Algorithmica (2007) 47: 399–420
DOI: 10.1007/s00453-006-0177-6

Algorithmica

Convex Drawings of 3-Connected Plane Graphs[1]

Nicolas Bonichon,[2] Stefan Felsner,[3] and Mohamed Mosbah[2]

Abstract. We use Schnyder woods of 3-connected planar graphs to produce convex straight-line drawings on a grid of size $(n-2-\Delta)\times(n-2-\Delta)$. The parameter $\Delta \geq 0$ depends on the Schnyder wood used for the drawing. This parameter is in the range $0 \leq \Delta \leq n/2-2$. The algorithm is a refinement of the face-counting algorithm; thus, in particular, the size of the grid is at most $(f-2)\times(f-2)$.

The above bound on the grid size simultaneously matches or improves all previously known bounds for convex drawings, in particular Schnyder's and the recent Zhang and He bound for triangulations and the Chrobak and Kant bound for 3-connected planar graphs. The algorithm takes linear time.

The drawing algorithm has been implemented and tested. The expected grid size for the drawing of a random triangulation is close to $\frac{7}{8}n \times \frac{7}{8}n$. For a random 3-connected plane graph, tests show that the expected size of the drawing is $\frac{3}{4}n \times \frac{3}{4}n$.

Key Words. Graph drawing, Schnyder wood, Plane graph, Convex drawing, Compact representation.

1. Introduction.

We investigate crossing-free straight-line drawings of planar graphs with the restriction that the vertices of the graph have to be located at integer grid points. The aim is to keep the area of an axis-aligned rectangle that covers the drawing as small as possible. It is known that a square of side length $n-2$, i.e., an $(n-2)\times(n-2)$ grid is enough to host every planar graph.

A drawing with the property that the boundary of every face (including the outer face) is a convex polygon is called a convex drawing. Convex drawings exist for every 3-connected planar graph. Again the aim is to keep the area of such a drawing as small as possible.

It is important to distinguish between convex drawings and strictly convex drawings. A drawing is strictly convex if every interior angle is less than 180° and every outer angle greater than 180°. In this paper we deal with convex drawings. The grid size for strictly convex drawings was recently studied by Rote [18]; he proves that an $O(n^{7/3}) \times O(n^{7/3})$ grid is enough for strictly convex drawings of planar graphs with n vertices. The construction is based on a convex drawing obtained via Schnyder woods.

1.1. *Previous Work*.

The question of whether every planar graph has a straight-line embedding on a grid of polynomial size was raised by Rosenstiehl and Tarjan [17]. Unaware of the problem Schnyder [20] constructed a barycentric representation which

[1] An extended abstract of this research was contributed to the *Proceedings of Graph Drawing* '04 [1].
[2] LaBRI, Université Bordeaux-1, 351 cours de la Libération, 33405 Talence Cedex, France. {bonichon,mosbah}@labri.fr.
[3] Technische Universität Berlin, Institut für Mathematik, MA 6-1, Straße des 17. Juni 136, 10623 Berlin, Germany. felsner@math.tu-berlin.de.

Received December 1, 2004; revised November 5, 2005. Communicated by J. Pach and F. Shahrokhi.
Online publication February 13, 2007.

translates to an embedding on the $(2n - 5) \times (2n - 5)$ grid. The first explicit answer to the question was given by de Fraysseix et al. [5], [6]. They construct straight-line embeddings on a $(2n - 4) \times (n - 2)$ grid and show that the embedding can be computed in $O(n \log n)$. De Fraysseix et al. also observed a lower bound of $(\frac{2}{3}n - 1) \times (\frac{2}{3}n - 1)$ for grid embeddings of the n vertex graph containing a nested sequence of $n/3$ triangles. It is conjectured that this is the worst case, i.e., that every planar graph can be embedded on the $(\frac{2}{3}n - 1) \times (\frac{2}{3}n - 1)$ grid. 4-Connected planar graphs with at least four vertices on the outer face can be drawn even more compactly. Work of He [12] and Miura et al. [16] shows that these graphs can be embedded on the $\frac{n}{2} \times \frac{n}{2}$ grid.

In his second paper [21] Schnyder proves the existence of an embedding on the $(n - 2) \times (n - 2)$ grid that can be computed in $O(n)$ time. In general Schnyder's result from [21] is still unbeaten. Lately, Zhang and He [26] used the minimum Schnyder wood of a triangulation to prove a bound of $(n - 1 - \Delta^{\circleddash}) \times (n - 1 - \Delta^{\circleddash})$, where $\Delta^{\circleddash}$ is the number of cyclic faces in the minimum Schnyder wood.

Though it is implicitly contained in Steinitz's characterization of 3-connected planar graphs as the skeleton graphs of three-dimensional polytopes, the existence of convex drawings for these graphs is known as Tutte's theorem. The idea for Tutte's proof [24], [25] is known as *spring-embedding*. Technically the embedding is obtained as a solution to a system of linear equations. Kant [13] has extended the approach of de Fraysseix et al. to construct convex drawings on the $(2n - 4) \times (n - 2)$ grid. The grid size was reduced to $(n - 2) \times (n - 2)$ by Chrobak and Kant [4]. Schnyder and Trotter [22] have worked on ideas for convex grid embeddings which are based on Schnyder woods. The basic approach was independently worked out by Di Battista et al. [7] and Felsner [8]. This results in convex grid drawings on the $(f - 1) \times (f - 1)$ grid, where f is the number of faces of the graph. In this paper this basic algorithm is used but the size of the required grid is reduced by some new ideas. Loosely speaking, some edges are eliminated which results in the reduction of f. This can be done until at most $n - \Delta$ faces remain. The eliminated edges can be reinserted in the resulting drawing on the $(n - 1 - \Delta) \times (n - 1 - \Delta)$ grid, with $\Delta \geq 0$. $\Delta \geq n - f$. The drawing procedure can be implemented to run in linear time. The algorithm has been implemented and integrated into the PIGALE library.[4]

1.2. *Organization of the Paper.* In the next section we introduce Schnyder woods. It is shown how to use Schnyder woods to obtain convex drawings of 3-connected planar maps. The lattice of Schnyder woods is discussed and a new operation called *merge* is introduced as a tool for transforming Schnyder woods and their underlying graphs.

Section 3 contains the generic drawing algorithm. It is shown that this algorithm produces convex drawings and the size of the grid required for the drawing is analyzed. The main ingredient of this analysis is a bound on the number of merges applicable to a Schnyder wood. In particular it is shown that starting with the Schnyder wood of a triangulation a sequence of $n - 4 + \Delta^{\circleddash} - \Delta^{\oplus}$ merge operations is admissible.

Section 4 adds some ideas for further reduction of the grid size. The first of these allows a decrease of the side length of the grid by one. This small reduction, however,

[4] PIGALE is an open-source library in which numerous planar graphs algorithms are implemented. It is developed and maintained by H. de Fraysseix and P. Ossona de Mendez. http://pigale.sourceforge.net.

is crucial to match Schnyder's $(n-2) \times (n-2)$ bound for planar triangulations. We present a second idea for further reducing the grid size. Basically, the improvement comes from disregarding some faces. Although the technique is appealing it has so far resisted our attempts of proving that it guarantees some non-zero gain. We adapt the method to produce compact convex drawings in the slightly more general case of internally 3-connected planar graphs.

Finally, we report some experimental results. Tests with the implementation allow guessing the average reduction in size obtained from the parameter $\Delta^{\ominus}$ or from disregarding some faces.

2. Schnyder Woods. Schnyder defined special colorings and orientations of the internal edges of a triangulation. In [20] and [21] he applied these Schnyder woods to characterize planar graphs via order dimension and to draw planar graphs on small grid sizes. Here we describe a generalization of Schnyder woods for 3-connected planar graphs. Such a generalization has been presented in [7] and [8]; in our exposition we follow [9].

A *planar map* M is a simple planar graph G together with a fixed planar embedding of G in the plane. A *suspension* M^σ of M is obtained as follows: Three different vertices from the outer face of M are specified and named a_1, a_2, a_3 in clockwise order. (For ease of visualization we identify the indices 1, 2, 3. Moreover, we assume a cyclic structure on the indices such that $i+1$ and $i-1$ are always defined.) At each of the three special vertices a_i, called *suspension vertices*, a half-edge reaching into the outer face is attached.

Let M^σ be a suspension of a planar map. A *Schnyder wood* is an orientation and coloring of the edges of M^σ with the colors 1, 2, 3 satisfying the following rules:

(W1) Every edge e is oriented by one or two opposite directions. The directions of edges are colored such that if e is bi-directed the two directions have distinct colors.
(W2) The half-edge at a_i is directed outward 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}. See Figure 1.
(W4) There is no interior face whose boundary is a directed cycle in one color.

