Grafy a počty - NDMI078 April 2009

Lecture 7

M. Loebl *J.-S. Sereni*

This lecture deals with some aspects of plane triangulations: we introduce 3-orientations and Schnyder woods, and we give an algorithm to build such structures. We mainly rely on the original paper by Schnyder [8] and the Diploma Thesis of Brehm [4].

1 3-orientations

Let G be a plane triangulation. An *inner vertex* of G is a vertex that is not incident with the outer face. An *outer vertex* is a vertex incident with the outer face. An *inner edge* is an edge that is not incident with the outer face (thus, an inner edge is incident with at least one inner vertex). An *outer edge* is an edge incident with the outer face.

Definition 1. Let G be a plane triangulation. A 3-*orientation* of G is an orientation of the inner edges of G in which every inner vertex has out-degree 3.

Let G be a plane triangulation with a 3-orientation. A straightforward consequence of the definition of a 3-orientation is that the outer vertices have out-degree 0. To see this, first note that a plane triangulation on n vertices has exactly $3n-6$ edges by Euler's formula. Among those, there are 3 outer edges and hence $3n-9$ inner edges. Now, the sum of the out-degrees of the inner vertices equals the number of inner edges, i.e. $3n-9$. Since there are $n-3$ inner vertices and each of them has out-degree 3, this yields the stated fact.

This simple observation extends to every triangle of G, not only the one bouding the outer face. Moreover, the converse is also true, i.e. if C is a cycle of G such that no arc inside C leaves a vertex of C, then C is a triangle (exercise).

2 Schnyder woods

A look at Figure 1 can help to read the following definition.

Definition 2. Given a plane triangulation $T=(V,E)$, a *Schnyder labelling* of T is a 3-orientation $\overrightarrow{T}$ of T along with a 3-colouring of the inner arcs of $\overrightarrow{T}$ such that

- every vertex has exactly one outgoing arc of each colour in $\overrightarrow{T}$;
- the colours of the outgoing arcs around a vertex always appear in the same clockwise order; and

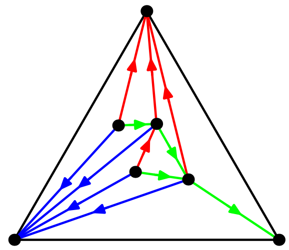

Figure 1: A Schnyder labelling of a plane triangulation.

- the incoming arcs of a given colour all appear between the outgoing arcs of the other two colours.

It turns out that *every plane triangulation has a Schnyder labelling* [8]. We prove this fact in the next section. To this end, we provide a linear-time algorithm to obtain a Schnyder labelling of any given plane triangulation.

Before doing so, let us note the important corollary that *the inner edges of any plane triangulation can be partitionned into three disjoint trees that span all the inner vertices, each rooted at a different outer vertex.* Indeed, a Schnyder labelling defines for every inner vertex a unique directed path to each of the three outer vertices (just follow the out-going red arcs, the out-going blue arcs and the out-going green arcs to obtain the three paths).

Such decompositions have a name: a 3-*tree-decomposition* of a plane triangulation is a partition of its inner edges into three trees, each being rooted at a different outer vertex and containing all inner vertices.

Thus, every 3-tree-decomposition defines for each inner vertex a unique path to each outer vertex. Directing every inner edge towards the father (in the tree to which it belongs) yields a 3-orientation of the considered triangulation.

We start with a preliminary straightforward lemma.

Lemma 1. *Let G be an outerplane graph. Either G is a triangle, or G has two non-adjacent vertices of degree at most* 2*.*

Proof. First, observe that it suffices to prove the result for 2-connected outerplane graphs (why?). So, let us assume that G is 2-connected. We define the *weak dual* T of G to be the graph obtained from the geometric dual G^* of G by removing the vertex corresponding to the outerface of G. The lemma is true if G is a cycle, so we assume that G is not a cycle. In particular, T has at least two vertices.

Let us show that T is a tree, which will yield the result. Indeed, T would have at least two leaves (since T has at least two vertices). Every leaf f of T corresponds to an inner face of G incident to at least one vertex of degree 2; furthermore, the vertices of degree 2 belonging to different such faces cannot be adjacent in G. Thus the lemma is proved provided we can show that T is a tree.

