Lattice Structures from Planar Graphs

Stefan Felsner

Technische Universität Berlin, Institut für Mathematik, MA 6-1
Straße des 17. Juni 136, 10623 Berlin, Germany
felsner@math.tu-berlin.de

Submitted: Dec 2, 2002; Accepted: Jan 23, 2004; Published: Feb 14, 2004

MR Subject Classifications: 05C10, 68R10, 06A07

Abstract.

The set of all orientations of a planar graph with prescribed outdegrees carries the structure of a distributive lattice. This general theorem is proven in the first part of the paper. In the second part the theorem is applied to show that interesting combinatorial sets related to a planar graph have lattice structure: Eulerian orientations, spanning trees and Schnyder woods. For the Schnyder wood application some additional theory has to be developed. In particular it is shown that a Schnyder wood for a planar graph induces a Schnyder wood for the dual.

1 Introduction

This work originated in the study of rigid embeddings of planar graphs and the connections with Schnyder woods. These connections were discovered by Miller [9] and further investigated in [3]. The set of Schnyder woods of a planar triangulation has the structure of a distributive lattice. This was independently shown by Brehm [1] and Mendez [10]. My original objective was to generalize this and prove that the set of Schnyder woods of a 3-connected planar graph also has a distributive lattice structure. The theory developed to this aim turned out to work in a more general situation. In the first half of this paper we present a theory of α-orientations of a planar graph and show that they form a distributive lattice. As noted in [4] this result was already obtained in the thesis of Mendez [10]. Another source for related results is a paper of Propp [13] where he describes lattice structures in the dual setting. The cover relations in Propp's lattices are certain *pushing-down* operations. These operations were introduced by Mosesian and further studied by Pretzel [11] as reorientations of diagrams of ordered sets.

The second part of the paper deals with special instances of the general result. In particular we find lattice structures on the following combinatorial sets related to a planar graph: Eulerian orientations, spanning trees and Schnyder woods. While the application to Eulerian orientations is rather obvious already the application of spanning trees requires some ideas. To connect spanning trees to orientations we introduce the completion of a plane graph which can be thought of as superposition of the primal and the dual which is planarized by introducing a new edge-vertex at every crossing pair of a primal edge with its dual edge. The lattice structure on spanning trees of a planar graph has been discovered in the context of knot theory by Gilmer and Litherland [5] and by Propp [13] as an example of his lattice structures. A closely related family of examples concerns lattices on matchings and more generally f-factors of plane bipartite graphs.

To show that the Schnyder woods of a 3-connected plane graph have a distributive lattice structure some additional theory has to be developed. We prove that a Schnyder wood for a planar graph induces a Schnyder wood for the dual. A primal dual pair of Schnyder woods can be embedded on a completion of the plane graph, i.e., on a superposition of the primal and the dual as described above. In the next step it is shown that the orientation of the completion alone allows to recover the Schnyder wood. As in the case of spanning trees the lattice structure comes from orientations of the completion.

2 Lattices of Fixed Degree Orientations

A *plane graph* is a planar graph $G = (V, E)$ together with a fixed planar embedding. In particular there is a designated outer (unbounded) face F^* of G. Given a mapping $\alpha : V \to \mathbb{N}$ an orientation X of the edges of G is called an *α-orientation* if α records the out-degrees of all vertices, i.e.,$\mathsf{outdeg}_X(v) = \alpha(v)$ for all $v \in V$. We call α feasible if α-orientation of G exists. The main result of this section is the following theorem.

Theorem 1. *Let G be a plane graph and $\alpha : V \to \mathbb{N}$ be feasible. The set of α-orientations of G carries an order-relation which is a distributive lattice.*

2.1 Reorientations and essential cycles

Let X be an α-orientation of G. Given a directed cycle C in X we let X^C be the orientation obtained from X by reversing all edges of C. Since the out-degree of a vertex is unaffected by the reversal of C the orientation X^C is another α-orientation of G. The plane embedding of G 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. Brief remark in passing: The transitive closure of the 'left of' relation is the order relation which makes the set of α-orientations of G a distributive lattice.

Let X and Y be α-orientations of G and let D be the set of edges with oppositional orientations in X and Y. Every vertex is incident to an even number of edges in D, hence, the subgraph with edge set D is Eulerian. If we impose the orientation of X on the edges of D the subgraph is a directed Eulerian graph. Consequently, the edge set D can be decomposed into simple cycles $C_1, .., C_k$ which are directed cycles of X. We restate a consequence of this observation as a lemma.

Lemma 1. *If $X \neq Y$ are α-orientations of G then for every edge e which is oppositionally directed in X and Y there is a simple cycle C with $e \in C$ and C is oppositionally directed in X and Y.*

An edge of G is *α-rigid* if it has the same direction in every α-orientation. Let $R \subseteq E$ be the set of α-rigid edges. Since directed cycles in X can be reversed, rigid edges never belong to directed cycles.

With $A \subset V$ we consider two sets of edges, the set $E[A]$ of edges with two ends in A, i.e., edges induced by A, and the set $E_{\mathsf{Cut}}[A]$ of edges in the cut $(A, \overline{A})$, i.e., the set of edges connecting a vertex on A to a vertex in the complement $\overline{A} = V \setminus A$.

Given A and a α-orientation X, then exactly $\sum_{v \in A} \alpha(v)$ edges have their tail in A. The number of edges incident to vertices in A is $|E[A]| + |E_{\mathsf{Cut}}[A]|$. The *demand* of A in X is the number of edges pointing from $\overline{A}$ into A.

Lemma 2. *For a set $A \subset V$ the demand is*

$$\mathsf{dem}_\alpha(A) = |E[A]| + |E_{\mathsf{Cut}}[A]| - \sum_{v \in A} \alpha(v)$$

In particular $\mathsf{dem}_\alpha(A)$ only depends on α and not on X.

By looking at demands we can identify certain sets of rigid edges. If for example $\mathsf{dem}_\alpha(A) = 0$, then all the edges in $E_{\mathsf{Cut}}[A]$ point away from A in every α-orientation and, hence, $E_{\mathsf{Cut}}[A] \subseteq R$ in this case. Symmetrically, if $\mathsf{dem}_\alpha(A) = |E_{\mathsf{Cut}}[A]|$, then all the edges in $E_{\mathsf{Cut}}[A]$ point towards A and again $E_{\mathsf{Cut}}[A] \subseteq R$.

Digression: Existence of α-orientations

For any graph $G = (V, E)$ and $\alpha : V \to \mathbb{N}$ there are two obvious necessary conditions on the existence of an α-orientation:

1. $\mathsf{dem}_\alpha(V) = 0$, i.e, $\sum_v \alpha(v) = |E|$,
2. $0 \leq \mathsf{dem}_\alpha(A) \leq |E_{\mathsf{Cut}}[A]|$ for all $A \subseteq V$.

It is less obvious that these two conditions are already sufficient for the existence of an α-orientation. This can be shown by a simple induction on the number of edges. If there is an $A \subset V$ with $\mathsf{dem}_\alpha(A) = 0$ then all edges in $E_{\mathsf{Cut}}[A]$ have to point away from A. Remove these edges, update α accordingly and apply induction to the components. If $\mathsf{dem}_\alpha(A) > 0$ for all $A \subset V$ then orient some edge arbitrarily. Remove this edge, update α accordingly and apply induction.

This simple proof has the disadvantage that it does not yield a polynomial algorithm to check the conditions and construct an α-orientation if the conditions are fulfilled. These requirements are matched by the following reduction to a flow-problem.

Start with an arbitrary orientation Z and let $\beta(v) = \mathsf{indeg}_Z(v)$. If $\beta(v) = \alpha(v)$ for all v then reversing the directions of all edges in Z yields an α-orientation. Otherwise we ask for a flow f subject to capacity constraints $0 \leq f(e) \leq 1$ for all directed edges $e \in Z$ and vertex constraints

$$\sum_{e \in out_Z(v)} f(e) - \sum_{e \in in_Z(v)} f(e) = \alpha(v) - \beta(v).$$

If such a flow exists, then there also exists an integral flow, i.e., $f(e) \in \{0,1\}$ for all e. Reversing the directions of those edges in Z which have $f(e) = 0$ yields an α-orientation. The existence of the flow is equivalent to the cut-conditions: For $A \subset V$ consider the amount of flow that has to go from A to $\overline{A}$. This amount is $\sum_{v\in A}\alpha(v) - \sum_{v\in A}\beta(v)$, The flow leaving A is constrained by the capacity of the cut, i.e., number of edges oriented from A to $\overline{A}$ in Z, this number is $|E[A]| + |E_{\mathsf{Cut}}[A]| - \sum_{v\in A}\beta(v)$. Thus the cut-condition $\sum_{v\in A}\alpha(v) - \sum_{v\in A}\beta(v) \le |E[A]| + |E_{\mathsf{Cut}}[A]| - \sum_{v\in A}\beta(v)$ is equivalent to $0 \le \mathsf{dem}_\alpha(A)$.

This ends the digression and we return to the study of α-orientations of a planar graph G. The set of vertices in the interior of a simple cycle C in G is denoted I_C. Of special interest to us will be cycles C with the property that $E_{\mathsf{Cut}}[I_C] \subseteq R$. In that case we say that *the interior cut of C is rigid*. This means that the orientation of all the edges connecting C to an interior vertex is fixed throughout all α-orientations. Note that the interior cut of a face cycle of G is always rigid because $E_{\mathsf{Cut}}[I_C] = \emptyset$ in this case.