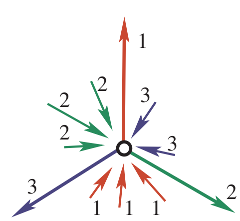

Fig. 1. Edge colorings[5] and orientations at a vertex.

[3] To see the colors visit the electronic versions at the authors' homepages.

A suspension is *internally* 3-*connected* if adding a new vertex v_∞ as the second endpoint for the three half-edges the graph obtained is planar and 3-connected.

FACT 1. *There is a Schnyder wood for* M^σ, *if and only if* M^σ *is the suspension* M^σ *internally* 3-*connected*.

The proof that only internally 3-connected suspensions admit a Schnyder wood is given by Theorem 5.1 of [15].

Given a Schnyder wood, let T_i be the set of edges colored i with the direction they have in this color. Since every internal vertex has outdegree one in T_i every v is the starting vertex of a unique i-path $P_i(v)$ in T_i.

FACT 2. *The digraph* T_i *is acyclic*; *even more*, T_i *is a tree with root* a_i.

2.1. *Convex Drawings via Face-Counting*. Schnyder and Trotter [22] had some ideas of using Schnyder woods for convex grid embeddings. The approach has been worked out in [7] and [8]. We describe the technique, omitting some details.

From the vertex condition (W3) it can be deduced 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 2.

FACT 3.

(a) $R_i(u) \subseteq R_i(v)$ *iff* $u \in R_i(v)$.
(b) $R_i(u) = R_i(v)$ *iff there is a path of bicolored edges in colors* $i-1$ *and* $i+1$ *connecting* u *and* v.
(c) *For all* u, v *there are* i *and* j *with* $R_i(u) \subset R_i(v)$ *and* $R_j(v) \subset R_j(u)$.

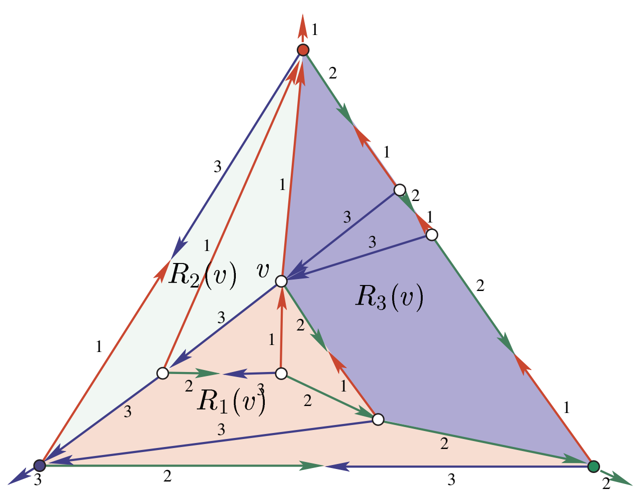

Fig. 2. A Schnyder wood and the regions of vertex v.

The *face-count* of a vertex v is the vector (v_1, v_2, v_3), where v_i is defined as

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

FACT 4. *For every edge* $\{u, w\}$ *and vertex* $v \neq u, w$ *there is a color* i *with* $\{u, w\} \in R_i(v)$; *hence*, $u_i \leq v_i$ *and* $w_i \leq v_i$.

Inclusion properties of the three regions of adjacent vertices imply:

FACT 5.

(a) *If edge* (u, v) *is uni-directed in color* i, *then* $u_i < v_i$, $u_{i-1} > v_{i-1}$ *and* $u_{i+1} > v_{i+1}$.
(b) *If* (u, v) *is directed in color* $i-1$ *and* (v, u) *in color* $i+1$, *then* $u_i = v_i$, $u_{i-1} > v_{i-1}$ *and* $u_{i+1} < v_{i+1}$.

Clearly, each vertex v has $v_1 + v_2 + v_3 = f - 1$, where f is the number of faces of M. Hence, we have a mapping of the vertices of the graph to the plane $T_f = \{(x_1, x_2, x_3)\colon x_1 + x_2 + x_3 = f - 1\}$ in $\mathbb{R}^3$. Connecting the points corresponding to adjacent vertices by the line segment between them yields a drawing $\mu(M)$ of M in the plane T_f.

The color and orientation of edges are nicely encoded in this drawing: Let v be a vertex with $\mu(v) = (v_1, v_2, v_3)$. The three lines $x_1 = v_1$, $x_2 = v_2$ and $x_3 = v_3$ partition the plane T_f into six wedges with apex $\mu(v)$. By Fact 5, the color and orientation of edges incident to v is determined by the wedge containing them; see Figure 3. In particular the bicolored edges are the edges supported by the lines defining the wedges.

THEOREM 1. *The drawing* $\mu(M)$ *is a convex drawing of* M *in* T_f. *Dropping the third coordinate yields a convex drawing of* M *on the* $(f-1) \times (f-1)$ *grid.*

PROOF (sketch).

- For every edge $\{u, w\}$ and vertex $v \neq u, w$ the point $\mu(v)$ is not contained in the segment $[\mu(u), \mu(v)]$ representing the edge.
- There are no crossing edges in $\mu(M)$, i.e., the embedding is planar. (This can be concluded from the observation that face-counting yields a weak barycentric embedding as defined by Schnyder [21].)

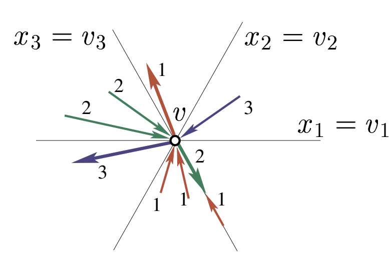

Fig. 3. Wedges and edges at a vertex v in the plane T_f.

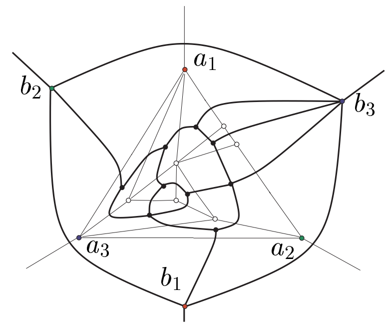

Fig. 4. The suspension dual of the example from Figure 2.

- The outgoing edges at a vertex (see Figure 3) guarantee that all interior angles of $\mu(M)$ are $\leq \pi$, i.e., the embedding is convex.
- Planarity and convexity are preserved by the projection of the drawing from T_f to the plane $x_3 = 0$. □

2.2. *The Lattice of Schnyder Woods*. In general the suspension M^σ of an internally 3-connected planar map will admit many Schnyder woods. Felsner [10] has shown that the set of all Schnyder woods of a given M^σ has the structure of a distributive lattice. As we will make use of some elements of this theory we recall some definitions and the main results.

Think of the three half-edges of M^σ as noncrossing infinite rays. These rays partition the outer face of M into three parts. The *suspension dual* M^{σ^*} of M^σ is the dual of this map. Thus M^{σ^*} has a triangle b_1, b_2, b_3 corresponding to the unbounded face of M. Half-edges reaching into the unbounded face of M^{σ^*} are attached to the three suspension vertices b_i. Figure 4 shows an example.

The *completion* $\widetilde{M^\sigma}$ of a plane suspension M^σ and its dual M^{σ^*} is obtained as follows: Superimpose M^σ and M^{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 $\{i, j, k\} = \{1, 2, 3\}$). At each crossing place a new vertex such that this new *edge vertex* subdivides the two crossing edges.

The completion $\widetilde{M^\sigma}$ is planar, every edge-vertex has degree four and there are six half-edges reaching into the unbounded face.

A 3-*orientation* of the completion $\widetilde{M^\sigma}$ of M^σ is an orientation of the edges of $\widetilde{M^\sigma}$ such that:

(O1) $\mathsf{outdeg}(v) = 3$ for all primal- and dual-vertices v.
(O2) $\mathsf{indeg}(v_e) = 3$ for all edge-vertices v_e (hence, $\mathsf{outdeg}(v_e) = 1$).
(O3) All half-edges are out-edges of their vertex.

THEOREM 2. *Let M^σ be a suspension of an internally* 3-*connected plane graph M. The following structures are in bijection*:

(1) *Schnyder woods of M^σ*.
(2) *Schnyder woods of the suspension dual M^{σ^*}*.
(3) 3-*Orientations of the completion $\widetilde{M^\sigma}$*.