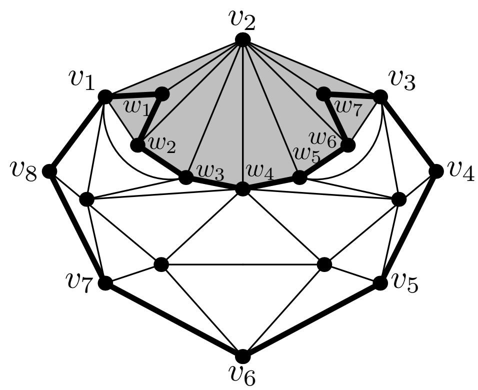

Figure 2: A near-triangulation: removing the vertex v_2 yields another near-triangulation.

One can realise that T is a tree using the duality between cycles of a graph and edge-cuts of its dual, as we saw in Lecture 1. We now give a proof using Euler's formula. Let N, E and F be the number of vertices, edges and faces of G, respectively. The graph G^* has F vertices and E edges. Let f be the vertex of G^* corresponding to the outerface of G. Since G is outerplanar, the degree of the vertex f in G^* is N. Consequently, the number of edges of T is $E - N = F - 2$ by Euler's formula. Since T has $F - 1$ vertices, this proves that T is a forest (and hence a tree, since G is connected). □

We repeatedly use the following observation (see Figure 2).

Observation 1. *Let $T = (V, E)$ be a 2-connected* near-triangulation *of the plane, i.e. T is a plane graph all of which inner faces are triangles (but the outer face may be a cycle of size greater than 3). Let $v_1, v_2, \ldots, v_k$ be the outer cycle of T, in clockwise order. Then, $T - v_2$ is a 2-connected near-triangulation with outer cycle $v_1, w_1, \ldots, w_d, v_3, \ldots, v_k$, where $w_1, \ldots, w_d$ are the inner vertices of T adjacent to v_2 in anti-clockwise order.*

Consider a plane triangulation $T = (V, E)$. We build a Schnyder labelling by inductively labelling and directing the inner edges. Let u, v and w be the three outer vertices of T, in clockwise order. At each step $i \geq 0$ of the algorithm, we consider a subgraph T_i of T, which is a 2-connected near-triangulation containing the edge uw. We will choose an outer vertex v_i of T_i different from u and w, such that v_i has exactly two neighbours that are outer vertices of T_i. Then, we orient and colour the edges of T_i incident to v_i, and set $T_{i+1} := T_i - v_i$. The exact procedure is described below, and illustrated in Figure 3.

First, we set $T_0 := T$ and $v_0 := v$. Then, **for** $i = 0$ **upto** $n-3$, we do the following.

The graph T_i is a subgraph of T, and T_i is a 2-connected near-triangulation containing the edge uw. Note that if we delete all the inner vertices of T_i, we obtain a

2-connected outerplane graph O_i. Thus, by Lemma 1, either O_i contains two non-adjacent vertices of degree 2, or O_i is a triangle. In both cases, O_i contains a vertex v_i of degree 2 that is different from u and w. Thus, in T_i, the vertex v_i has exactly two neighbours x_i and y_i that are outer vertices of T_i. We design by x_i the one that is closer to u on the outer cycle of T_i; and thus y_i is the one closer to w.

We orient all inner edges incident to v_i towards v_i, and we colour them red. We orient the edge v_ix towards x and colour it blue. Then, we orient the edge v_iy towards y and colour it green. These two steps are ignored when $i = 0$, since it would amount to colour and orient the outer edges of T, which is not needed for us. Last, we set $T_{i+1} := T_i - v_i$. If $i < n-3$, the graph T_{i+1} is a 2-connected near-triangulation of the plane by Observation 1, which contains the edge uw. We end the "for" loop.

Let us pause here to make some remarks.

- Another termination condition, which does not involve the number of vertices, is to end as soon as $i > 0$, $x_i = u$ and $y_i = w$.

- The algorithm is linear in time, even if no embedding is given since determining whether a graph is planar and finding a planar embedding can be done in linear time (consult [3, 5, 6, 10] to read more about this interesting topic).

- The procedure can be used to determine all possible Schnyder labellings of a plane triangulation (see the original paper [8]). In particular, there usually are more than one choice for the vertex v_i: according to how the choice is made the obtained Schnyder labelling can have some desirable properties, e.g. no clockwise cycle (see the thesis of Brehm [4]).