Definition 1. *A cycle C of G is an* essential cycle *if*

- *C is simple and chord-free,*
- *the interior cut of C is rigid, i.e., $E_{\mathsf{Cut}}[I_C] \subseteq R$,*
- *there exists an α-orientation X such that C is a directed cycle in X.*

With lemmas 3–6 we show that with reorientations of essential cycles we can commute between any two α-orientations. In fact reorientations of essential cycles represent the cover relations in the 'left of' order on α-orientations.

A cycle C has a *chordal path* in X if there is a directed path consisting of edges interior to C whose first and last vertex are vertices of C. We allow that the two end vertices of a chordal path coincide.

Lemma 3. *If C has no chordal path in some α-orientation X, then the interior cut of C is rigid.*

Proof. Assume that C has no chordal path in some α-orientation X. Let A be the set of vertices which are reachable in X by a directed path starting from C with an edge pointing into the interior of C. The definition of A and the assumption that C has no chordal path in X imply that $A \subseteq I_C$ and all edges in the cut $(A, \overline{A})$ are directed toward A in X, i.e., $\mathsf{dem}_\alpha(A) = |E_{\mathsf{Cut}}[A]|$. Let $B = I_C \setminus A$, the definition of A and $\mathsf{dem}_\alpha(A) = |E_{\mathsf{Cut}}[A]|$ imply that $\mathsf{dem}_\alpha(B) = 0$. This implies that the interior cut of C is rigid. □

If C is a directed cycle the implication from the previous lemma is in fact an equivalence (Lemma 4). This provides us with a nice criterion for deciding whether a directed cycle is essential.

Lemma 4. *Let C be a directed cycle in an α-orientation X. The interior cut of C is rigid iff C has no chordal path in X.*

Proof. A chordal path P of C in X can be extended to a cycle C' by adding some edges of the directed cycle C. Reversing C' in X yields another α-orientation X'. The orientation of the first edge e of P is different in X and X'. The edge e belongs to the interior cut of C, hence, the interior cut of C is not rigid. ☐

Lemma 5. *If C and C' are essential cycles, then either the interior regions of the cycles are disjoint or one of the interior regions is contained in the other and the two cycles are vertex disjoint. See Figure 1 for an illustration.*

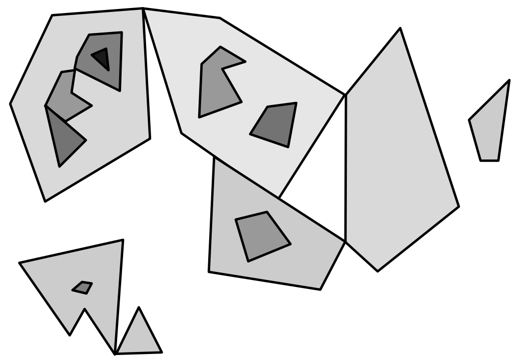

Figure 1: Interiors of essential cycles are disjoint or contained with disjoint borders.

Proof. In all other cases an edge e of one of the cycles, say C', would connect a vertex on C to an interior vertex of C. Since C is essential and e belongs to the interior cut of C edge e is rigid. However, e belongs to C' which is essential, therefore, there is an α-orientation X such that C' is directed in X. Let $X^{C'}$ be the orientation obtained from X by reversing C'. The two orientations show that e is not rigid; contradiction. ☐

Corollary 1. *Let e be and edge and F an incident face in G, then there exists at most one essential cycle C with $e \in C$ and $F \subseteq \mathsf{Int}(C)$.*

Lemma 6. *If C is a cycle which is directed in X, then X^C can also be obtained by a sequence of reversals of essential cycles.*

Proof. We show that as long as C is not essential we find cycles C_1 and C_2 such that $X^C = (X^{C_1})^{C_2}$ and both C_i are less complex than C so that we can apply induction.

If C is not simple we cut C at a vertex which is visited multiply to obtain C_1 and C_2.

If C has a chord e. Suppose that e is oriented as $e = (v, u)$ in X. Decompose C into a path P_1 from u to v and a path P_2 from v to u. Let C_1 be P_1 together with e. After reversing C_1 the reoriented edge e together with P_2 forms a cycle C_2 which is admissible for reorientation. Clearly $X^C = (X^{C_1})^{C_2}$.

If C is simple, chord-free and directed in X, but not essential, then the interior cut of C is not rigid. Lemma 3 implies that C has a chordal path P in X. Let u and v be the end-vertices of P on C and let P be directed from v to u. As in the previous case decompose C into a path P_1 and P_2. Again $C_1 = P_1 \cup P$ and $C_2 = P_2 \cup P$ are two less complex cycles with $X^C = (X^{C_1})^{C_2}$. ☐

Let C be a simple cycle which is directed in X as a ccw-cycle. If C_1 and C_2 are constructed as in the proof of the lemma then C_1 is is a ccw-cycle in X and C_2 is a ccw-cycle in X^{C_1}. This suggests a stronger statement:

Lemma 7. *If C is a simple directed ccw-cycle in X, then X^C can also be obtained by a sequence of reversals of essential cycles from ccw to cw. Moreover, the set of essential cycles involved in such a sequence is the unique minimal set such that the interior regions of the essential cycles cover the interior region of C.*

Proof. The proof of Lemma 6 provides a set of essential cycles such that all the reorientations are from ccw to cw. Furthermore the interior regions of these essential cycles are disjoint and cover the interior region of C.

It remains to prove the uniqueness. An edge on the boundary of the union of a set of essential cycles (viewed as topological discs) is only contained in one of the cycles (Corollary 1) and will therefore change its orientation in the sequence of reorientations. Consequently, this boundary is just C and all the essential cycles involved are contained in the interior of C.

A similar consideration shows that the interiors are disjoint. If the interiors of two of the cycles are not disjoint, then (Lemma 5) one of them is contained in the other, call the larger one C'. For the set of all essential cycles contained in C' we again observe: An edge on the boundary of the union of this set is only contained in one of the cycles and will change its orientation in the sequence of reorientations. Therefore such an edge interior to C' has to belong to C which is impossible. □

2.2 Interlaced flips in sequences of flips

A *flip* is the reorientation of an essential cycle from ccw to cw. A *flop* is the converse of a flip, i.e., the reorientation of an essential cycle from cw to ccw.

A *flip sequence on* X is a sequence $(C_1, .., C_k)$ of essential cycles such that C_1 is flipable in X, i.e., C_1 is a ccw-cycle of X, and C_i is flipable in $X^{C_1 \ldots C_{i-1}}$ for $i = 2, .., k$.

Recall from Corollary 1 that an edge e is contained in at most two essential cycles. If we think of e as directed, then there can be an essential cycle $C^{l(e)}$ left of e and another essential cycle $C^{r(e)}$ right of e.

Lemma 8. *If $(C_1, .., C_k)$ is a flip sequence on X then for every edge e the essential cycles $C^{l(e)}$ and $C^{r(e)}$ alternate in the sequence, i.e., if $i_1 < i_2$ with $C_{i_1} = C_{i_2} = C^{l(e)}$ then there is a j with $i_1 < j < i_2$ and $C_j = C^{r(e)}$. The same holds with left and right exchanged.*

Proof. Let F be the face with $e \in F$ and $F \subset \mathsf{Int}(C^{l(e)})$. If F is left of e in the current orientation then $C^{l(e)}$ may be flipable but $C^{r(e)}$ is clearly not a ccw-cycle and, hence, not flipable. □

Lemma 9. *For every edge e there is a $t_e \in \mathbb{N}$ such that for all α-orientations X a flip sequence on X implies at most t_e reorientations of e.*

Proof. If e is not contained in an essential cycle, then e is rigid and $t = 0$. Let C^1 be an essential cycle containing e, choose a point $x \in \mathsf{Int}(C^1)$ and consider a horizontal ray ℓ from x to the right. Ray ℓ will leave $\mathsf{Int}(C^1)$ at an edge e^1, let C^2 be the essential cycle on the other side of e^1. Further right ℓ will leave $\mathsf{Int}(C^2)$ at an edge e^2, let C^3 be the essential cycle on the other side of e^2. Repeat the construction until ℓ leaves $\mathsf{Int}(C^s)$ at e^s and this edge has no essential cycle on the other side. Such an s exists since ℓ emanates into the unbounded face of G which is not contained in the interior of an essential cycle.

Now we apply Lemma 8 backwards for every pair C^i, C^{i-1}. Since C^s is flipped at most once in any flip-sequence we find that C^{s-1} is flipped at most twice, C^{s-2} is flipped at most three times and so on. Hence, C^1 is flipped at most s times. With Lemma 8 this bound implies that edge e is reoriented at most $2s + 1$ times in any sequence of flips. □

Lemma 10. *The length of any flip sequence is bounded by some $t \in \mathbb{N}$ and there is a unique α-orientation $X_{\min}$ with the property that all cycles in $X_{\min}$ are cw-cycles.*

Proof. The number of essential cycles of G is finite. It can e.g. be bounded by the number of faces of G. For each essential cycle there is a finite bound for the number of times it can be flipped in a flip sequence Lemma 9. This makes a finite bound on the length of any flip-sequence.

Let X be an arbitrary α-orientation and consider a maximal sequence of flips starting at X. Let Y be the α-orientation reached through this sequence of flips. If Y would contain a ccw-cycle then by Lemma 7 there is an essential ccw-cycle and hence a possible flip. This is a contradiction to the maximality of the sequence, hence, Y is an α-orientation without ccw-cycles. By Lemma 1 there can be only one α-orientation without ccw-cycles, denoted $X_{\min}$. In particular a maximal sequence of flips starting in an arbitrary α-orientation X always leads to $X_{\min}$. □

From this lemma it follows that the 'left of' relation is acyclic. We now adopt a more order theoretic notation and write $Y \prec X$ if Y can be obtained by a sequence of flips starting at X. We summarize our knowledge about this relation.