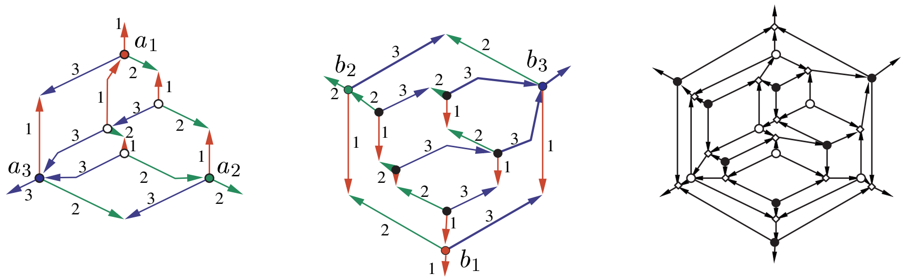

Fig. 5. The bijections for Theorem 2.

The bijections are illustrated in Figure 5. The proof of this theorem is given in [10]. The proof of Lemma 2 contains a piece of detail about the bijections.

The lattice structure of Schnyder woods is best understood by looking at 3-orientations: Let X be a 3-orientation and let C be a directed cycle of X. Reverting the orientation of all edges of C yields another 3-orientation X^C. If C is a simple directed cycle it has a connected interior and we can speak of the clockwise and the counterclockwise order of C. Define $X \succ X^C$ if C is a clockwise-directed cycle in X. The transitive closure $\succ^*$ of this relation is an order relation on the set of 3-orientations.

THEOREM 3. *The relation $\succ^*$ is the order relation of a distributive lattice on the set of 3-orientations of the completion $\widetilde{M^\sigma}$ of a suspension M^σ of an internally 3-connected planar map. The unique minimum 3-orientation contains no clockwise directed cycles.*

In view of Theorem 2 a suspension M^σ has unique minimum Schnyder wood S_{Min}. Figure 6 shows two sub-structures that are impossible in S_{Min}.

- A uni-directed edge incoming at v in color $i+1$ such that the counterclockwise next edge is bi-directed, outgoing at v in color $i-1$ and incoming in color i.
- A clockwise triangle of uni-directed edges, having colors $i, i+1, i+2$ in this clockwise order.

An algorithm to compute S_{Min} has been described and analyzed by Fusy et al. [11]. The result is the following:

THEOREM 4. *Let M^σ be a suspended 3-connected planar map. The minimal Schnyder wood S_{Min} of M^σ can be computed in linear time.*

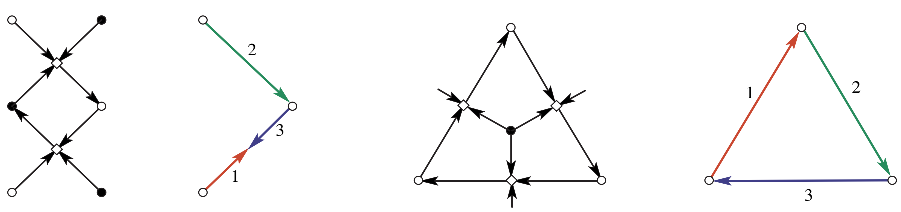

Fig. 6. Two types of clockwise cycles in 3-orientations and the corresponding sub-structures of Schnyder woods.

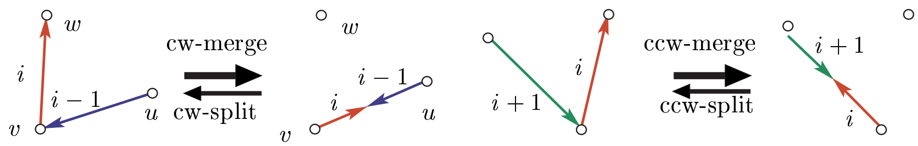

Fig. 7. Clockwise and counterclockwise merge and split.

2.3. *Merging and Splitting*. The operations *merge* and *split* introduced in this section operate on Schnyder woods and the underlying graph. Merge and split can be seen as inverse operations, corresponding to the deletion and insertion of an edge.

Given a Schnyder wood, a *knee at vertex* v is an ordered pair of uni-directed edges adjacent at an angle of v such that the first of the edges is incoming and the second outgoing at v. Knees come in two kinds: if the in-edge of the knee is the clockwise neighbor of the out-edge at v we speak of a *cw-knee*; otherwise, if the in-edge of the knee is the counterclockwise neighbor of the out-edge it is a *ccw-knee*.

Let (u, v), (v, w) be a knee at v. Suppose that the color of (v, w) is i; by the vertex condition the color of (u, v) is $i - 1$ if it is a cw-knee and $i + 1$ if it is a ccw-knee. The *merge of the knee* consists of the deletion of the out-edge (v, w) while making (u, v) a bi-directed edge outgoing at v in color i and incoming in the same color as before. Depending on the type of the knee we distinguish between clockwise and counterclockwise merge operations. Figure 7 illustrates the definition.

LEMMA 1. *Let S be a Schnyder wood*; *the coloring and orientation of edges after merging a knee is again a Schnyder wood*.

PROOF. The first three conditions (W1), (W2) and (W3) of Schnyder woods obviously remain true after the merge. Instead of arguing for (W4) we use the bijection with 3-orientations (see Theorem 2) from [10] (the idea of the bijection is shown in Figure 12). The merge of the knee corresponds to the deletion of an edge-vertex and the merge of two face-vertices as shown in Figure 8. The result of the merge is again a 3-orientation. □

A *split of a bi-directed edge* is the inverse operation of a merge. A split, however, is not determined by the choice of a bi-directed edge. The bi-directed edge can only split into one of its adjacent faces; this corresponds to the choice for the split of being

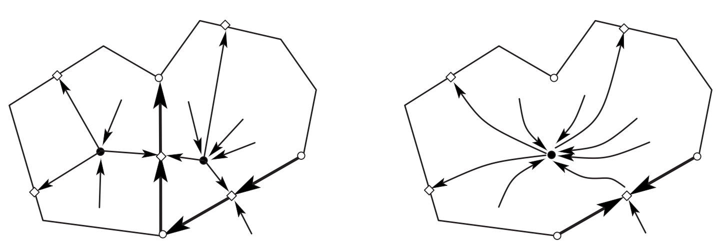

Fig. 8. A cw-merge in a 3-orientation.

the inverse of a cw-merge or a ccw-merge, i.e., a cw-split or a ccw-split. This choice determines the resulting color of the edge. If this choice is fixed, there can still be several choices for the second endpoint of the split edge.

At this point we stay with the remark that every bi-directed edge can be split; actually, it can be split by a cw-split or a ccw-split. To see that this should be true look at Figure 8 from right to left.

In the context of this paper we only need one very specific type of split. The *short cw-split* is the inverse of a cw-merge with the additional property that (u, w) is an edge, i.e., u, v, w form a triangle.

3. The Drawing Algorithm. Let M be a 3-connected planar map with n vertices and f faces. The steps of the drawing algorithm with input M are the following:

(A1) Choose three vertices from the outer face for the suspension M^σ.
(A2) Compute the minimum Schnyder wood S_{Min} for M^σ and let $S_0 = S_{\mathsf{Min}}$.
(A3) Compute a maximal cw-merge sequence $S_0 \to S_1 \to \cdots \to S_k$ of Schnyder woods, i.e., S_{i+1} is obtained from S_i by a cw-merge and S_k contains no cw-knee.
(A4) Use face-counting to draw S_k on the $(f - k - 1) \times (f - k - 1)$ grid.
(A5) Reinsert all edges that have been deleted by merge operations into the drawing from the previous step.

With Figure 9 we illustrate step (A3) of the algorithm.

Note from the example that the Schnyder woods of a merge sequence may correspond to suspended maps that are only internally 3-connected.

The gray triangle in the left part of Figure 9 contains a ccw-knee which disappears with a cw-merge (see the right part of the figure). An important fact for the analysis of our algorithm is that cw-merges never make a cw-knee disappear.

3.1. *The Drawing Is Convex*

THEOREM 5. *Reinserting all the edges that have been deleted by a sequence of cw-merge operations into the drawing of S_k obtained in step* (A4) *keeps the drawing planar and convex.*

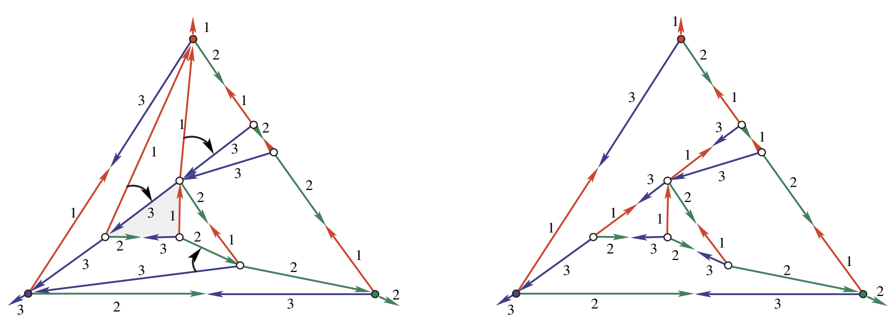

Fig. 9. On the left a Schnyder wood S_0; cw-knees are indicated by arcs. On the right the final Schnyder wood of a merge sequence.

