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 Schnyder woods and orthogonal surfaces. Schnyder woods and the face counting approach have important applications in graph drawing and dimension theory. Orthogonal surfaces explain the 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 dimension three. Our proof yields a companion linear time algorithm for the construction of the three linear orders that realize the face lattice.
Coplanar orthogonal surfaces are in correspondance 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 [15, 16] 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 [1], [2], [5], [13] use Schnyder woods, and more references can be found in [6].

In [15], 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 [17], [3] or [7].

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 [4] 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 [3] was long and technical, and Felsner gave a simpler one in [7]. Consecutively, Felsner [8] showed that every Schnyder wood of a 3-connected planar graph is supported by a rigid orthogonal surface (Theorem 5). An orthogonal surface is called rigid if it supports a unique graph, see Figures 3 and 4.b. By a result of Miller [14], Felsner's result implies Theorem 2. In Section 3 we present an intuitive proof of Theorem 5, that leads to a simple linear time algorithm for the computation of the rigid surface. The reduction to topological sorting it uses is simpler and more efficient than the constructions that are implicit in the other proofs. 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 4. 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 we refer the reader to [9].

As mentioned above, Section 3 deals with rigid orthogonal surfaces. Section 4 is concerned with coplanar surfaces, that is orthogonal surfaces with the property that all generating minima lie on some plane. 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. At the end of the section, we give an example of a Schnyder wood that has no supporting orthogonal surface which is simultaneously rigid and coplanar. We conclude with a related open problem.

Some proofs are omitted in this paper, others are considerably shortened. Complete proofs can be found in the full version [10].

2 Basics on Schnyder Woods and Orthogonal Surfaces

All the proofs omitted in the this section can be found in [9], [8] or [7]. A *planar map* M is a simple planar graph G together with a fixed planar embedding of

G. 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 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[1].

(W1) Every edge e is oriented in one directin 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 outdegree 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}.

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

We will sometimes 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 outdegree 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 $i = 1, 2, 3$.*

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 1.

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.*

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

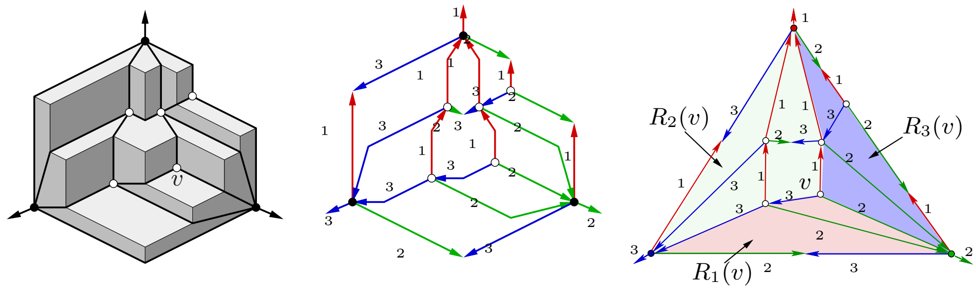

Fig. 1. An orthogonal surface, the induced Schnyder wood and the regions of a vertex v in this Schnyder wood. The small numbers correspond to edge colors.

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

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

Later we will use this vector to construct orthogonal surfaces supporting a given Schnyder wood. In that context $\{(v_1, v_2, v_3) \mid v \in V\}$ will be the generating set for the surface.

Another tool needed from the theory of Schnyder woods is the *edge split*. The following lemma from [2] describes the generic 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. 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.*

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 2 illustrates the operation.

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.*

We now introduce orthogonal surfaces and review some facts that we will need in the sequel. 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

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

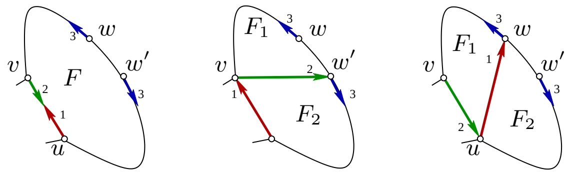

Fig. 2. The two possible types of splits of a non-special bidirected red-green edge uv in F. The numbers in the figure correspond to the edge colors.