Corollary 2. *The relation $\prec$ is an order relation with a unique minimal element $X_{\min}$.*

2.3 Flip-sequences and potentials

With the next series of lemmas we investigate properties of sequences of flips that lead from X to $X_{\min}$. It will be shown that any two such sequences contain the same essential cycles.

Lemma 11. *Suppose $Y \prec X$ and let C be an essential cycle. Every sequence $S = (C_1, \ldots, C_k)$ of flips that transforms X into Y contains the same number of flips at C.*

Proof. We recycle the proof technique used in Lemma 9. Let $C = C^1$, choose a point $x \in \mathsf{Int}(C^1)$ and consider a horizontal ray ℓ from x to the right. Let $C^1, \ldots, C^s$ be the sequence of essential cycles defined by ℓ, that is, C^i and C^{i+1} share an edge e^i and $e^s \in C^s$ has no essential cycle on its other side.

For the essential cycle C^i let $z_S(C^i) = |\{j : C_j = C^i\}|$ be the number of occurrences of C^i in the sequence S. Since C^i and C^{i+1} share an edge it follows from Lemma 8 that $|z_S(C^i) - z_S(C^{i+1})| \leq 1$ and $z_S(C^s) \leq 1$.

Let D be the set of edges with different orientations in X and Y. If $e^i \notin D$ then e^i is reoriented an even number of times by S. There are only two essential cycles available to reorient e^i (Corollary 1) these cycles are C^i and C^{i+1}. Since $|z_S(C^i) - z_S(C^{i+1})|$ is even and at most one it follows that $z_S(C^i) = z_S(C^{i+1})$ for all $e^i \notin D$.

An edge $e^i \in D$ is reoriented an odd number of times. There remain two cases either $z_S(C^i) = z_S(C^{i+1}) + 1$ or $z_S(C^i) = z_S(C^{i+1}) - 1$. The decision which case applies depends on the orientation of e^i in X. If C^i is left of the directed edge e^i in X then C^i is ccw and C^{i+1} is cw in X. This implies that the first flip of C^i precedes the first flip of C^{i+1} in every flip sequence that starts with X. Therefore, $z_S(C^i) = z_S(C^{i+1}) + 1$ in this case. If, however, C^{i+1} is left of the directed edge e^i in X then $z_S(C^i) = z_S(C^{i+1}) - 1$.

These rules show that X and Y uniquely determine $z_S(C^1) = z_S(C)$. A possible way to express the value is $z_S(C) = \big|\{e^i : e^i \in D$ and in X edge e^i is crossing ℓ from below$\}\big| - \big|\{e^i : e^i \in D$ and in X edge e^i is crossing ℓ from above$\}\big|$. $\square$

For a given α let $\mathcal{E} = \mathcal{E}_\alpha$ be the set of all essential cycles. Given an α-orientation X there is a flip sequence S from X to $X_{\min}$. For $C \in \mathcal{E}$ let $z_X(C)$ be the number of times C is flipped in a flip sequence S. The previous lemma shows that this independent of S and hence a well defined mapping $z_X : \mathcal{E} \to \mathbb{N}$. Moreover, if $X \neq Y$ then $z_X \neq z_Y$.

Definition 2. *An* α-potential *for G is a mapping $\wp : \mathcal{E}_\alpha \to \mathbb{N}$ such that*

- $|\wp(C) - \wp(C')| \leq 1$, *if C and C' share an edge e.*
- $\wp(C) \leq 1$, *if there is an edge $e \in C$ such that C is the only essential cycle to which e belongs.*
- *If $C^{l(e)}$ and $C^{r(e)}$ are the essential cycles left and right of e in $X_{\min}$ then* $\wp(C^{l(e)}) \leq \wp(C^{r(e)})$.

Lemma 12. *The mapping $z_X : \mathcal{E}_\alpha \to \mathbb{N}$ associated to an α-orientation X is an α-potential.*

Proof. The first two properties are immediate from the alternation property shown in Lemma 8. For the third property consider a flip-sequence S from X to $X_{\min}$. The orientation of e in $X_{\min}$ implies that the last flip affecting e is a flip of $C^{r(e)}$. With Lemma 8 this implies $\wp(C^{l(e)}) \leq \wp(C^{r(e)})$. $\square$

Lemma 13. *For every α-potential $\wp : \mathcal{E}_\alpha \to \mathbb{N}$ there is an α-orientation X with $z_X = \wp$.*

Proof. We define an orientation $X_\wp$ of the edges of G as follows.

- If e is not contained in an essential cycle then $X_\wp(e) = X_{\min}(e)$, i.e., the orientation of e in X equals the orientation of e in $X_{\min}$ (these are the rigid edges).
- If e is contained in one essential cycle C^e, then $X_\wp(e) = X_{\min}(e)$ if $\wp(C^e) = 0$ and $X_\wp(e) \neq X_{\min}(e)$ if $\wp(C^e) = 1$.

- If e is contained in two essential cycles $C^{l(e)}$ which $C^{r(e)}$ are left and right of e in $X_{\min}$, then $X_\wp(e) = X_{\min}(e)$ if $\wp(C^{l(e)}) = \wp(C^{r(e)})$ and $X_\wp(e) \neq X_{\min}(e)$ if $\wp(C^{l(e)}) \neq \wp(C^{r(e)})$.

It remains to show that $X_\wp$ is indeed an α-orientation. This is proven by induction on $\wp(\mathcal{E}) = \sum_{C\in\mathcal{E}} \wp(C)$.

If $\wp(\mathcal{E}) = 0$ then $X_\wp(e) = X_{\min}(e)$ for all e and $X_\wp$ is an α-orientation.

If $\wp(\mathcal{E}) > 0$ let m be the maximum value taken by $\wp$. Let R_m be the union of the interiors $\mathsf{Int}(C)$ of all the essential cycles C with $\wp(C) = m$. Let ∂R_m be the boundary of R_m. The third property of a potential implies that in $X_{\min}$ every edge $e \in \partial R_m$ has R_m on its right side. Therefore, ∂R_m decomposes into simple cycles which are cw in $X_{\min}$ and ccw in $X_\wp$. Let B be one of these cycles in ∂R_m. By Lemma 7 there is a unique subset $\mathcal{E}_B$ of $\mathcal{E}$ such that the flip of C is equivalent to flipping each member of $\mathcal{E}_B$.

Define $\wp^* : \mathcal{E} \to \mathbb{N}$ by $\wp^*(C) = \wp(C) - 1$ if $C \in \mathcal{E}_B$ and $\wp^*(C) = \wp(C)$ if $C \in \mathcal{E} \setminus \mathcal{E}_B$. We claim that $\wp^*$ is a potential. To prove this we have to check the properties of the definition for all edges. For edges that are not contained in B these properties for $\wp^*$ immediately follow from the properties for $\wp$. For $e \in B$ the definition of B implies $\wp(C^{l(e)}) = \wp(C^{r(e)}) - 1$. Since $C^{l(e)} \in \mathcal{E}_B$ and $C^{r(e)} \notin \mathcal{E}_B$ this shows $\wp^*(C^{l(e)}) = \wp^*(C^{r(e)})$.

By induction the orientation $X_{\wp^*}$ corresponding to the potential $\wp^*$ by the above rules is an α-orientation. The orientations $X_{\wp^*}$ and $X_\wp$ only differ on the edges of the directed cycle B which is cw in $X_{\wp^*}$ and ccw in $X_\wp$. Therefore, the outdegree of a vertex in $X_\wp$ equals its outdegree in $X_{\wp^*}$. This proves that $X_\wp$ is an α-orientation. Along the same inductive line it also follows that $z_{X_\wp} = \wp$. □

With Lemma 12 and Lemma 13 we have established a bijection between α-orientations and α-potentials. The following lemma completes the proof of Theorem 1.

Lemma 14. *The set of all α-potentials $\wp : \mathcal{E} \to \mathbb{N}$ with the dominance order $\wp \prec \wp'$ if $\wp(C) \leq \wp'(C)$ for all $C \in \mathcal{E}$ is a distributive lattice. Join $\wp_1 \vee \wp_2$ and meet $\wp_1 \wedge \wp_2$ of two potentials $\wp_1$ and $\wp_2$ are given by $(\wp_1 \vee \wp_2)(C) = \max\{\wp_1(C), \wp_2(C)\}$ and $(\wp_1 \wedge \wp_2)(C) = \min\{\wp_1(C), \wp_2(C)\}$ for all $C \in \mathcal{E}$.*

Proof. The fact that max and min fulfill the distributive laws is a folklore result. Therefore, all that has to be shown is that $\wp_1 \vee \wp_2$ and $\wp_1 \wedge \wp_2$ are potentials. Consider an edge e and and the essential cycles $C^{l(e)}$ and $C^{r(e)}$. From $\wp_i(C^{l(e)}) \leq \wp_i(C^{r(e)})$ for $i = 1, 2$, it follows that $(\wp_1 \vee \wp_2)(C^{l(e)}) \leq (\wp_1 \vee \wp_2)(C^{r(e)})$. If $(\wp_1 \vee \wp_2)(C^{r(e)}) = \wp_i(C^{r(e)})$ then $(\wp_1 \vee \wp_2)(C^{l(e)}) \geq \wp_i(C^{l(e)}) \geq \wp_i(C^{r(e)}) - 1$ hence $|(\wp_1 \vee \wp_2)(C^{r(e)}) - (\wp_1 \vee \wp_2)(C^{l(e)})| \leq 1$. This shows that the join $\wp_1 \vee \wp_2$ is a potential. The argument for the meet is similar. □

Corollary 3. *Let G be a plane graph and $\alpha : V \to \mathbb{N}$ be feasible. The following sets carry isomorphic distributive lattices*

- *The set of α-orientations of G.*
- *The set of α-potentials $\wp : \mathcal{E}_\alpha \to \mathbb{N}$.*
- *The set of Eulerian subdigraphs of a fixed α-orientation X.*

3 Applications

Distributive lattices are beautiful and well understood structures and it is always nice to identify a distributive lattice on a finite set $\mathcal{C}$ of combinatorial objects. Such a lattice structure may then be exploited in theoretical and computational problems concerning $\mathcal{C}$.