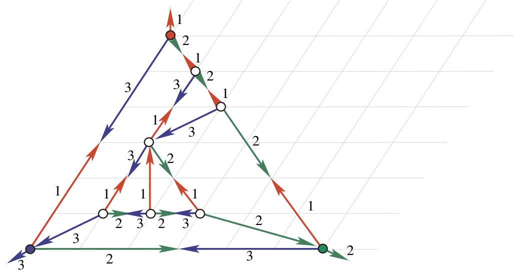

Fig. 10. The example graph with $n = f = 9$ drawn on the 6×6 grid.

The drawing steps of the algorithm ((A4) and (A5)) are illustrated in Figure 10. Essential for the proof of the theorem is the following lemma:

LEMMA 2. *Given a Schnyder wood of a suspended map M^σ, let F be an interior face. The orientation and coloring of edges around F obey the following rule* (*see Figure* 11):

- *In clockwise order the types of edges at the boundary of the face can be described as follows* (*in case of bi-directed edges the clockwise color is noted first*): *One edge from the set* {1-cw, 3-ccw, 1-3}, *any number* (*may be* 0) *of edges* 2-3, *one edge from the set* {2-cw, 1-ccw, 2-1}, *any number of edges* 3-1, *one edge from the set* {3-cw, 2-ccw, 3-2}, *any number of edges* 1-2.

PROOF. There is a bijection between Schnyder woods of M^σ and the dual M^{σ^*} (Theorem 2). This bijection can be constructed edge by edge; the rule is shown in Figure 12.

Given this rule the statement of the lemma is equivalent to the vertex condition (W3) at the vertex v_F dual to face F. □

In Figure 11 the faces are drawn with a surrounding triangle. It is one of the features of drawing via face counting, as described in Section 2.1, that all the vertices of a face sit on the boundary of such a triangle. Planarity of the drawing implies that there are no vertices in the interior of the face. Even more is true: there are no vertices in the (open) interior of the bounding triangle for the face.

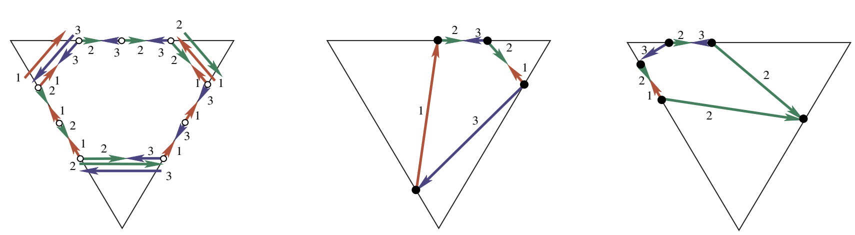

Fig. 11. The generic structure of a face as described by Lemma 2 and two concrete instances.

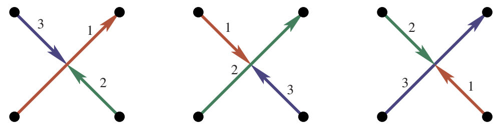

Fig. 12. Rule for orientation and coloring of dual edges.

Since we will make use of the shape of a face in the drawing we include a proof.

LEMMA 3. *Given a suspended map M^σ with a Schnyder wood, let $\mu(M)$ be the face-count drawing of M in the plane $x_1 + x_2 + x_3 = f - 1$ and let F be an interior face of M. Then the vertices of F are placed on the boundary of a triangle with sides $x_1 = c_1$, $x_2 = c_2$ and $x_3 = c_3$ as shown in Figure* 11.

PROOF. Lemma 2 gives information about colorings and orientation of edges around F. Fact 5 and Figure 3 state how edges of a given color and orientation are embedded with respect to their incident vertices. In combination this implies the statement of the lemma. □

PROOF OF THEOREM 5. Consider a merge operation performed during step (A3) of the algorithm. The bi-directed edge e resulting from the merge is an edge of S_k. For concreteness let us assume that $e = (u, v)$ was originally colored 3 and the merge was a cw-merge at v. It follows that the edge $e' = (v, w)$ that was merged into e was colored 1 (this is the situation shown in the left part of Figure 7).

In the drawing of S_k consider the face F which is to the left of (v, u) and let ∇ be the bounding triangle for F (Lemma 3). Given the colors of the bicolored edge e, it follows from Lemma 3 that v and u belong to the boundary line of ∇ with equation $x_2 = c_2$; moreover, $u_1 > v_1$. Vertex w also belongs to F. The edge e' removed in the merge was entering w between the 2-outgoing and the 3-outgoing edge in clockwise order. Therefore, F is contained in the region $R_1(w)$.

Suppose that the area of F is in the region $R_1(w)$ with respect to the final Schnyder wood S_k of the merge sequence. In this case vertex w is placed on the boundary line of ∇ with equation $x_1 = c_1$. The positions of v and w in the triangle ∇ imply that the straight edge (v, w) can be added to the drawing.

The alternative is that F does not remain in the region $R_1(w)$. This can only happen if there is a later merge using a knee at w that contains F in its angle. Since this is a cw-merge it must merge the 2-outgoing edge at w into a 1-incoming. The area of F is in the region $R_3(w)$ after this merge. Further merges cannot change this situation. In this case vertex w is placed on the boundary line of ∇ with equation $x_3 = c_3$. From the positions of v and w in the triangle ∇ we can again conclude that the straight edge (v, w) can be added to the drawing.

A given face F of the drawing of S_k may be the host for more than one reinsertion of a merge-edge. We can argue that all these reinsertions can be done without conflict as follows: The boundary of face F is a cycle C in the planar map M. Edges that have disappeared while transforming the Schnyder wood S_0 corresponding to M to S_k are chords of C. In M all these chords are drawn without crossings in the interior of C. In the drawing of S_k the cycle C is the boundary of the convex face F. Hence, the chords can be reinserted without crossings.

It remains to verify the convexity of all faces after reinsertion of edges. The reinsertion partitions a convex face F into pieces using a set of non-crossing straight edges, each connecting two points on the boundary of F. From the definition of a convex set it follows that all the pieces, i.e., faces of M, are convex as well. ☐

3.2. *The Number of Merges*. Essential for the grid-size required for the drawing produced by the algorithm is the length k of the merge sequence computed in step (A3). The main result in this section is a lower bound for k in terms of easily recognizable substructures of the initial Schnyder wood S computed in step (A2) of the algorithm.

As a warm-up we consider the case where M is a triangulation and S is an arbitrary Schnyder wood of M. Consider the $(2n-4)-4$ triangles of S which are bounded by three uni-directed edges (only external edges are bi-directed). These triangles can be partitioned into two classes: Class one are those with at least two clockwise-oriented edges on the boundary and class two are those with at least two counterclockwise edges on the boundary. Suppose that the number C_1 of triangles of class one is the larger one or that C_1 and C_2 are equal, i.e., $C_1 \geq n-4 \geq C_2$. In a triangle T of class one there is a knee of two consecutive clockwise edges of T, this knee is a candidate for a clockwise merge. Since every edge is clockwise only for one of its neighboring triangles these C_1 merges can be performed independently. It follows that starting from S there is a merge sequence of length $k \geq C_1 \geq n-4$. This estimate yields drawing of triangulations on grids of size at most $(f-(n-4)-1) \times (f-(n-4)-1) = (n-1) \times (n-1)$.

Let again M be a triangulation with a Schnyder wood S. We aim for a more precise estimate for the number of merges that can be applied to S. Consider a cw-knee at v, let i be the color of the incoming edge (u, v) of the knee. The edge (u, v) is a witness that v is an inner vertex of the tree T_i of i-colored edges in S. Conversely, if v is an inner vertex of T_i, then the outgoing edge in color $i+1$ together with its adjacent incoming edge in color i form a cw-knee. For fixed S this proves a bijection between inner vertices of T_i and cw-knees with an incoming edge of color i. The number of cw-knees thus is $\sum_i \mathsf{inner}(T_i)$. If a uni-directed edge (v_1, v_2) participates at two different cw-knees, then both v_1 and v_2 are vertices of a cw-knee. It follows that the triangle to the right of (v_1, v_2) is a clockwise triangle of S. A clockwise triangle contributes three cw-knees that are pairwise incompatible. If $\Delta_S^{\circlearrowright}$ is the number of clockwise triangles of S, then the number of cw-merges that can be performed with initial Schnyder wood S is at least $\sum_i \mathsf{inner}(T_i) - 2\Delta_S^{\circlearrowright}$. This leads to Proposition 1 which makes use of the following counts for a Schnyder wood S of a plane triangulation:

- $\Delta_S^{\circlearrowright}$ is the number of clockwise triangles of S.
- $\Delta_S^{\circlearrowleft}$ be the number of counterclockwise triangles of S.

