets*, Handbook of Combinatorics, Vol I, Graham, Grötschel, Lovász (eds) (1995), pp. 433–480.