Usually the cover relation in the lattice $\mathcal{L}_\mathcal{C}$ corresponds to some minor modification (*move*) in the combinatorial object. In our example the moves are reorientations of essential cycles (flips and flops). In most cases it is easy to find all legal moves that can be applied to a given object from $\mathcal{C}$. In our example finding the applicable moves corresponds to finding the directed essential cycles of an α-orientation. This task is easy in the sense that it can be accomplished in time polynomial in the size of the plane graph G. By the fundamental theorem of finite distributive lattices: there is a finite partially ordered set $P_\mathcal{C}$ such that the elements of $\mathcal{L}_\mathcal{C}$, i.e., the objects in $\mathcal{C}$, correspond to the order ideals (down-sets) of $P_\mathcal{C}$. The moves operating on the objects in $\mathcal{C}$ can be viewed as elements of $P_\mathcal{C}$. If $\mathcal{C}$ is the set of α-orientations the elements of $P_\mathcal{C}$ thus correspond to essential cycles, however, a single essential cycle may correspond to several elements of $P_\mathcal{C}$, Figure 2 illustrates this effect. The elements of $P_\mathcal{C}$ can be shown to be in bijection to the flips on a maximal chain from $X_{\max}$ to $X_{\min}$ in $\mathcal{L}_\mathcal{C}$. Consequently, in the case of α-orientations of G the order $P_\mathcal{C}$ has size polynomial in the size of G and can be computed in time polynomial in the size of G.

We explicitly mention three applications of a distributive lattice structure on a combinatorial set $\mathcal{C}$ before looking at some specific instances of Theorem 1.

- Any two objects in $\mathcal{C}$ can be transformed into each other by a sequence of moves. Proof: Every element of $\mathcal{L}_\mathcal{C}$ can be transformed into the unique minimum of $\mathcal{L}_\mathcal{C}$ by a sequence (chain) of moves. Reversing the moves in one of the two chains gives a transformation sequence for a pair of objects.
- All elements of $\mathcal{C}$ can be generated/enumerated with polynomial time complexity per object. The idea is as follows: Assign different priorities to the elements of $P_\mathcal{C}$. Use these priorities in a tree search (e.g., depth-first-search) on $\mathcal{L}_\mathcal{C}$ starting in the minimal element. An object is output/count only when visited for the first time, i.e., with the lexicographic minimal sequence of moves that generate it.
- To generate an element of $\mathcal{C}$ from the uniform distribution a Markov chain combined with the coupling from the past method can be used. This very elegant approach gives a process that stops itself in the perfect uniform distribution. Although this stop can be observed to happen quite fast in many processes of the described kind, only few of these processes have been analyzed satisfactorily. For more on this subject we recommend the work of Propp and Wilson [12] and [14].

3.1 Eulerian orientations

Let G be a plane graph, such that every vertex v has even degree $d(v)$. An Eulerian orientation of G is an orientation with $\mathsf{indeg}(v) = \mathsf{outdeg}(v)$ for every vertex v. Hence, Eulerian orientations are just the α-orientations with $\alpha(v) = \frac{d(v)}{2}$ for all $v \in V$. By Theorem 1 the Eulerian orientations of a planar graph form a distributive lattice.

To understand and work with the distributive lattice of Eulerian orientations of a plane graph it is useful to know the set of essential cycles. At first observe that there are no rigid edges, this follows from the fact that reversing all edges of an Eulerian orientation yields again an Eulerian orientation. This implies that all the essential cycles have to be face-cycles of bounded faces. To show that every bounded cycle is essential we note that Eulerian orientations can be constructed by iteratively orienting a cycle and removing it from the graph. This procedure can start with any face-cycle and each of its two orientations. This shows that for every face-cycle there are Eulerian orientations that difer just in the orientation of that face-cycle, i.e., face-cycles of bounded faces are essential.

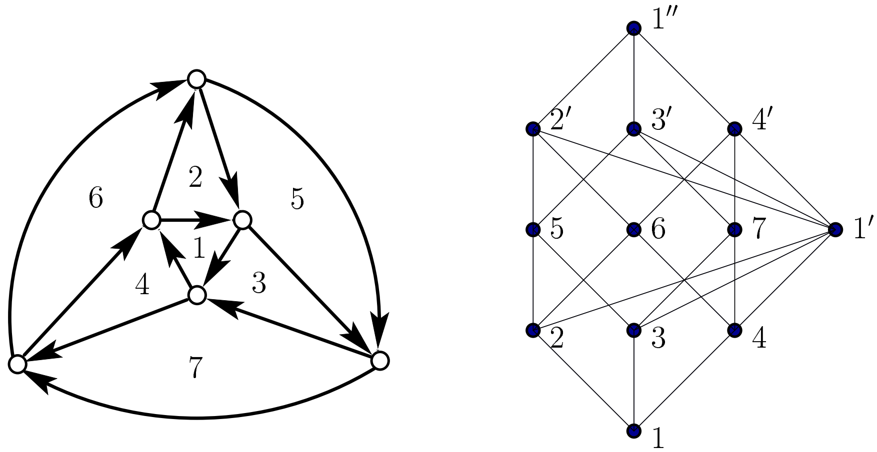

Figure 2: Left: A graph G with its minimal Eulerian orientation and a labeling of the faces. Right: The ordered set P such that the set of ideals of P is the lattice of Eulerian orientations of G.

3.2 The primal dual completion of a plane graph

For later applications we need the primal dual completion of a plane graph G. With G there is the dual graph G^*, the primal dual completion $\widetilde{G}$ of G is constructed as follows: Superimpose plane drawings of G and G^* such that only the corresponding primal dual pairs of edges cross. The completion $\widetilde{G}$ is obtained by adding a new vertex at each of these crossings. The construction is illustrated in Figure 3. If G has n vertices, m edges and f faces, then the corresponding numbers $\widetilde{n}$, $\widetilde{m}$ and $\widetilde{f}$ for $\widetilde{G}$ can be expressed as follows:

- $\widetilde{n} = n + m + f$. We denote the vertices of $\widetilde{G}$ originating in vertices of G, G^* and crossings of edges as *primal-vertices*, *dual-vertices* and *edge-vertices*.
- $\widetilde{m} = 4m$.
- $\widetilde{f} = 2m$: This follows since every face of $\widetilde{G}$ is a quadrangle with a primal- and a dual-vertex at opposite corners and edge-vertices at the remaining corners. Thus,

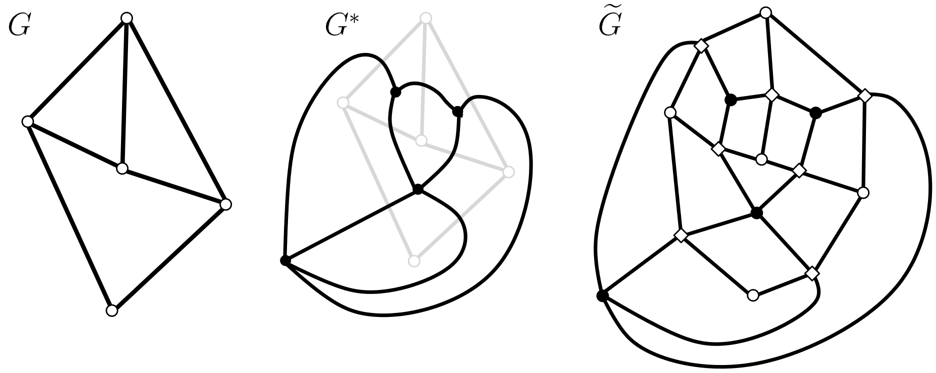

Figure 3: A plane graph G with its dual G^* and completion $\widetilde{G}$.

there is a bijection between angles of G and faces of $\widetilde{G}$. The number of angles of G is $\sum_v d(v) = 2m$.

There is a subtlety with the notion of the dual and, hence, of the completion when the connectedness of G is too small. If G has a bridge then $\widetilde{G}$ has multiple edges. In general, however, the completion is at least as well behaved as G:

- If G is connected and bridgeless $\implies$ $\widetilde{G}$ is 2-connected.
- If G is 2-connected $\implies$ $\widetilde{G}$ is 3-connected.

Completions of planar graphs have a nice characterization.

Proposition 1. *Let H be 2-connected plane graph, H is the completion of plane graph G iff the following three conditions hold:*

1. *All the faces of H are quadrangles, in particular H is bipartite.*
2. *In one of the two color classes of H all vertices have degree four.*

Proof. (sketch) It is immediate that the completion of a planar graph has the properties listed. For the converse first identify the edge-vertices as the color class of H consisting of degree four vertices. The other vertices are split into primal and dual vertices. Define two vertices as equivalent if they are opposite neighbours of an edge-vertex. If two vertices of a four-face fall into the same class, then there is a chain of 'opposite' vertices connecting them. Arguing with such a chain enclosing a minimum number of faces leads to a contradiction. Finally, show that the graph on one of the classes of vertices has H as completion. □

3.3 Spanning trees

We show that there is a bijection between the spanning trees of a planar graph $G = (V, E)$ and the α-orientations of the completion $\widetilde{G}$ of G for a certain α. Together with Theorem 1 this implies:

Theorem 2. *There is a distributive lattice of orientations of $\widetilde{G}$ which induces a distributive lattice on the spanning trees of a planar graph G.*

After having obtained this result we found that it was already known. Gilmer and Litherland [5] arrive at such a lattice on spanning trees in the context of knot theory. They also point out the equivalence to Kaufmann's *Clock Theorem*. Propp [13] describes a large class of distributive lattices related to orientations of graphs. If G is planar then the lattice of α-orientations of G is isomorphic to a Propp lattice of the dual G^*. Propp discovered lattices on spanning trees as a special case of his theory.