PROPOSITION 1. *Let S be a Schnyder wood with $\Delta_S^{\circlearrowright}$ clockwise and $\Delta_S^{\circlearrowleft}$ counterclockwise triangles. The number of cw-merges applicable in a merge sequence starting with S is at least $n-4-\Delta_S^{\circlearrowright}+\Delta_S^{\circlearrowleft}$.*

PROOF. We have already proven the lower bound $\sum_i \mathsf{inner}(T_i) - 2\Delta_S^{\oplus}$ for the number of cw-merges applicable in a merge sequence starting with S. Thus it is enough to prove

$$\sum_i \mathsf{inner}(T_i) = n - 4 + \Delta_S^{\oplus} + \Delta_S^{\ominus}.$$

This formula is due to Bonichon et al. [3]. The following simple double counting proof was found by Lin et al. [14].

Each tree T_i spans all n vertices of the graph. However, it is easier if we disregard the three special vertices, such that $\mathsf{inner}(T_i) + \mathsf{leaves}(T_i) = n - 3$ for each T_i.

If v is a leaf in T_i, then outgoing edges at v in colors $i-1$ and $i+1$ are adjacent and the triangle containing both of them is not cyclic, neither clockwise nor counterclockwise. Conversely, a triangle that is not cyclic has a unique source vertex v with two outgoing edges. If these edges have colors $i-1$ and $i+1$, then v is a leaf in T_i. This proves a bijection between leaves of all colors and non-cyclic triangles in S. Hence, $\sum_i \mathsf{leaves}(T_i) = 2n - 5 - \Delta_S^{\oplus} - \Delta_S^{\ominus}$. Combining the formulas we have $\sum_i \mathsf{inner}(T_i) = 3(n-3) - \sum_i \mathsf{leaves}(T_i) = n - 4 + \Delta_S^{\oplus} + \Delta_S^{\ominus}$. □

Combining Theorem 5 and Proposition 1 we find that a triangulation with Schnyder wood S can be drawn on a grid of size $(n - 1 + \Delta_S^{\oplus} - \Delta_S^{\ominus}) \times (n - 1 + \Delta_S^{\oplus} - \Delta_S^{\ominus})$. An interesting special case of this bound (Corollary 1) was first obtained by Zhang and He [26] with a different method. Let again S_{Min} be the minimum Schnyder wood in the lattice and recall that $\Delta_{S_{\mathsf{Min}}}^{\oplus} = 0$.

COROLLARY 1. *A planar triangulation with n vertices has a straight-line drawing on a grid of size* $(n - 1 - \Delta_{S_{\mathsf{Min}}}^{\ominus}) \times (n - 1 - \Delta_{S_{\mathsf{Min}}}^{\ominus})$.

To estimate the number of merges that can be applied to a Schnyder wood S of a non-triangulated map we introduce two parameters:

- $\Delta_S^{\ominus}$ is the number of faces, with a counterclockwise edge in each of the three colors. (These edges are not required to be uni-directed.)
- $\Delta_S^{\hat\oplus}$ counts the number of clockwise triangles of uni-directed edges plus patterns of the following type: a uni-directed edge incoming at v in color $i+1$ such that the counterclockwise next edge around v is bi-directed, outgoing at v in color $i-1$ and incoming in color i; see Figure 13.

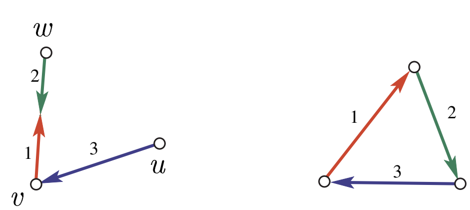

Fig. 13. The two patterns counted by $\Delta^{\hat\oplus}$.

THEOREM 6. *Let S be a Schnyder wood of a* 3*-connected planar map*. *The number of cw-merges that can be applied to S is at least $f - n + \Delta^{\ominus} - \Delta^{\oplus}$*.

PROOF. The proof for the special case where the outer face is a triangle is somewhat simpler. We first deal with this situation.

The bound is obtained from the bound of Proposition 1 as follows: Starting from S we construct a triangulation and a corresponding Schnyder wood S' such that S can be obtained from S' by a sequence of k cw-merges. Proposition 1 gives a bound k' for the number of cw-merges applicable to S'. The difference $k' - k$ is a bound for the number of cw-merges applicable to S.

Consider an internal face F in S. At the boundary of F there are three special edges, these are the dual edges of the outgoing edges of the dual vertex v_F. The special edges may be separated by bi-directed paths (see Lemma 2). The edges of the bi-directed boundary path of F in colors i and $i+1$ are split such that the edge in color $i+1$ points to the clockwise last vertex of the bi-directed path of F in colors $i-1$ and i. This can be achieved by a sequence of cw-splits. See Figure 14.

Applying the construction to all bounded faces of S yields an *inner* triangulation and a corresponding Schnyder wood S'. The following observations are crucial:

- The number of cw-merges applicable to S' is $k' \geq n - 4 + \Delta^{\ominus}_{S'} - \Delta^{\oplus}_{S'}$ (Proposition 1).
- The original Schnyder wood S can be obtained from S' via a sequence of k cw-merges.
- Each merge reduces the number of faces by one; hence, $k = (2n - 4) - f$.
- $\Delta^{\ominus}_{S'} = \Delta^{\ominus}_{S}$ and $\Delta^{\oplus}_{S'} = \Delta^{\oplus}_{S}$. This invariant is true for a single cw-split of the type described above; the result follows by induction (Figure 14 shows some examples).

Therefore, the number of cw-merges applicable to S is $k' - k \geq f - n + \Delta^{\ominus}_{S} - \Delta^{\oplus}_{S}$.

Now suppose that the outer face contains more than three vertices. Starting from S we produce a Schnyder wood S^* with an outer triangle; this is done with some external splits, as shown in Figure 15. Let the number of external splits required to get from S to S^* be $t \in \{0, 1, 2, 3\}$.

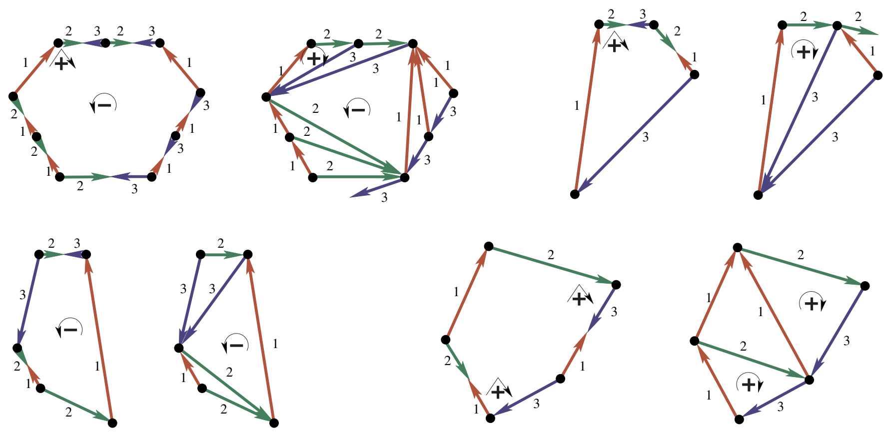

Fig. 14. Four examples for the triangulation of a face.

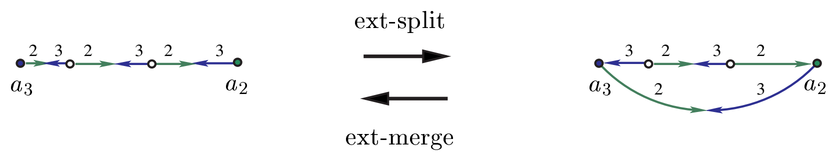

Fig. 15. External merge and split.