The boundary $\mathfrak{S}_\mathcal{V}$ of $\langle \mathcal{V} \rangle$ is the *orthogonal surface* generated by $\mathcal{V}$, see Figure 1.

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. Figure 1 shows an axial orthogonal surface. These definitions have been proposed by Miller [14] who, essentially, also observed the following theorem.

Theorem 3. *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.*

An embedding of a Schnyder wood into an orthogonal surface is shown in Figure 1. A proof of the theorem can be found in [9].

Since every orthogonal arc leaving a vertex is occupied by an edge, every angle is completely contained in a *flat*. Basically, flats are the connected regions of constant gray-value in our drawings of orthogonal surfaces. To make this precise, let H be the plane $x_i = h$ and $\tilde{F}_1, \ldots, \tilde{F}_\ell$, the connected components of the interior of $H \cap \mathfrak{S}$. The topological closures $F_1, \ldots, F_\ell$ of these components are i-flats of height h. The i-flat of $v \in V$ is denoted by $F_i(v)$.

Given Theorem 3, it is natural to ask questions about *existence* and *uniqueness* of geodesic embeddings. A surface with three orthogonal arcs meeting in a single point does not support a Schnyder wood, see Figure 3.a. We call surfaces with a such a pattern degenerate. All other orthogonal surfaces support a

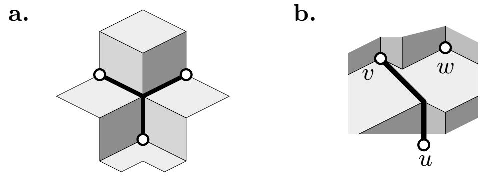

Fig. 3. a. A degenerate pattern. **b.** A non-rigid edge (u, v).

Schnyder wood. In Figure 3.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 $\mathfrak{S}_{\mathcal{V}}$ is a *rigid surface* if all its elbow geodesics are rigid. For an example, see the left part Figure 4, where the inner blue edge is not rigid.

3 Rigid Orthogonal Surfaces via Flat Shifting

We set Theorem 5 into context before we give a new proof. Miller [14] observed that a rigid orthogonal surface supports exactly one Schnyder wood and proved:

Theorem 4. *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 [14] (see also [8]) this implies the Brightwell-Trotter Theorem (Theorem 2).

Proposition 1. *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 region of M. If α_F is the join of the vertices of F, then $w \in F \Leftrightarrow w \leq \alpha_F$.*

We give a new proof of the following result by Felsner [8], who answered a question by Miller with this extension of Theorem 4.

Theorem 5. *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 will now give a sketch of the proof. Lemmas 6 and 7 are part of this sketch. 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 via 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 4 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 there is a vertex $w \in V$ such that $F_j(w) = F_j(u)$ and $w_i > u_i$, then $F_i(v) < F_i(w)$ in Γ_i.

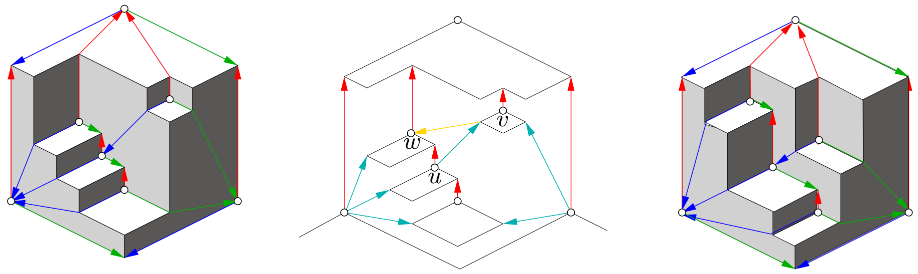

Fig. 4. A Schnyder wood S on a non-rigid surface, the corresponding relation Γ_1, and S on a rigid surface. For Γ_1 a-relations correspond to red arrows, p-relations to cyan arrows, and the only r-relation to a golden arrow.