Let $T \subseteq E$ be the set of edges of a spanning tree of G. If T^* is the set of dual edges of non-tree edges (edges in $E \setminus T$), then T^* is the set of edges of a spanning tree of the dual graph G^*. This is the natural bijection between the spanning trees of G and G^*.

With a spanning tree T of G we associate an orientation of $\widetilde{G}$. First we select two special root vertices for $\widetilde{G}$, a primal-vertex v_r and a dual-vertex v_r^*. Now T and the corresponding dual tree T^* are thought of as directed trees in which every edge points towards the primal- respectively dual-root. The direction of edge $e = (u, w) \in T \cup T^*$ is passed on to the edges (u, v_e) and (v_e, w) in $\widetilde{G}$, where v_e is the edge-vertex of $\widetilde{G}$ corresponding to edge e. All the remaining edges of $\widetilde{G}$ are oriented so that they point away from their incident edge-vertex. Figure 4 illustrates the construction. The orientation thus

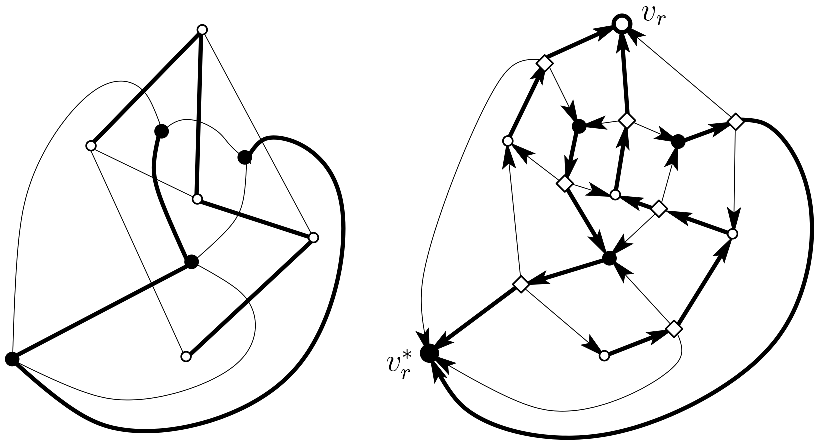

Figure 4: A pair of spanning trees for G and G^* and the corresponding orientation of the completion $\widetilde{G}$ with roots v_r and v_r^*.

obtained is an α_T-orientation for the following α_T:

- $\alpha_T(v_r) = 0$ and $\alpha_T(v_r^*) = 0$, i.e., the roots have outdegree zero.
- $\alpha_T(v_e) = 3$ for all edge-vertices v_e.
- $\alpha_T(v) = 1$ for all primal- and dual- non-root vertices v.

A pair of root vertices v_r and v_r^* is *legal* if both are incident to some face of $\widetilde{G}$.

Proposition 2. *The spanning trees of a planar graph G are in bijection to the α_T-orientations of $\widetilde{G}$ with a legal pair of root-vertices.*

Proof. We have described an orientation of $\widetilde{G}$ corresponding to a spanning tree of G. This orientation is an α_T-orientation and the mapping from spanning trees to α_T-orientations is injective.

Let X be an α_T-orientation, from X we obtain a set $S_X \subset E$ of edges as follows: At every edge-vertex v_e look at the unique incoming edge. Put e into S_X iff this incoming edge emanates from a primal-vertex. Since $\alpha_T(v) = 1$ for all primal-vertices, save the primal-root, the set S_X contains $n-1$ edges. We claim that S_X contains no cycles. Given the claim it follows that S_X is a spanning tree of G. Hence, the mapping from spanning trees to α_T-orientations is surjective and this completes the proof that the mapping is a bijection.

It remains to verify the claim that S_X contains no cycle. Suppose C is a cycle in S_X. Let $\widetilde{C}$ be the corresponding cycle in $\widetilde{G}$, i.e., for each edge $e = (u, w)$ in C there are two edges u, v_e and v_e, w in $\widetilde{C}$. We assume that $\widetilde{C}$ is a simple cycle and define the interior $\mathsf{Int}(\widetilde{C})$ so that the root vertices v_r and v_r^* are exterior, this can be done since the pair of roots is legal and $v_r \notin C$.

Let H be the subgraph of G obtained by eliminating all vertices and edges from G which are in the exterior of C, i.e., H consists of C and together with the part of G interior to C. Let the length of C be l and suppose that H has p vertices, $q+1$ faces and k edges. Eulers's formula for H implies that $p - k + (q+1) = 2$.

Consider the subgraph $\widetilde{H}$ of $\widetilde{G}$ obtained by eliminating all vertices and edges from G which are in the exterior of $\widetilde{C}$. Note that $\widetilde{H}$ has p primal-vertices, q dual-vertices, k edge-vertices and that the length of $\widetilde{C}$ is $2l$. We count the edges of $\widetilde{H}$ is two ways: Every edge-vertex has 4 incident edges in $\widetilde{G}$, each of the l edge-vertices on $\widetilde{C}$ has exactly one edge in the exterior of $\widetilde{C}$, hence $|E(\widetilde{H})| = 4k - l$. Counting the oriented edges of $\widetilde{H}$ at the vertices where they originate we count 1 for every primal- and dual-vertex (including the primal-vertices on $\widetilde{C}$!) and 3 for every interior edge-vertex, while the l edge-vertices on $\widetilde{C}$ only have outdegree two in $\widetilde{H}$. This makes $|E(\widetilde{H})| = p + q + 3k - l$. Subtracting the two expressions for $|E(\widetilde{H})|$ we obtain $p - k + q = 0$. This contradiction to the Euler formula for H completes the proof. □

Figure 5 shows the distributive lattice of the spanning trees of a graph with two different choices of the primal-root. The dual-root for both examples is the dual-vertex corresponding to the outer face.

When studying the set of spanning trees of a graph we loose nothing with the assumption that the graph is two edge connected, a bridge belongs to every spanning tree anyway. Let G be a simple planar two edge connected graph. Fix a primal- and a dual-root and consider the set of α_T-orientations of $\widetilde{G}$. It is easy to see that the only rigid edges are the edges incident to one of the roots which are necessarily oriented toward the root. Therefore, the set of essential cycles for α_T is contained in the set of faces of $\widetilde{G}$ which have no root vertex on the boundary. Actually, these two sets coincide. This simple characterization of essential cycles helps in understanding the lattice structure on the spanning trees of G. We illustrate this by explaining the cover relation $T \prec T'$ between two trees: The two trees only differ in one edge $T' = T - e + e'$ and there is a vertex $v \neq v_r$ such that e

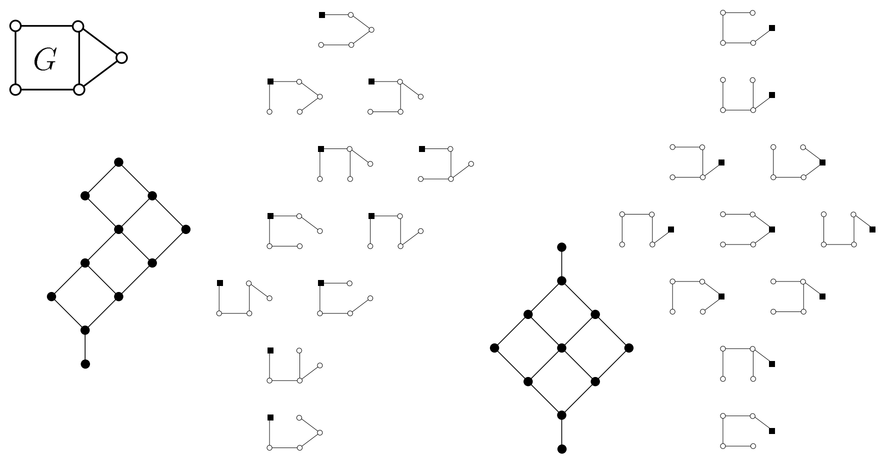

Figure 5: A graph and two distributive lattices for its spanning trees.

is the first edge of the $v \to v_r$ path in T and e' is the first edge of the $v \to v_r$ path in T'. Moreover, in the clockwise ordering of edges around v edge e' is the immediate successor of e and the angle between e and e' at v belongs to the interior of the unique cycle of $T + e'$ (this last condition is based on the choice of v_r^* as the dual-vertex corresponding to the unbounded face of G). The characterization is illustrated in Figure 6.

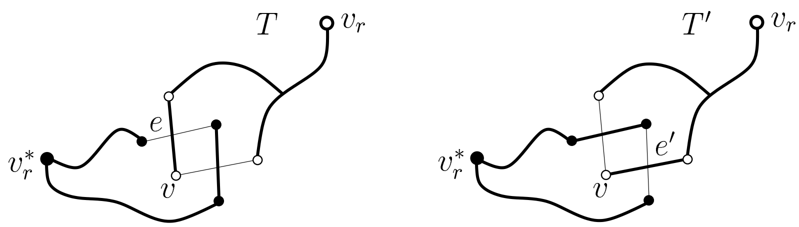

Figure 6: A typical flop between spanning trees $T \prec T'$ and their duals.