We know that S^* admits $k^* \geq f^* - n + \Delta^{\circlearrowleft}_{S^*} - \Delta^{\triangle}_{S^*}$ cw-merges. The goal is to derive the inequality $k \geq f - n + \Delta^{\circlearrowleft}_{S} - \Delta^{\triangle}_{S}$ by comparing the corresponding parts of the two formulae:

- The face numbers f and f^* of S and S^* are related by $f^* = f + t$.
- By comparing knees it is obvious that $k^* \geq k$. Let $t_1 = k^* - k$ and note that $t_1 \leq t$ since every knee of S^* that does not correspond to a knee of S is of the form shown in Figure 16(a). The orientation of edges at the suspension vertices a_i makes knees at these vertices impossible.
- It is easy to verify $\Delta^{\triangle}_{S^*} = \Delta^{\triangle}_{S}$.
- It remains to compare $\Delta^{\circlearrowleft}_{S^*}$ and $\Delta^{\circlearrowleft}_{S}$. It is obvious that $\Delta^{\circlearrowleft}_{S} \geq \Delta^{\circlearrowleft}_{S^*}$. A face F contributing to $\Delta^{\circlearrowleft}_{S}$ but not to $\Delta^{\circlearrowleft}_{S^*}$ is of the form shown in Figure 16(b). The crucial fact is that such a face cannot contribute to the difference in the count of cw-knees. Hence, $\Delta^{\circlearrowleft}_{S} - \Delta^{\circlearrowleft}_{S^*} \leq t - t_1$.

Together this shows that S admits at least $f - n + \Delta^{\circlearrowleft}_{S} - \Delta^{\triangle}_{S}$ cw-merges. □

Given an arbitrary Schnyder wood the contribution of $\Delta^{\circlearrowleft} - \Delta^{\triangle}$ in the above formula may well be negative. However, the choice of $S = S_{\mathsf{Min}}$ guarantees that $\Delta^{\triangle} = 0$. (Figure 6 shows that the two configurations counted by $\Delta^{\triangle}$ are impossible, since they imply a clockwise cycle in the completion S_{Min}.) The findings of this section can be summarized as follows:

THEOREM 7. *A* 3-*connected planar map* M *with* n *vertices has a convex drawing on a grid of size* $(n - 1 - \Delta^{\circlearrowleft}_{S_{\mathsf{Min}}}) \times (n - 1 - \Delta^{\circlearrowleft}_{S_{\mathsf{Min}}})$, *where* $\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}} \geq 0$ *is the number of faces with a counterclockwise edge in each color in* S_{Min}. *Such a drawing can be computed in linear time*.

PROOF. From the previous lemmas, it only remains to prove that all the cw-merges can be done in linear time. The algorithm that makes all these cw-merges is quite simple: for each inner vertex v in the tree T_i, such that its parent edge in T_{i+1} is uni-directed and the

Fig. 16. (a) A cw-knee of S^* which is not a cw-knee of S, exemplified at vertex a_3. (b) A face F contributing to $\Delta^{\circlearrowleft}_{S}$ but not to $\Delta^{\circlearrowleft}_{S^*}$.

edge toward its rightmost child in T_i is also uni-directed, cw-merge the outgoing edge colored $i+1$ with the edge toward the rightmost child of v. Since $\Delta^{\oplus} = 0$ the edge toward the rightmost child is not mergeable. So all the cw-merges are independent in S and can be performed in one run. □

In order to show that the parameter $\Delta^{\ominus}$ can be up to $n/2-2$, we consider M^σ the suspension of a 3-connected n-vertex cubic planar map. Let S be the minimal Schnyder wood of M^σ. The dual of S is a Schnyder wood of a triangulation. Every edge of the dual is uni-directed except the three external ones. Hence, except the four faces that are adjacent to an external edge, every face has at least three bi-directed edges respectively colored $1-2$, $2-3$ and $3-1$. These three edges are the ones corresponding to the three uni-directed outgoing edges in the dual. Each of these $n/2-2$ faces contributes to the parameter $\Delta^{\ominus}$ and to the parameter $\Delta^{\oplus}$. Moreover we can observe that for any Schnyder wood S' obtained from S applying cw-splits, we also have $\Delta^{\ominus}_{S'} = n/2-2$.

4. Improvements and Limitations. Our ambition was to design an algorithm for convex drawings of 3-connected planar graphs which at least matches all known algorithms for this task. Theorem 7 shows that we are very close. Still, there is Schnyder's $(n-2)\times(n-2)$ bound for triangulations which is not completely matched by $(n-1-\Delta^{\ominus}_{S_{\mathrm{Min}}})\times(n-1-\Delta^{\ominus}_{S_{\mathrm{Min}}})$ since there are triangulations with $\Delta^{\ominus}_{S_{\mathrm{Min}}} = 0$. An example of such a triangulation is shown in Figure 17.

It is indeed the case that with an algorithm that computes the positions of vertices just by face-counting the graph of Figure 17(a) requires a grid of size $(n-1)\times(n-1)$. This can be verified as follows:

- The suspension vertices a_1 and a_3 both embed on the y-axis.
- Every internal vertex of the triangulation is connected with both a_1 and a_3. Since there are no crossing edges, no two internal vertices can have the same x-coordinate.
- The two available grid points on the grid-line $x = n-2$ are used by the edges of the outer triangle. Therefore, this line contains no vertex.

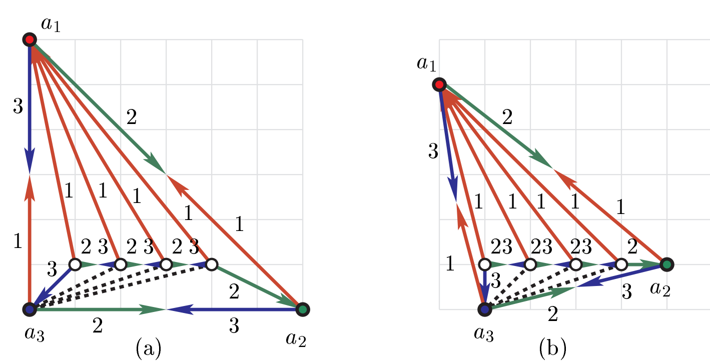

Fig. 17. (a) A stacked triangulation on the $(n-1)\times(n-1)$ grid. (b) The same graph drawn with the improved method.

Together this shows that n vertical grid-lines are necessary. In the next subsection we propose a modified drawing strategy that circumvents this obstruction. The effect of the modification is that the outer triangle is tilted; see Figure 17(b).

4.1. *From $n-1$ to $n-2$.* In the standard algorithm, the face-count of the vertices a_1 is $(f-k-1, 0, 0)$, a_2 is $(0, f-k-1, 0)$ and a_3 is $(0, 0, f-k-1)$. In order to reduce the grid size, we change the coordinates of these vertices as follows: $a_1 \rightarrow (f-k-2, 0, 1)$, $a_2 \rightarrow (1, f-k-2, 0)$ and $a_3 \rightarrow (0, 1, f-k-2)$. The effect on the drawing is that a_1 is moving down by one unit, a_2 is moving one unit to the left and one unit up, while a_3 is moving one unit to the right. Figure 17(b) shows an example.

Below we show that the resulting drawing is convex and planar. This yields:

THEOREM 8. *A 3-connected planar map M with n vertices has a convex drawing on a grid of size $(n-2-\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}}) \times (n-2-\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}})$, where $\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}} \geq 0$ is the number of faces with a counterclockwise edge in each color in S_{Min}.*

Figure 23 at the end of the article presents an example of produced drawings.

In the proof of the theorem, we use the following lemma. It may be interesting to observe that again the validity of this lemma depends on our choice of only using cw-merges to produce S_k.

LEMMA 4. *In the face-count of S_k there is no vertex with the coordinates $(f-k-2, 0, 1)$, $(1, f-k-2, 0)$ or $(0, 1, f-k-2)$.*

PROOF. Assume that there is a vertex v with the face-count $(1, f-k-2, 0)$. Since $|R_1(v)| = 1$ there is a unique face F in $R_1(v)$. Since $|R_3(v)| = 0$ vertex v is on the bi-directed path between a_1 and a_2. It follows that the degree of a_2 in S_k is 2. The path $P_3(v)$ from v to a_3 uses a part of the boundary of F to get from v to the bi-directed path between a_2 and a_3. A cw-split applied to an edge of F would have to split one of the edges of $P_3(v)$. This is impossible since these edges have color 3 in ccw direction. We conclude that $deg(a_2) = 2$ in the original Schnyder wood S_0. This is in contradiction to the assumption that G is 3-connected. The two other cases are symmetrical. □

PROOF OF THEOREM 8. The previous lemma shows that the new coordinates of the suspension vertices do not coincide with other vertices.

The next step is to show that after the change in the coordinates of the suspension vertices the drawing of S_k remains convex. Consider a face F containing a_2. With the new coordinates, a_2 is interior to the triangle associated with F by Lemma 3; see Figure 18. From Section 3.1 we know that the (open) interior of the bounding triangle of F contains no vertices. This shows that the modified drawing is free of crossings. The convexity of the new drawing of face F can also be read off from Figure 18.