Let us show that the structure built by the algorithm is indeed a Schnyder labelling of T. To this end, let us state some invariants that are satisfied throughout the procedure. We first give some definitions. Let r be an outer vertex of T_i different from u and w. Say that the outer face of T_i is bounded by the (non-necessarliy induced) cycle $r_0 = u, r_1, \ldots, r_k = w$. For $i \notin \{0, k\}$, the *left neighbour* of r_i is r_{i-1}, and the right neighbour of r_i is r_{i+1}. If a and b are two edges incident with a vertex r, by *clockwise between a and b* we mean: start from a and walk clockwise towards b. We use *anti-clockwise between* in an analoguous way.

During the procedure, the following holds at the end of each step $i \geq 0$.

- the edges of T_{i+1} are neither coloured nor oriented;

- every vertex v_j with $j < i$ has exactly one blue out-going arc and one green out-going arc. All edges couter-clockwise between those two arcs are red and in-coming;

- every outer vertex r of T_{i+1} has exactly one out-going red arc $\vec{a}$ except u and w; further letting x and y be r's left and right neigbhours on the outer face of T_{i+1} respectively, all edges anti-clockwise between $\vec{a}$ and rx are green in-coming arcs, and all edges clockwise between $\vec{a}$ and ry are blue in-coming arcs.

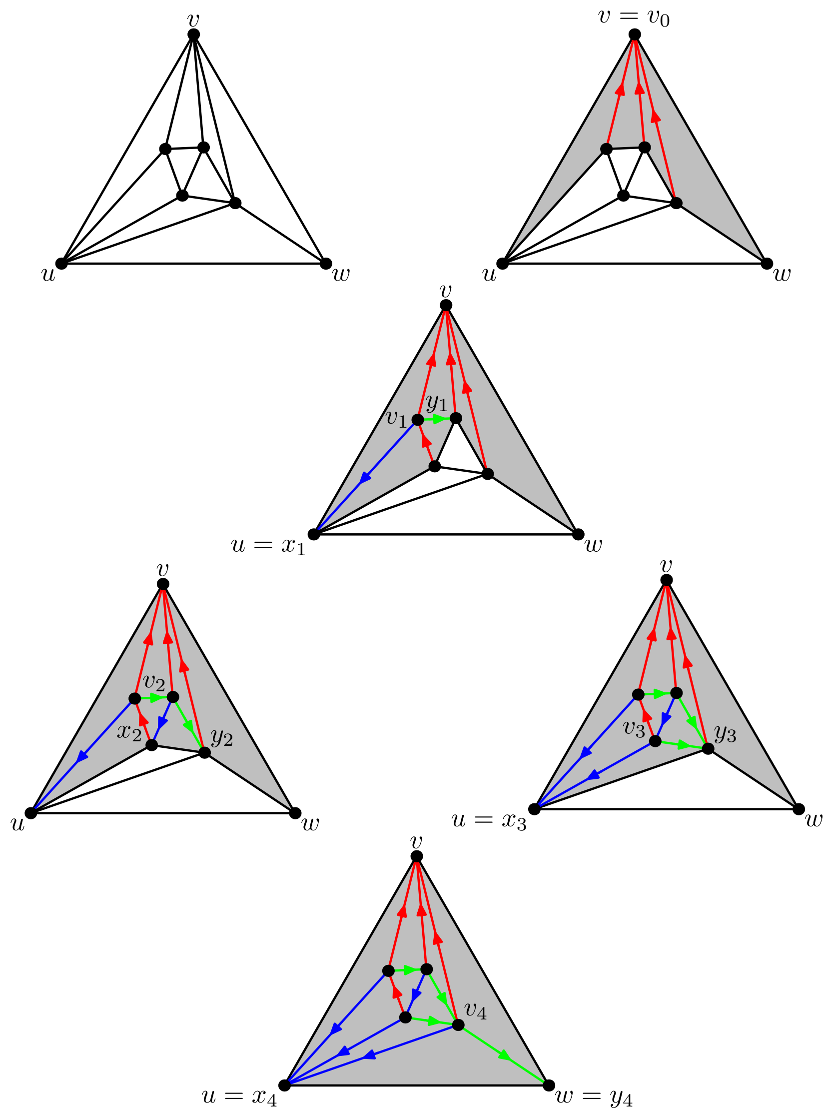

Figure 3: Schnyder's algorithm.

Assuming these invariants hold, the procedure yields a Schnyder labelling of T since each inner vertex of T is chosen as one of the vertices v_i for some $i \in \{1, 2, \ldots, n-3\}$, and the out-going blue and green arcs of v_i are its left and right neighbours on the outerface of T_{i+1}, respectively.