3.4 Matchings and f-factors

Given a function $f : V \to \mathbb{N}$ an *f-factor* of $G = (V, E)$ is a subgraph H of G such that $d_H(v) = f(v)$ for all $v \in V$. A perfect matching is a 1-factor, i.e, an f-factor for $f \equiv 1$. Let G be bipartite with bipartization (U, W). There is a bijection between f-factors of G and orientations of G with $\mathsf{outdeg}(u) = f(u)$ for $u \in U$ and $\mathsf{indeg}(w) = f(w)$ for $w \in W$. That is a bijection between f-factors and α_f-orientations where $\alpha_f(u) = f(u)$ for $u \in U$ and $\alpha_f(w) = d_G(w) - f(w)$ for $w \in W$. Together with Theorem 1 this implies:

Theorem 3. *For a plane bipartite graph G the distributive lattice of α_f–orientations of $\widetilde{G}$ induces a distributive lattice on the f-factors of G.*

Again, this result was already obtained by Propp [13] as a special case of his theory. Propp also points out that even in the case of perfect matchings there may exist rigid edges, i.e., edges that are in all perfect matching or in none. Therefore, there can be essential cycles which are not just faces of the graph.

The completion $\widetilde{G}$ of a plane graph G is a bipartite graph. Choose a primal-vertex v_r and a dual-vertex v_r^* from $\widetilde{G}$ and remove them, let $\widetilde{G}_r$ be the remaining graph. $\widetilde{G}_r$ has perfect matchings. The perfect matchings of $\widetilde{G}_r$ are in bijection with the spanning trees of G. A proof can be given by comparing the α-orientations of $\widetilde{G}$ corresponding to matchings and spanning trees. A special case of this bijection is Temperley's bijection between spanning trees and matchings of square grids. The general correspondence between trees and matchings has been exploited in [6]. Another recent source for the above theorem is [8].

3.5 Schnyder woods

Let G be a plane graph and let a_1, a_2, a_3 be three different vertices in clockwise order from the outer face of G. The *suspension* G^σ of G is obtained by adding a half-edge that reaches into the outer face to each of the three special vertices a_i. The *closure* G^σ_∞ of a suspension G^σ is obtained by adding a new vertex v_∞, this vertex is used as second endpoint for the three half-edges of G^σ.

Schnyder [15], [16] introduced edge orientations and equivalent angle labelings for planar triangulations. He used this structures for a remarkable characterization of planar graphs in terms of order dimension. The incidence order P_G of a graph $G = (V, E)$ is the order on $V \cup E$ with relations $v < e$ iff $v \in V$, $e \in E$ and $v \in e$. Schnyder proved: A graph G is planar $\iff$ the dimension of its incidence order is at most 3. Another important application of Schnyder's labelings is a proof that every planar n vertex graph admits a straight line drawing on the $(n-1) \times (n-1)$ grid.

De Fraysseix and de Mendez [4] prove a bijection between Schnyder labelings of a planar triangulation G and 3-orientations of G^σ_∞, i.e., α-orientations with $\alpha(v) = 3$ for every regular vertex and $\alpha(v_\infty) = 0$. Based on the bijection with 3-orientations de Mendez [10] and Brehm [1] have shown that the set of Schnyder labelings of a planar triangulation G has the structure of a distributive lattice. This result stimulated the research that lead to Theorem 1. The proof of the general theorem given in the first part of this paper is widely based on ideas that are already contained in the cited proofs of the special case.

In [2] the concept of Schnyder labelings was generalized to 3-connected planar graphs. It was also shown that like the original concept the generalization yields strong applications in the areas of dimension theory and graph drawing. The following definition of Schnyder woods is taken from [3] where it is also shown that Schnyder woods and Schnyder labelings are in bijection.

Let G^σ be the suspension of a 3-connected plane graph. A *Schnyder wood* rooted at a_1, a_2, a_3 is an orientation and labeling of the edges of G^σ with the labels $1, 2, 3$ (alternatively: red, green, blue) satisfying four rules. On the labels we assume a cyclic structure so that $i + 1$ and $i - 1$ is defined for all i.

(W1) Every edge e is oriented by one or two opposing directions. The directions of edges are labeled such that if e is bioriented the two directions have distinct labels.

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

(W3) Every vertex v has one outgoing edge in each label. The edges e_1, e_2, e_3 leaving v in labels 1,2,3 occur in clockwise order. Each edge entering v in label i enters v in the clockwise sector from e_{i+1} to e_{i-1}. (See Figure 7).

(W4) There is no interior face whose boundary is a directed cycle in one label.

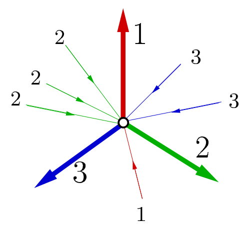

Figure 7: Edge orientations and labels at a vertex.

Unlike in the case of planar triangulations, the labeling of edges of a Schnyder wood can not be recovered from the underlying orientation, Figure 8 shows an example. However, orientations of an appropriate primal dual completion of a suspended plane graph are in bijection to Schnyder woods (this will be shown in Proposition 4).

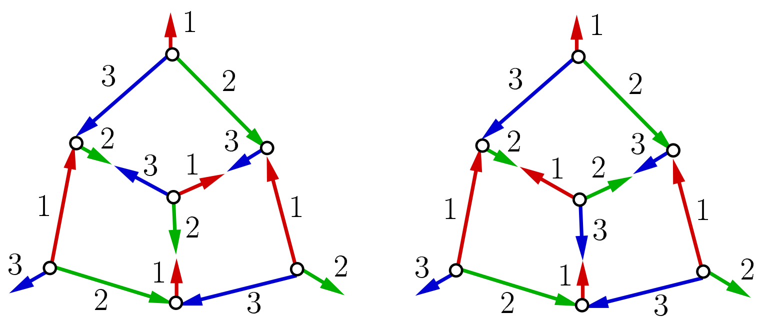

Figure 8: Two different Schnyder woods with the same underlying orientation.

We first show that Schnyder woods of a suspended plane graph are in bijection with Schyder woods of a properly defined dual. Figure 9 exemplifies the duality. Actually, the figure illustrates much more: With the primal and dual graphs and Schnyder woods it also shows a corresponding orthogonal surface. We include this figure for two reasons. The duality between primal and dual Schnyder woods becomes nicely visible on the surface. Moreover, it was in this context of geodesic embeddings of planar graphs on orthogonal surfaces that the duality was first observed by Miller [9]. For details on geodesic embeddings and the connections with Schnyder woods we refer to [9] and [3].

Recall that the definition of the suspension G^σ of a plane graph G was based on the choice of three vertices a_1, a_2, a_3 in clockwise order from the outer face of G. The *suspension dual* G^{σ^*} is obtained from the dual G^* by some surgery: The dual-vertex corresponding to the unbounded face is replaced by a triangle with vertices b_1, b_2, b_3 which are the three special vertices for the Schnyder woods on G^{σ^*}. More precisely, let A_i be

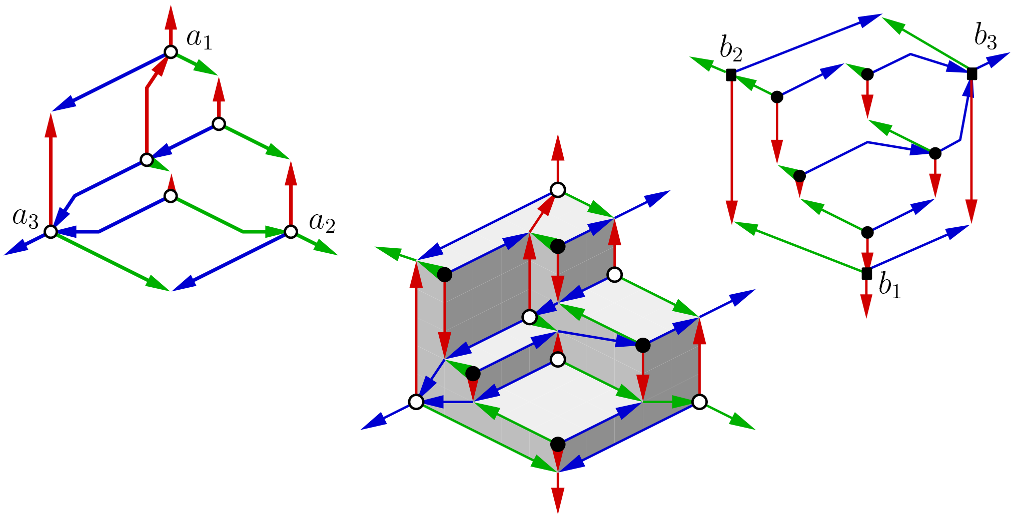

Figure 9: A suspended graph G^σ with a Schnyder wood, a corresponding embedding and the dual Schnyder wood.

the set of edges on the arc of the outer face of G between vertices a_j and a_k, with $\{i,j,k\} = \{1,2,3\}$. Let B_i be the set of dual edges to the edges in A_i, i.e., $B_1 \cup B_2 \cup B_3$ is the set of edges containing the vertex v^*_∞ of G^* which corresponds to the unbounded face of G. Exchange v^*_∞ by b_i at all the edges of B_i and add three edges $\{b_1, b_2\}, \{b_2, b_3\}, \{b_1, b_3\}$ and an half-edge that reaches into the triangle face $\{b_1, b_2, b_3\}$ to each of the special vertices b_i. The resulting graph is the suspension dual G^{σ^*} of G. Figure 9 illustrates the definition.