The reinsertion of edges can be performed without introducing crossings or non-convex faces. This follows from the proof of Theorem 5 which also works for faces of the types shown in Figure 18. □

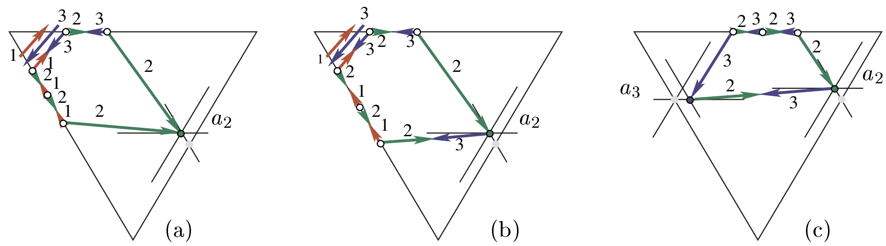

Fig. 18. (a) Generic structure of a face containing one external vertex and no external edges. (b) Generic structure of a face containing one external vertex and one external edge. (c) Generic structure of a face containing two external vertices.

Note that if ccw-merges are allowed to produce S_k, then the result becomes false. In Figure 18(c) it could be the case that a_2 is on the top edge of the triangle, i.e., the edge colored 2 pointing to a_2 is horizontal. Applying a ccw-split to one of the horizontal double edges 2-3 would then produce overlapping edges.

We conclude this subsection with an application of Theorem 8 to internally 3-connected planar maps.

A planar map M is internally 3-connected iff adding a new vertex v^+ connected to all vertices of the outer face yields a 3-connected map M^+. Thomassen [23] proves the following characterization: *M has a (strictly) convex drawing if and only if M is internally 3-connected.*

A drawing is called *internally convex* if all bounded faces are convex; the outer face, however, may be ragged. Chrobak and Kant [4] adapt their algorithm so that internally 3-connected graphs are drawn internally convex on the $(n-1) \times (n-2)$ grid. With our approach we can reduce the grid size.

Let M be internally 3-connected with n vertices. First, extend M to M^+ by adding a new vertex v^+ connected it to all vertices of the outer face of M. Let $S = S_{\mathsf{Min}}$ be the minimal Schnyder wood of the suspension of M^+ with $v^+ = a_1$. Since M^+ has $n+1$ vertices, Theorem 8 guarantees a drawing of M^+ on the $(n - 1 - \Delta_S^{\circlearrowleft}) \times (n - 1 - \Delta_S^{\circlearrowleft})$ grid. Since the outer face of M^+ is a triangle, a_1 is the only vertex on the highest horizontal and on the leftmost vertical grid-line. Therefore we only have to remove a_1 to prove:

COROLLARY 2. *An internally 3-connected map M with n vertices can be drawn internally convex on the $(n - 2 - \Delta_S^{\circlearrowleft}) \times (n - 2 - \Delta_S^{\circlearrowleft})$ grid, where S is a minimal Schnyder wood of M^+.*

4.2. *Ignoring Faces.* In this subsection we propose another technique that can be used to reduce the size of a drawing. The idea is to ignore some of the faces for the count of faces in regions. More formally, define weights $w: \mathcal{F}_b \to \{0, 1\}$ for all bounded faces $F \in \mathcal{F}_b$. This extends linearly to weights for the regions of vertices. Map a vertex v to the point $(w(R_2(v)), w(R_1(v)))$ in the plane and add the edges as line segments. If we are lucky, then this results in a convex drawing on a grid with side-length less than $f - 1$. Figure 19 shows an example.

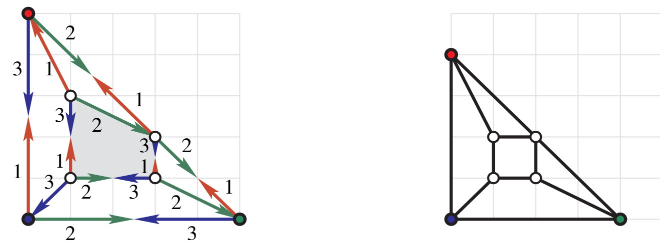

Fig. 19. Ignoring the gray face in the count yields the more compact drawing on the right.

Analyzing the proof of the drawing algorithm it can be concluded that the following property is just what is needed to guarantee a convex drawing:

- If u and v are vertices and $R_i(u) \subsetneq R_i(v)$, then $w(R_i(u)) < w(R_i(v))$.

This condition can be used to identify a face whose weight can be set to zero without spoiling the convexity of the drawing. The drawback with the condition is that we know of no really efficient test. Below we give a proposition with a sufficient criterion that is much easier to verify.

A face F is called *pointed in color* i if there is a vertex $x \in F$ such that the two boundary edges of F that are incident to x are both directed toward x in color i. Figure 20 illustrates the concept.

PROPOSITION 2. *If the weight of non-pointed faces is set to zero, then the drawing is still convex.*

PROOF. We verify the above condition for the weight $w\colon \mathcal{F}_b \to \{0, 1\}$ with $w(F) = 0$ for all non-pointed faces F. Let u and v be a pair of vertices with $R_i(u) \subsetneq R_i(v)$. Let $P_{i+1}(u)$ (resp., $P_{i+1}(v)$), the path of color $i+1$ from u (resp., v) to a_{i+1}. There are two cases to consider:

- $P_{i+1}(v)$ merge at a face pointed in color $i+1$ that is contained in $R_i(v)$ but not in $R_i(u)$. Since the weight of this face is non-zero we obtain $w(R_i(u)) < w(R_i(v))$ as wanted.
- $P_{i+1}(u)$ is a subpath of $P_{i+1}(v)$. In this case, however, paths $P_{i-1}(u)$ and $P_{i-1}(v)$ merge at a face pointed in color $i-1$ so that the inequality $w(R_i(u)) < w(R_i(v))$ also holds for this case. □

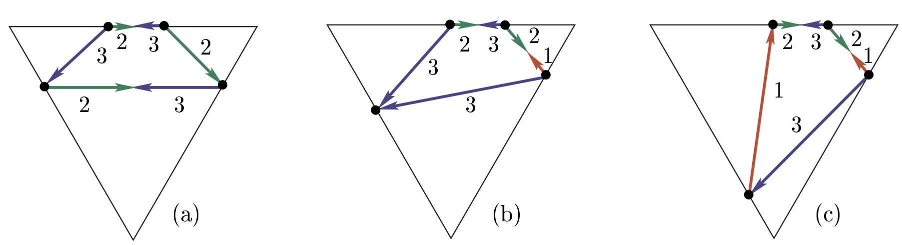

Fig. 20. (a) A face pointed in colors 2 and 3. (b) A face pointed in color 3. (c) A non-pointed face.

In their recent paper about drawings of triangulations Zhang and He [26] use a similar idea of ignoring faces when counting faces in regions.

In conjunction with our drawing algorithm the technique of "ignoring faces" can only be applied with care: A drawing of S_k obtained with a count that ignores non-pointed faces is convex (Proposition 2) but the faces need not obey the generic structure (Lemma 2); therefore, the reinsertion of edges may cause an overlap of edges.

4.3. *Experimental Results*. Theorem 8 gives a bound for the size of a grid that accommodates a given 3-connected planar graph. This bound depends on the parameter $\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}}$ which can be equal to zero in "bad" cases. Bonichon et al. [2] have analyzed the asymptotic average value of $\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}}$ over triangulations with n vertices. They prove that $E(\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}}) = n/8 + o(n)$. Hence, the number of merges that can be applied when starting with the S_{Min} of a random triangulation is of order $9n/8$ and the expected size of the grid required for the drawing is $7n/8 \times 7n/8$.

For general 3-connected planar graphs there is no theory about the expected size of $\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}}$. Still it is possible to make some experiments. For this purpose, we have first tested our drawing algorithm on random planar maps generated uniformly over m-edge 3-connected planar maps. The generator used is due to Schaeffer [19]. We generated some 6000 maps with m edges where m goes from 200 up to 20000. For each map we have computed the parameter $\Delta^{\circlearrowleft}$, the number of faces and the grid size of the drawing of the graph obtained with the algorithm of Theorem 8 (see Figure 21). We can first observe that the expected value of $\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}}$ is $m/8$. Using Euler's formula, we also see that the expected number of vertices is equal to the expected number of faces of a random m-edge 3-connected planar map: $m/2$. Consequently, the expected size of the grid required for the drawing is $3n/4 \times 3n/4$.