It remains to prove that the invariants stated above hold. This is true when $i = 0$: at the end of step 0, only the inner edges incident to $v_0 = v$ are oriented and/or coloured, thus no edge of $T_1 = T - v$ is oriented or coloured. Further, an outer vertex of T_1 distinct from u and w is necessarily an inner neighbour of v in T and hence has an out-going red arc. The other statements hold trivially.

Suppose now that the invariants are satisfied until step $i-1$ for some $i \in \{1, \ldots, n-3\}$, and let us show that they still hold at the end of step i. First, all edges that get oriented and coloured during step i are incident to v_i. Thus, since $T_{i+1} = T_i - v_i$ we deduce that no edge of T_{i+1} is coloured or oriented at the end of step i. Further, a vertex received a new incoming red arc during step i if and only if it is an inner vertex of T_i adjacent to v_i. So, if r is an outer vertex of T_{i+1} distinct from u and w, then either r is an inner vertex of T_i adjacent to v_i, or r is an outer vertex of T_i (distinct from u and w). In the former case, r had exaclty one out-going red arc at the end of step $i-1$ and receives no such arc during step i. In the latter case, no edge incident to r was coloured or oriented at the end of step $i-1$, and r received exactly one out-going red arc during step i.

During step i, only v_i receives an out-going blue or green arc, and v_i receives exactly one such arc. Moreover, v_i does not receive any blue or green in-coming arc at any step j with $j \geq i$, since during step j only x_j and y_j receive such arcs, and then they belong to the outer cycle of T_j and are distinct from v_j.

As a result, when v_i receives its out-going blue arc, it already received all its in-coming green arcs. Moreover, all the in-coming green arcs are anti-clockwise between the out-going red arc and the out-going blue arc (recall that the red arc out-going form v_j goes to an outer vertex of T_{i-1}). A similar argument shows that all blue-incoming arcs are clockwise between the out-going red arc and the out-going green arc. Last, v_i receives incoming red arcs only during step i, and all the arcs come from inner vertices of T_i. Since x_i and y_i are outer vertices of T_i, we infer that all these arcs are clockwise between the green arc $v_i \to y_i$ and the blue arc $v_i \to x_i$, as wanted.

□

We conclude with a (non-exhaustive and short) list of possible applications of Schnyder labellings.

- Schnyder's theorem [9], stating that a graph G is planar if and only if its associated poset $P(G)$ has dimension at most 3 ($P(G)$ is defined as the partial order with ground set $V(G) \cup E(G)$, and $x < y$ if and only if $x \in V(G)$, $y \in E(G)$ and y is incident to x in G).
- Graph drawing [2].
- Coding [1, 7].

References

[1] O. Bernardi and N. Bonichon. Intervals in Catalan lattices and realizers of triangulations. *J. Combin. Theory Ser. A*, 116(1):55–75, 2009.

[2] N. Bonichon, S. Felsner, and M. Mosbah. Convex drawings of 3-connected plane graphs. *Algorithmica*, 47(4):399–420, 2007.

[3] K. S. Booth and G. S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using *PQ*-tree algorithms. *J. Comput. System Sci.*, 13(3):335–379, 1976. Working Papers presented at the ACM-SIGACT Symposium on the Theory of Computing (Albuquerque, N. M., 1975).

[4] E. Brehm. 3-*Orientations and Schnyder* 3-*Tree-Decompositions*. Diploma thesis, Freie Universität Berlin, 2000.

[5] N. Chiba, T. Nishizeki, S. Abe, and T. Ozawa. A linear algorithm for embedding planar graphs using *PQ*-trees. *J. Comput. System Sci.*, 30(1):54–76, 1985.

[6] H. de Fraysseix, P. Ossona de Mendez, and P. Rosenstiehl. Trémaux trees and planarity. *Internat. J. Found. Comput. Sci.*, 17(5):1017–1029, 2006.

[7] D. Poulalhon and G. Schaeffer. Optimal coding and sampling of triangulations. In *Automata, languages and programming*, volume 2719 of *Lecture Notes in Comput. Sci.*, pages 1080–1094. Springer, 2003.

[8] W. Schnyder. Embedding planar graphs on the grid. In *SODA*, pages 138–148, 1990.

[9] Walter Schnyder. Planar graphs and poset dimension. *Order*, 5(4):323–343, 1989.

[10] W.-K. Shih and W.-L. Hsu. A new planarity test. *Theoret. Comput. Sci.*, 223(1-2):179–191, 1999.