Proposition 3. *Let G^σ be a suspension of a 3-connected plane graph G. There is a bijection between the Schnyder woods of G^σ and the Schnyder woods of the suspension dual G^{σ^*}.*

Proof. In [3] it is shown that Schyder woods of G^σ are in bijection to Schnyder angle labelings. These are labelings of the angles of G^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) *Rule of vertices:* The labels of the angles at each vertex form, in clockwise order, a nonempty interval of 1's, a nonempty interval of 2's and a nonempty interval of 3's.

(A3) *Rule of faces:* The labels of the angles at each interior 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.

There is an obvious one-to-one correspondence between the angles of G^σ and the inner angles of G^{σ^*}. This correspondence yields an exchange between the rule of vertices and the rule of faces. Therefore, any Schnyder labeling of G^σ is a Schnyder labeling of G^{σ^*} and vice versa. This is exemplified in Figure 10. □

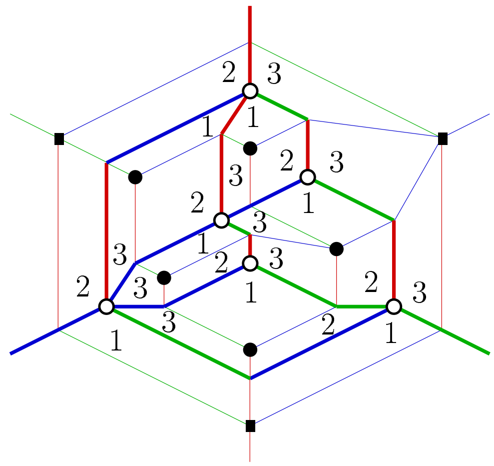

Figure 10: Bold edges show a suspended graph G^σ, light edges correspond to G^{σ^*}. The Schnyder angle labeling shown is valid for both graphs.

Notable is the connection between Schnyder labelings and orthogonal embeddings: The angle labeling in Figure 10 corresponds to the shades in the orthogonal embedding of the same graph in Figure 9.

We now define the completion of a plane suspension G^σ and its dual G^{σ^*}. Superimpose G^σ and G^{σ^*} so that exactly the primal dual pairs of edges cross (the half edge at a_i has a crossing with the dual edge $\{b_j, b_k\}$, for $\{i,j,k\} = \{1,2,3\}$). The common subdivision of each crossing pair of edges by a new edge-vertex gives the *completion* $\widetilde{G^\sigma}$. The completion $\widetilde{G^\sigma}$ is planar and has six half-edges reaching into the unbounded face. Similar to the closure of a suspension we define the closure $\widetilde{G^\sigma_\infty}$ of $\widetilde{G^\sigma}$ by adding a new vertex v_∞ which is the second endpoint of the six half-edges.

A pair of corresponding Schnyder woods on G^σ and G^{σ^*} induces an orientation of $\widetilde{G^\sigma_\infty}$ which is an α-orientation for the following α_S:

- $\alpha_S(v) = 3$ for all primal- and dual-vertices v.
- $\alpha_S(v_e) = 1$ for all edge-vertices v_e.
- $\alpha_S(v_\infty) = 0$ for the special closure vertex v_∞.

The outdegree of a primal- or dual-vertex is three by Schnyder wood axiom (W3). The outdegree of v_∞ is zero since all six half edges of $\widetilde{G^\sigma}$ are directed outwards (W2). For the edge-vertices we first recall the translation from a Schnyder labeling to a Schnyder wood as given in [2]: If an edge has different angular labels i and j at its primal- or dual-endvertex it is oriented away from this vertex in label k. Lemma 1 of [2] says that in a Schnyder labeling the four angles around an edge are labeled $i, i-1, i+1, i$ in clockwise order for some $i \in \{1,2,3\}$. This shows that at the corresponding edge-vertex v_e only the edge between the two i labeled angles is outgoing, the other three edges are incoming, i.e., $\alpha_S(v_e) = 1$.

Proposition 4. *The Schnyder woods of a planar suspension G^σ are in bijection with α_S-orientations of $\widetilde{G^\sigma_\infty}$.*

Proof. With a Schnyder wood of G^σ we have a Schnyder labeling of angles and thus a dual Schnyder wood on G^{σ^*}. As already shown this pair induces an α_S-orientation of $\widetilde{G^\sigma_\infty}$.

It remains to recover a unique pair of Schnyder woods from any given α_S-orientation of $\widetilde{G^\sigma_\infty}$. The orientation of edges is clearly given, but for Schnyder woods the edges also need labels. Given an edge e we define the *straight path* of e by a simple rule.

Straight path rule: Upon entering an edge-vertex on edge e_h continue on the other side, i.e., traverse the other half of the underlying primal or dual edge. When entering a primal- or dual-vertex v on e_h and e_h is directed towards v continue with the opposite outgoing edge. If e_h is outgoing at v the continuation depends on the outgoing edge at the edge-vertex v_e on the other side of e_h. If the outgoing edge at v_e points to the right of the straight path choose the right outgoing edge at v and if the outgoing edge at v_e points to the left continue the straight path with the left outgoing edge at v. The rule is illustrated in Figure 11.

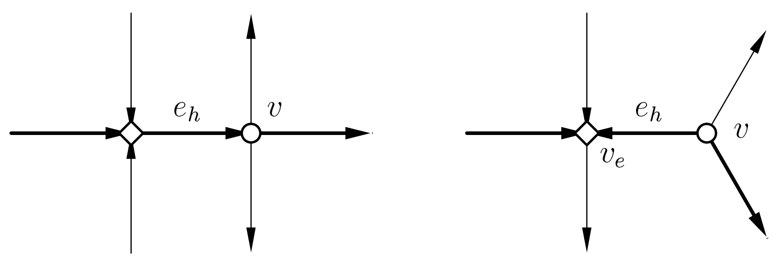

Figure 11: Illustrating the straight path rule, bold edges indicate the straight path.

A straight path may reach the special vertex v_∞. There is no continuation at this vertex and the straight path ends. The next lemma shows that every straight path ends that way.

Lemma 15. *Let X be an α_S-orientation of $\widetilde{G^\sigma_\infty}$ and $e \in X$ be a directed edge. The straight path whose first edge is e leads to v_∞ where it ends.*

Proof. Suppose some straight path does not end at v_∞ then it has to run in into a cycle. In that case we find a simple cycle $v_1, v_2, \ldots, v_{2k}, v_1$ as part of the straight path starting with $e_h = (v_1, v_2)$. Every second vertex on the cycle is an edge-vertex, omitting these vertices yields a simple cycle C' in G^σ or in G^{σ^*}. Assume that C' is a cycle in G^σ and let H be the planar graph induced by C' and its interior (the interior of a cycle is defined relative to v_∞ which is exterior). H has k vertices on its exterior cycle, let r be the number of inner vertices and s be the number of inner faces of H. Hence, $|V(H)| = k + r$ and $|F(H)| = s + 1$. By Euler's Formula $|E(H)| = k + r + s - 1$.

Let $\widetilde{H}$ be the graph induced by C and its interior in $\widetilde{G^\sigma_\infty}$. Since this graph is closely related to H we know that it has $k + r + s - 1$ edge-vertices. The degree of k of these edge-vertices is only 3 since they sit on the outer cycle C of $\widetilde{H}$. Counting edges of $\widetilde{H}$ at edge-vertices shows that this number is $|E(\widetilde{H})| = 3k + 4(r + s - 1)$.

A second count of the edges of $\widetilde{H}$ can be done by considering the orientation X and counting the edges at their source vertices. $\widetilde{H}$ has $r+s$ interior primal- and dual-vertices, each of these has outdegree 3 in the α_S-orientation X. There are $r + s - 1$ interior edge-vertices, each with outdegree 1. Therefore, the number of edges of $\widetilde{H}$ emanating from interior vertices is $4(r+s)-1$. On the outer cycle C there are $2k$ edges. It remains to count those edges pointing from a vertex on cycle C to the interior. Looking at Figure 11 we observe that together an edge-vertex and its subsequent primal-vertex on C have exactly one edge pointing into the interior of C. Altogether there are k such pairs on C. The total number of edges thus is $|E(\widetilde{H})| = 4(r + s) - 1 + 3k$.

Comparing the two counts for $|E(\widetilde{H})|$ we obtain the contradiction $-4 = -1$. This shows that the straight path rule never leads into a cycle and proves the lemma. □

The lemma is the basis for a simple rule for assigning a label (color) to every edge of an α_S-orientation of $\widetilde{G^\sigma_\infty}$. With the directed edge e consider the straight path. The last vertex visited by this path before its end in v_∞ is one of a_1, a_2, a_3, in case the underlying edge of e is a primal edge, or one of b_1, b_2, b_3, if the underlying edge of e is a dual edge. Take the index of this last vertex as the label for e, e.g., if the straight path starting with e has (a_2, v_∞) as its last edge the label of e is 2.

The claim is that mapped back to G^σ and G^{σ^*} this labeling and orientation gives a pair of Schnyder woods on these graphs. For the verification of the axioms for Schnyder woods we need another lemma.

Lemma 16. *Let X be an α_S-orientation of $\widetilde{G^\sigma_\infty}$. If x is a primal- or dual-vertex and p and q are two straight paths leaving x on different edges then p and q meet only at v_∞.*

Proof. We assume that x is a primal-vertex. Suppose that after leaving x the straight path p shares a vertex $y \neq v_\infty$ with q. Let y be the first such vertex on p. With an edge-vertex every straight path also contains both adjacent primal-vertices. Therefore, y has to be a primal-vertex.