Our second experimental result is obtained on 3-connected cubic planar maps. For the class of maps, the analysis of $\Delta^{\circlearrowleft}_{S_{\mathsf{Min}}}$ is quite trivial: All the edges except three (one connected to each suspension vertex) are bicolored, so no merge is possible. In order to reduce the grid size, we have experimented with ignoring non-pointed faces

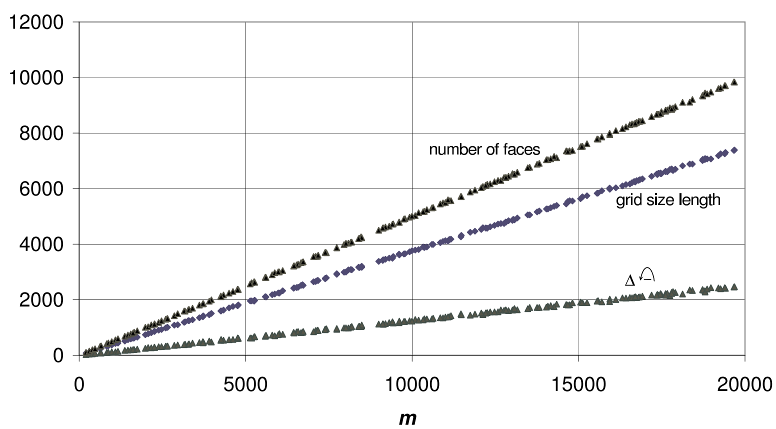

Fig. 21. Experimental results for 3-connected planar graphs of different sizes.

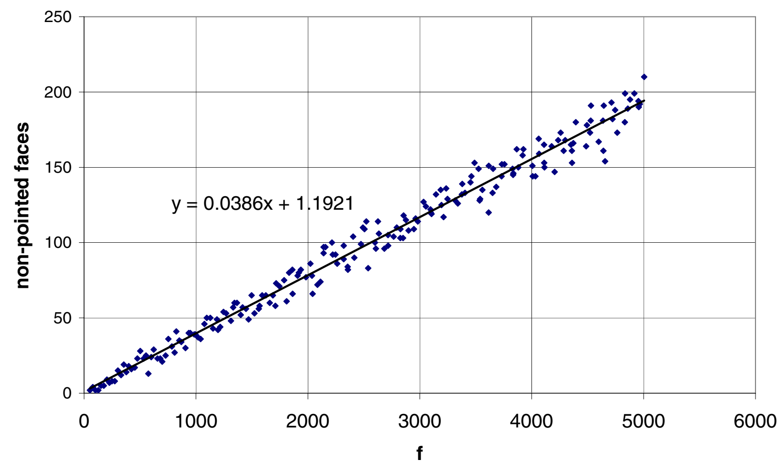

Fig. 22. Number of non-pointed faces of 200 uniform random cubic planar graphs of different sizes.

as described in Section 4.2. Random cubic graphs are easily obtained as duals of random triangulations. Once again, we used the uniform random triangulation generator due to Schaeffer [19]. Figure 22 shows that approximately 3.8% of the faces are non-pointed. Hence, for these graphs the expected size of the grid required for the drawing is $(0.481n) \times (0.481n)$.

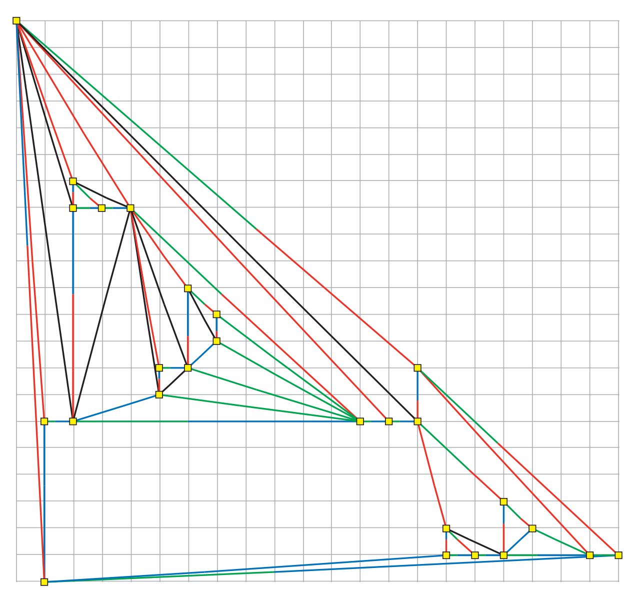

Fig. 23. An example of drawing obtained by our algorithm (Theorem 8). The present graph has 26 vertices, 57 edges, 33 faces. It is drawn on a 21×21 grid.

References

[1] N. Bonichon, S. Felsner, and M. Mosbah, Convex drawings of 3-connected planar graphs (extended abstract), in *Proc*. *GD* '04, vol. 3383 of Lecture Notes Comput. Sci., Springer-Verlag, Berlin, 2004, pp. 60–70.

[2] N. Bonichon, C. Gavoille, N. Hanusse, D. Poulalhon, and G. Schaeffer, Planar graphs, via well-orderly maps and trees, in *Proc*. *WG* '04, vol. 3353 of Lecture Notes Comput. Sci., Springer-Verlag, Berlin, 2004, pp. 270–284.

[3] N. Bonichon, B. Le Saëc, and M. Mosbah, Wagners theorem on realizers, in *Proc*. *ICALP* '02, vol. 2380 of Lecture Notes Comput. Sci., Springer-Verlag, Berlin, 2002, pp. 1043–1053.

[4] M. Chrobak and G. Kant, Convex grid drawings of 3-connected planar graphs, *Internat*. *J*. *Comput*. *Geom*. *Appl*., **7** (1997), 211–223.

[5] H. de Fraysseix, J. Pach, and R. Pollack, Small sets supporting Fary embeddings of planar graphs, in *Proc*. 20*th Annu*. *ACM Sympos*. *Theory Comput*., 1988, pp. 426–433.

[6] H. de Fraysseix, J. Pach, and R. Pollack, How to draw a planar graph on a grid, *Combinatorica*, **10** (1990), 41–51.

[7] G. Di Battista, R. Tamassia, and L. Vismara, Output-sensitive reporting of disjoint paths, *Algorithmica*, **23** (1999), 302–340.

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

[9] S. Felsner, *Geometric Graphs and Arrangements*, Vieweg Verlag, Braunschweïg, 2004.

[10] S. Felsner, Lattice structures from planar graphs, *Electron*. *J*. *Combin*., **11** R15 (2004), 1–24.

[11] E. Fusy, D. Poulalhon, and G.Schaeffer, Coding, counting and sampling 3-connected planar graphs, in *Proc*. 16*th ACM–SIAM Sympos*. *Discrete Algorithms*, 2005, pp. 690–699.

[12] X. He, Grid embeddings of 4-connected plane graphs, *Discrete Comput*. *Geom*., **17** (1997), 339–358.

[13] G. Kant, Drawing planar graphs using the canonical ordering, *Algorithmica*, **16** (1996), 4–32.

[14] C. Lin, H. Lu, and I. Sun, Improved compact visibility representation of planar graph via Schnyder's realizer, in *Proc*. *STACS* '03, vol. 2607 of Lecture Notes Comput. Sci., Springer-Verlag, Berlin, 2003, pp. 14–25.

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

[16] K. Miura, S. Nakano, and T. Nishizeki, Grid drawings of 4-connected plane graphs, *Discrete Comput*. *Geom*., **26** (2001), 73–87.

[17] P. Rosenstiehl and R. E. Tarjan, Rectilinear planar layouts and bipolar orientations of planar graphs, *Discrete Comput*. *Geom*., **1** (1986), 343–353.

[18] G. Rote, Strictly convex drawings of planar graphs, in *Proc*. 16*th ACM–SIAM Sympos*. *Discrete Algorithms*, 2005, pp. 728–734.

[19] G. Schaeffer, Random sampling of large planar maps and convex polyhedra, in *Proc*. 31*st Annu*. *ACM Sympos*. *Theory Comput*. (*Atlanta*, *GA*, 1999), 1999, pp. 760–769.

[20] W. Schnyder, Planar graphs and poset dimension, *Order*, **5** (1989), 323–343.

[21] W. Schnyder, Embedding planar graphs on the grid, in *Proc*. 1*st ACM–SIAM Sympos*. *Discrete Algorithms*, 1990, pp. 138–148.

[22] W. Schnyder and W. T. Trotter, Convex embeddings of 3-connected plane graphs, *Abstracts of the Amer*. *Math*. *Soc*., **13** (1992), 502.

[23] C. Thomassen, Planarity and duality of finite and infinite planar graphs, *J*. *Combin*. *Theory Ser*. *B*, **29** (1980), 244–271.

[24] W. T. Tutte, Convex representations of graphs, *Proc*. *London Math*. *Soc*., **10** (1960), 304–320.

[25] W. T. Tutte, How to draw a graph, *Proc*. *London Math*. *Soc*., **13** (1963), 743–768.

[26] H. Zhang and X. He, Compact visibility representation and straight-line grid embedding of plane graphs, *Proc*. *WADS* '03, vol. 2748 of Lecture Notes Comput. Sci., Springer-Verlag, Berlin, 2003, pp. 493–504.