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 5 as it justifies why the flat shifts (i.e. r-relations) can be combined to obtain a rigid surface.

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 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 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 $\{u, u_1\}$, $\{u_1, u_2\}$, ..., $\{u_{k-1}, u_k\}$ 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. The a-relations and p-relations are mapped as above while with an r-relation $F(u) < F(v)$ we associate the blue or green edge (v, u) which was introduced into S' by a split.

The idea is to show that a cycle C in Γ_1 would induce a cycle C' in $T_1 \cup T_2^{-1} \cup T_3^{-1}$ using the inverse of the mapping from edges to relations described above. Here the T_i, $i \in \{1, 2, 3\}$, are the respective trees of S' and the existence of C' yields a contradiction to Lemma 1. Note that 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 does not yield gaps in the intended cycle C' because vertices on the same flat are connected by a path of green-blue bidirected edges. The contradiction shows that Γ_i is acyclic. $\square$

Let $\mathfrak{S}$ be the orthogonal surface supporting S which is generated by the face counting vectors (c.f. Theorem 3). 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 on L_i, more formally to $\alpha_{F_i} = |\{F_i' \in \mathcal{F}_i : F_i' < F_i \text{ in } L_i\}|$. With V we associate a set of points $\mathcal{V}_\alpha = \{(\alpha_{F_1(v)}, \alpha_{F_2(v)}, \alpha_{F_3(v)}) \mid v \in V\} \subset \mathbb{R}^3$. We will outline the rest of the proof of Theorem 5, the first step is to prove 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 $u_i' < v_i'$.*

Lemma 7 is the key to proving the following four statements, which complete the proof of Theorem 5: $\mathcal{V}_\alpha$ is an antichain in $\mathbb{R}^3$, $\mathfrak{S}_{\mathcal{V}_\alpha}$ is non-degenerate, $\mathfrak{S}_{\mathcal{V}_\alpha}$ supports the Schnyder wood S, and $\mathfrak{S}_{\mathcal{V}_\alpha}$ is rigid. This completes the proof sketch for Theorem 5.

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

Proposition 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 stored. 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. Treat repel relations associated to green unidirected edges analogously. Finally, contract all blue-green edges from S 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 is implied by Theorem 5. $\square$

Theorem 6. *Let P be a 3-polytope with n vertices. Then, a Brightwell-Trotter realizer for P can be computed in $O(n)$ time.*

Proof. As Fusy et al. [11] show, a Schnyder wood S for the edge graph of P can be computed in $O(n)$ time. Alternatively this can be done with the help of Kant's algorithm [12] as well. From S construct a rigid orthogonal surface $\mathfrak{S}$, in time $O(n)$ using Proposition 2. Then, $\mathfrak{S}$ induces a Brightwell-Trotter realizer of P by Proposition 1. $\square$

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 $v_1 + v_2 + v_3 = c$. Schnyders classic approach of drawing graphs using the face-count vectors $\{(v_1, v_2, v_3)|v \in V\}$ yields a subclass of all coplanar surfaces, see Figure 1. 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. 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 fail as the drawings need not be crossing-free.

Theorem 7. *Let $\mathfrak{S}$ be a coplanar orthogonal surface supporting a Schnyder wood S. 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 $i \in \{1, 2, 3\}$*

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

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 3 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 sketch for Theorem 7. Let $\mathfrak{S}$ be a coplanar orthogonal surface and S a Schnyder wood induced by $\mathfrak{S}$. Let where $\{i, j, k\} = \{1, 2, 3\}$. We define $t_i = (a_j)_i = (a_k)_i$, which is possible since $F_i(a_j) = F_i(a_k)$ for the suspension vertices a_1, a_2, a_3 of S. First, shift the surface by (t_1, t_2, t_3), 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.