Corresponding to p and q there are paths p' and q' from x to y in G^σ. The arcs from x to y on p' and q' bound an interior region R. As in the proof of Lemma 15 we consider the restriction G_R of G^σ to R. Let r be the number of inner vertices, s be the number of inner faces and k be the number of vertices on the exterior cycle of G_R. Hence, $|V(G_R)| = k+r$ and $|F(G_R)| = s + 1$. By Euler's Formula $|E(G_R)| = k + r + s - 1$.

For the subgraph $\widetilde{G_R}$ of $\widetilde{G^\sigma_\infty}$ corresponding to G_R we find $|E(\widetilde{G_R})| = 3k+4(r+s-1)$. This count is obtained by considering edge-vertices and their incident edges in $\widetilde{G_R}$.

A second count of the edges of $\widetilde{G_R}$ is done by counting the edges at their source vertices in the orientation X. This is again done as in the proof of Lemma 15. The difference is that on the arcs of p' and q' between x and y we may only find $k - 2$ pairs consisting of an edge-vertex and the subsequent primal-vertex on the straight path. Still we get the estimate $|E(\widetilde{G_R})| \geq 4(r + s) - 1 + 3k - 2$.

The two counts for $|E(\widetilde{G_R})|$ yield the contradiction $-4 \geq -3$. □

We are ready now to complete the proof of Proposition 4 by verifying the Schnyder wood axioms. The most interesting axiom is (W3): Since X is an α_S-orientation a vertex v has exactly three outgoing edges. The labeling rule and Lemma 16 imply that the labels of these three edges are different. Since the special vertices a_1, a_2, a_3 are clockwise on the outer cycle of G it follows, again from Lemma 16, that the edges e_1, e_2, e_3 leaving v in labels 1,2,3 occur in clockwise order. The labeling rule, which was derived from the straight path rule, implies that an edge entering v in the sector between e_{i+1} and e_{i-1} has label i. This shows that (W3) is valid.

Axiom (W2) is trivial. The orientation part of (W1) is obvious. The labels of a bidirected edge are different, otherwise the incident vertices would have at least two outgoing edges in one label, a contradiction to (W3). A face whose border is directed in one label would imply a contradiction to Lemma 15, therefore, (W4) must also hold. This completes the proof of the proposition. $\square$

Combining Proposition 4 and Theorem 1 we obtain the main result of this section.

Theorem 4. *The set of Schnyder woods of a planar suspension G^σ form a distributive lattice.*

In the case of Schnyder woods a full characterization of all possible essential cycles seems to be a complex task. Unlike in the case of spanning trees or Eulerian orientations it is not enough to consider faces of $\widetilde{G^\sigma_\infty}$ as candidates for essential cycles.

We first discuss the case of a planar triangulation G. In this case all the inner vertices of the suspension dual G^{σ^*} are of degree 3. Therefore, all the edges of $\widetilde{G^\sigma_\infty}$ that correspond to G^{σ^*} are rigid, they point from the dual-vertex to the edge-vertex. This very special property made it possible to investigate the lattice of Schnyder woods of a planar triangulation without reference to duality (see [10] and [1]). The usual essential cycles in this case are triangular faces of G, this corresponds to the union of three faces of $\widetilde{G^\sigma_\infty}$ sharing a vertex of degree 3. For 4-connected triangulations these are all essential cycles. However, if G has separating triangles these also are essential cycles.

If G is not triangulated the structure of essential cycles for the α_S-orientations of $\widetilde{G^\sigma_\infty}$ can be more complex, in Figure 12 we display three examples. From the shown examples further examples can be obtained with the following surgery: Choose a degree three vertex v in the interior of the essential cycle. Remove v and its three edges this leaves an empty 6-gon T. Choose a completed plane graph $\widetilde{H^\sigma}$ and paste this graph into the triangular face by identifying the outer 6-gon of $\widetilde{H^\sigma}$ with T such that edge-vertices go on edge-vertices. With the next lemma we show that still in some sense the essential cycles cannot be too complicated.

Lemma 17. *Let G^σ be suspended plane graph. The possible length of essential cycles for α_S-orientations of $\widetilde{G^\sigma_\infty}$ are 4, 6, 8, 10 and 12.*

Proof. Let C be an essential cycle. Since $\widetilde{G^\sigma_\infty}$ is bipartite every second vertex on C is an edge-vertex, therefore, the length $|C|$ of C is even. If $|C| = 4$ then C is the boundary of a face. If $|C| > 4$ then the interior of C contains vertices. If C is essential all the

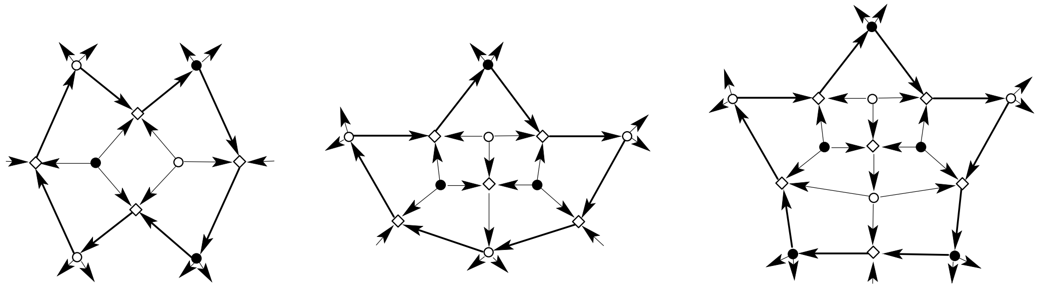

Figure 12: Bold edges show non-trivial essential cycles for α_S-orientations.

edges in the interior cut $E_{\mathsf{Cut}}[I_C]$ are directed towards C: Suppose an edge e in $E_{\mathsf{Cut}}[I_C]$ is directed away from C. The straight path starting with e will eventually leave C thus giving a chordal path for C. By Lemma 4 this is in contradiction to C being essential.

Consider an edge-vertex v_e on C. We claim that either one ore two edges in $E_{\mathsf{Cut}}[I_C]$ are incident to v_e. Suppose that $E_{\mathsf{Cut}}[I_C]$ contains no edge incident to v_e. Consider the 4-face F interior of C with v_e as one of its corners. Let v' be the edge-vertex diagonally opposite to v_e on F. Recall that $\alpha_S(v') = 1$ and all edges in $E_{\mathsf{Cut}}[I_C]$ are directed towards C. It is impossible to satisfy these two conditions simultaneously.

We can thus classify the edge-vertices on C as either *straight*, if both neighbors on C are primal-vertices or both are dual-vertices, or *reflex*, if they have two incident edges in the interior cut of C. Let a be the number of reflex vertices on C.

As in previous proofs we consider the graph H obtained by restricting $\widetilde{G^\sigma_\infty}$ to C and its interior. Let $|C| = 2k$ and let there be r primal-vertices, s dual-vertices and t edge-vertices in the interior of C. Considering edges of H at their source in an α_S-orientation we find that $|E(H)| = 3(r+s)+t+2k$. Since $|V(H)| = r+s+t+2k$ Euler's formula implies that the number of faces of H is $2(r+s)+2$. Excluding the unbounded face the faces of H are 4-faces, therefore, $2|E(H)| = 4(2(r+s)+1)+2k$. Yet another count of edges of H is by counting them at their incident edge-vertex. Since there are a reflex vertices on C this gives $|E(H)| = 4t+3k+a$. Elementary algebra yields $2k = a+6$. Together with the obvious inequality $a \leq k$ this gives $k \leq 6$ and completes the proof. $\square$

Acknowledgement

Thanks to Patrice Ossona de Mendez for valuable conversations which got me thinking about the possibility of generalizing the result from Schnyder woods to general orientations of planar graphs.

References

[1] E. Brehm, *3-orientations and Schnyder 3–tree–decompositions*, Diplomarbeit, Freie Universität Berlin, 2000.
`www.inf.fu-berlin.de/~felsner/Diplomarbeiten/brehm.ps.gz`

[2] S. Felsner, *Convex drawings of planar graphs and the order dimension of 3-polytopes*, Order 18 (2001), pp. 19–37.

[3] S. Felsner, *Geodesic embeddings and planar graphs*, Order 20 (2003), pp. 135–150.

[4] H. de Fraysseix and P.O. de Mendez, *On topological aspects of orientations*, Discr. Math. 229 (2001), pp. 57–72.

[5] P.M. Gilmer and R.A. Litherland, *The duality conjecture in formal knot theory*, Osaka J. Math. 23 (1986), pp. 229–247.

[6] R. Kenyon, J. Propp and D.B. Wilson, *Trees and matchings*, Electr. J. Comb. 7, No.1, R25, 34 p (2000).

[7] S. Khuller, J. Naor and P. Klein, *The lattice structure of flow in planar graphs*, SIAM J. Discr. Math. 6 (1993), pp. 477–490.

[8] P.C.B. Lam and H. Zhang, *A distributive lattice on the set of perfect matchings of a plane bipartite graph*, Order 20 (2003), pp. 13–29.

[9] E. Miller, *Planar graphs as minimal resolutions of trivariate monomial ideals*, Documenta Math. 7 (2002), pp. 43–90.

[10] P.O. de Mendez, *Orientations bipolaires*, PhD Thesis, Paris, 1994.

[11] O. Pretzel, *On reorienting graphs by pushing down maximal vertices*, Order 3 (1986), pp. 135–153.

[12] J. Propp, *Generating random elements of finite distributive lattices*, Electr. J. Comb. 4, No.2, R15, 12 p. (1997).

[13] J. Propp, *Lattice structure for orientations of graphs*, Manuscript 1993. `www.math.wisc.edu/~propp/orient.html`

[14] J. Propp and D.W. Wilson, *Exact sampling with coupled Markov chains and applications to statistical mechanics*, Rand. Struct. and Alg. 9 (1996), pp. 223–252.

[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.