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

$$r_i(v)_F = 1 \text{ if } F \in R_i(v) \text{ and } r_i(v)_F = 0 \text{ otherwise.}$$

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

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

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

Claim 2. Let e_F be the $(f-1)$-dimensional row vector with a single one at the position corresponding to the face F. Then, e_F is in the span of the region-face incidence vectors $\{r_i(v) \mid i \in \{1,2,3\}, v \in V\}$.

Proof sketch. For the proof we distinguish several cases: the boundary of F is a directed cycle, two special edges of the same color are unidirected in opposite directions, two special edges of different colors are unidirected in opposite directions. There are several subcases to be distinguished. We present details for only one case here, the other proofs are similar.

If 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 $w_{i-1} = v_i$). Say $e_1 = (v_1, w_1)$, $e_2 = (w_2, v_2)$, are two unidirected of the same color in opposite directions.

We treat the case that $w_1 = v_2$ and e_3 is directed as (w_3, v_3), (this includes the case where e_3 is bidirected) explicitly. The left of Figure 5 shows the situation with $i = 1$.

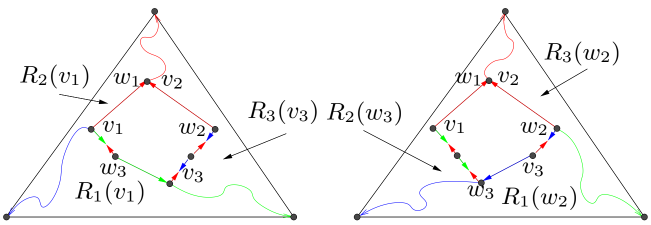

Fig. 5. Faces without directed cycle and $w_1 = v_2$.

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

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

The case that e_3 is directed as (v_3, w_3) is shown in the right part of Figure 5.△

Claims 1 and 2 together imply that the linear system $(*)$ has rank $f-1$ and hence a unique solution. □

Next we show how to obtain an efficient algorithm that computes the representation of Theorem 7 for a given orthogonal surface $\mathfrak{S}$.

Theorem 8. *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 first sketch how to extract the Schnyder wood S from $\mathfrak{S}$. The algorithm scans $\mathfrak{S}$ from bottom to top with a sweep plane P orthogonal to the x_1-axis. 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. Figure 6.a shows a snapshot of the intersection of P with $\mathfrak{S}$. The order in which the sweep considers the generators is the lexicographic order on (x_1, x_2). When a minimum v

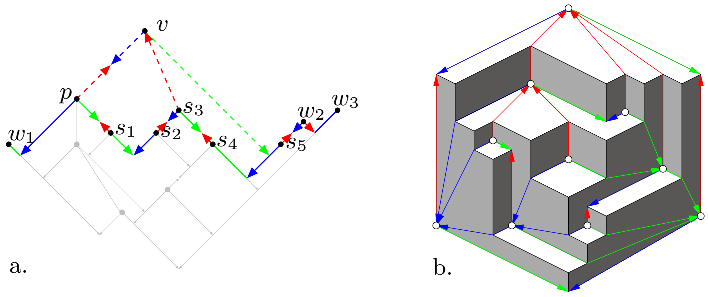

Fig. 6. Part a shows the projection of the explored part of an orthogonal surface on the sweep plane. The dotted arrows represent the new edges added with v, the other colored arrows are the sweep front. The gray line segments and vertices are the part of the surface that has already been explored. Part b illustrates Proposition 3, it shows a rigid but not coplanar surface.

is added, its blue outgoing edge and incoming red edges are added as well. The green outgoing edge is added only if it is not a green-blue bidirected edge, in this case it is added when the next minimum is treated. This procedure builds the Schnyder wood step by step, and needs $O(n \log n)$ time when the sweep front is implemented as a dynamic search tree.

The second part of the algorithm is the computation of the face weights. After normalizing all coordinate vectors the faces are now considered one by one. When considering a face F, we first determine the type of F. Based on the proof of Theorem 7 we distinguish twenty such types.

The weight of F can be computed from certain region weights of boundary vertices of F. The required region weights are the coordinates of these vertices. If for example we are in the case shown in Figure 5 then the weight of F is $c-(w_2)_1-(w_3)_2-(w_2)_3$ where c is the constant obtained through normalization.

When scanning a face F we touch each edge only once and every edge lies in at most two inner faces. This implies a runtime of $O(n)$. □

Face counting produces coplanar surfaces supporting a given Schnyder wood S. In Section 3 we have seen how to construct a rigid surface supporting S. Coplanarity and rigidity are useful properties for an orthogonal surface. It is natural to ask whether every Schnyder wood has a supporting surface with both properties. The answer to this question is negative, the proof of Proposition 3 can be found in [10].

Proposition 3. *The Schnyder wood shown in Figure 6.b cannot be embedded on a rigid and simultaneously coplanar surface.*

5 Conclusions

We conclude our investigations of the connections between Schnyder woods and orthogonal surfaces with an open problem of a flavor similar to Theorem 7. Let S be a Schnyder wood induced by an orthogonal surface $\mathfrak{S}$. From Proposition 1 it follows that the bounded faces of S are in bijection with the maxima of $\mathfrak{S}$. We refer to the set of minima and maxima as $\mathcal{V} \cup \mathcal{F}$. For $p \in \mathfrak{S}$ we define its *height* as $h(p) = p_1 + p_2 + p_3$.

Problem. Given $\mathfrak{S}$, do the Schnyder Wood S and the vector $h = (h(v))_{v \in \mathcal{V} \cup \mathcal{F}}$ of heights uniquely determine $\mathfrak{S}$?

We can prove this in the case where the underlying graph is a stacked triangulation. With computer's help we have verified that the answer is affirmative for small triangulations with up to twelve vertices.

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. I. Bárány and G. Rote, *Strictly convex drawings of planar graphs*, 2005. arXiv:cs.CG/0507030.
2. N. Bonichon, S. Felsner, and M. Mosbah, *Convex drawings of 3-connected planar graphs*, in Graph Drawing (Proc. GD '04), vol. 3383 of Lecture Notes in Comput. Sci., 2004, pp. 60–70.
3. G. Brightwell and W. T. Trotter, *The order dimension of convex polytopes*, SIAM J. Discrete Math., 6 (1993), pp. 230–245.
4. G. Brightwell and W. T. Trotter, *The order dimension of planar maps*, SIAM J. Discrete Math., 10 (1997), pp. 515–528.
5. G. Di Battista, R. Tamassia, and L. Vismara, *Output-sensitive reporting of disjoint paths*, Algorithmica, 23 (1999), pp. 302–340.
6. S. Felsner. `http://www.math.tu-berlin.de/~felsner/Schnyder.bib`.
7. S. Felsner, *Convex drawings of planar graphs and the order dimension of 3-polytopes*, Order, 18 (2001), pp. 19–37.
8. S. Felsner, *Geodesic embeddings and planar graphs*, Order, 20 (2003), pp. 135–150.
9. S. Felsner, *Geometric Graphs and Arrangements*, Vieweg Verlag, 2004.
10. S. Felsner and F. Zickfeld, *Schnyder woods and orthogonal surfaces*, 2006. `http://www.math.tu-berlin.de/~felsner/swaos.pdf`.
11. 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.
12. G. Kant, *Drawing planar graphs using the lmc-ordering*, in Proc. 33rd IEEE Sympos. on Found. of Comp. Sci., 1992, pp. 101–110.
13. 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.
14. E. Miller, *Planar graphs as minimal resolutions of trivariate monomial ideals*, Documenta Math., 7 (2002), pp. 43–90.
15. W. Schnyder, *Planar graphs and poset dimension*, Order, 5 (1989), pp. 323–343.
16. W. Schnyder, *Embedding planar graphs on the grid*, in Proc. 1st ACM-SIAM Sympos. Discrete Algorithms, 1990, pp. 138–148.
17. W. T. Trotter, *Combinatorics and Partially Ordered Sets: Dimension Theory*, The Johns Hopkins University Press, 1992.