anyt
[entry]anyt/global/

Chapter 14

Embedding planar graphs on a grid

By Sariel Har-Peled, April 8, 2018[1]

Version: 0.1

This is an early draft of a new chapter. Read at your own peril.

> It's getting them wrong that is living, getting them wrong and wrong and wrong and then, on careful reconsideration, getting them wrong again. That's how we known we're alive: we're wrong. Maybe the best thing would be to forget being right or wrong about people and just go along for the ride. But if you can do that - well, lucky you.
>
> American Pastoral, Philip Roth

Here, we show that every planar graph G with n vertices can be embedded as a straight-line graph with the vertices being on the grid, and furthermore, the grid is of size $[\![n]\!] \times [\![n]\!]$.

14.1. Canonical ordering and embedding on the grid

14.1.1. Canonical ordering

For a cycle C in a graph, a ***chord*** is an edge that connects two vertices that are not consecutive in C.

Lemma 14.1.1. *Let G be a simple planar graph that is embedded in the plane, and let $C = \mathfrak{r} = v_1, \ldots, v_k = \mathfrak{g}$ be a cycle of G (there are no repeated vertices in C). Then, there exist vertices $u, u' \in V(G) \setminus \{\mathfrak{r}, \mathfrak{g}\}$, such that u (resp., u') is not adjacent to any inside (resp., outside) chord of C.*

Proof: Here we consider the case of inner chords, as the outer case follows by a symmetric argument. If there are no chords in C then the claim is obvious. Otherwise, pick indices j, i such that $j > i$, $v_j v_i$ is a chord and $j - i$ is minimal. Since $v_j v_i$ is a chord, we have that $j - i \geq 2$. As such, let ℓ be any value between i and j, and observe that v_ℓ can not be adjacent to an inner chord, since such a chord must connect to a vertex of $C \setminus \{v_i, \ldots, v_j\}$, but such a chord must intersect the edge $v_i v_j$ by planarity. ■

We remind the reader that a *triangulation* is a maximal planar graph (i.e., all its faces are triangles).

Definition 14.1.2. Let G be a triangulation with n vertices in the plane, with its outer face vertices being $\mathfrak{r}, \mathfrak{g}$, and $\mathfrak{b}$. An ordering of the vertices $v_1 = \mathfrak{r}, v_3 = \mathfrak{g}, \ldots, v_n = \mathfrak{b}$ is a ***canonical ordering*** if the following conditions hold, for all $k > 2$:

(A) The induced subgraph $G_{k-1} = G_{\{v_1 \cup \ldots \cup v_{k-1}\}}$ is 2-connected, and its outer face is a cycle C_{k-1} containing the edge $\mathfrak{r}\mathfrak{g}$.
(B) v_k is in the (interior of the) outer face of G_{k-1}.
(C) v_k has at least two neighbors in G_{k-1}.
(D) v_k neighbors in G_{k-1} appears consecutively along the path $C_{k-1} - \mathfrak{r}\mathfrak{g}$.

An example of canonical ordering is depicted in Figure 14.1.

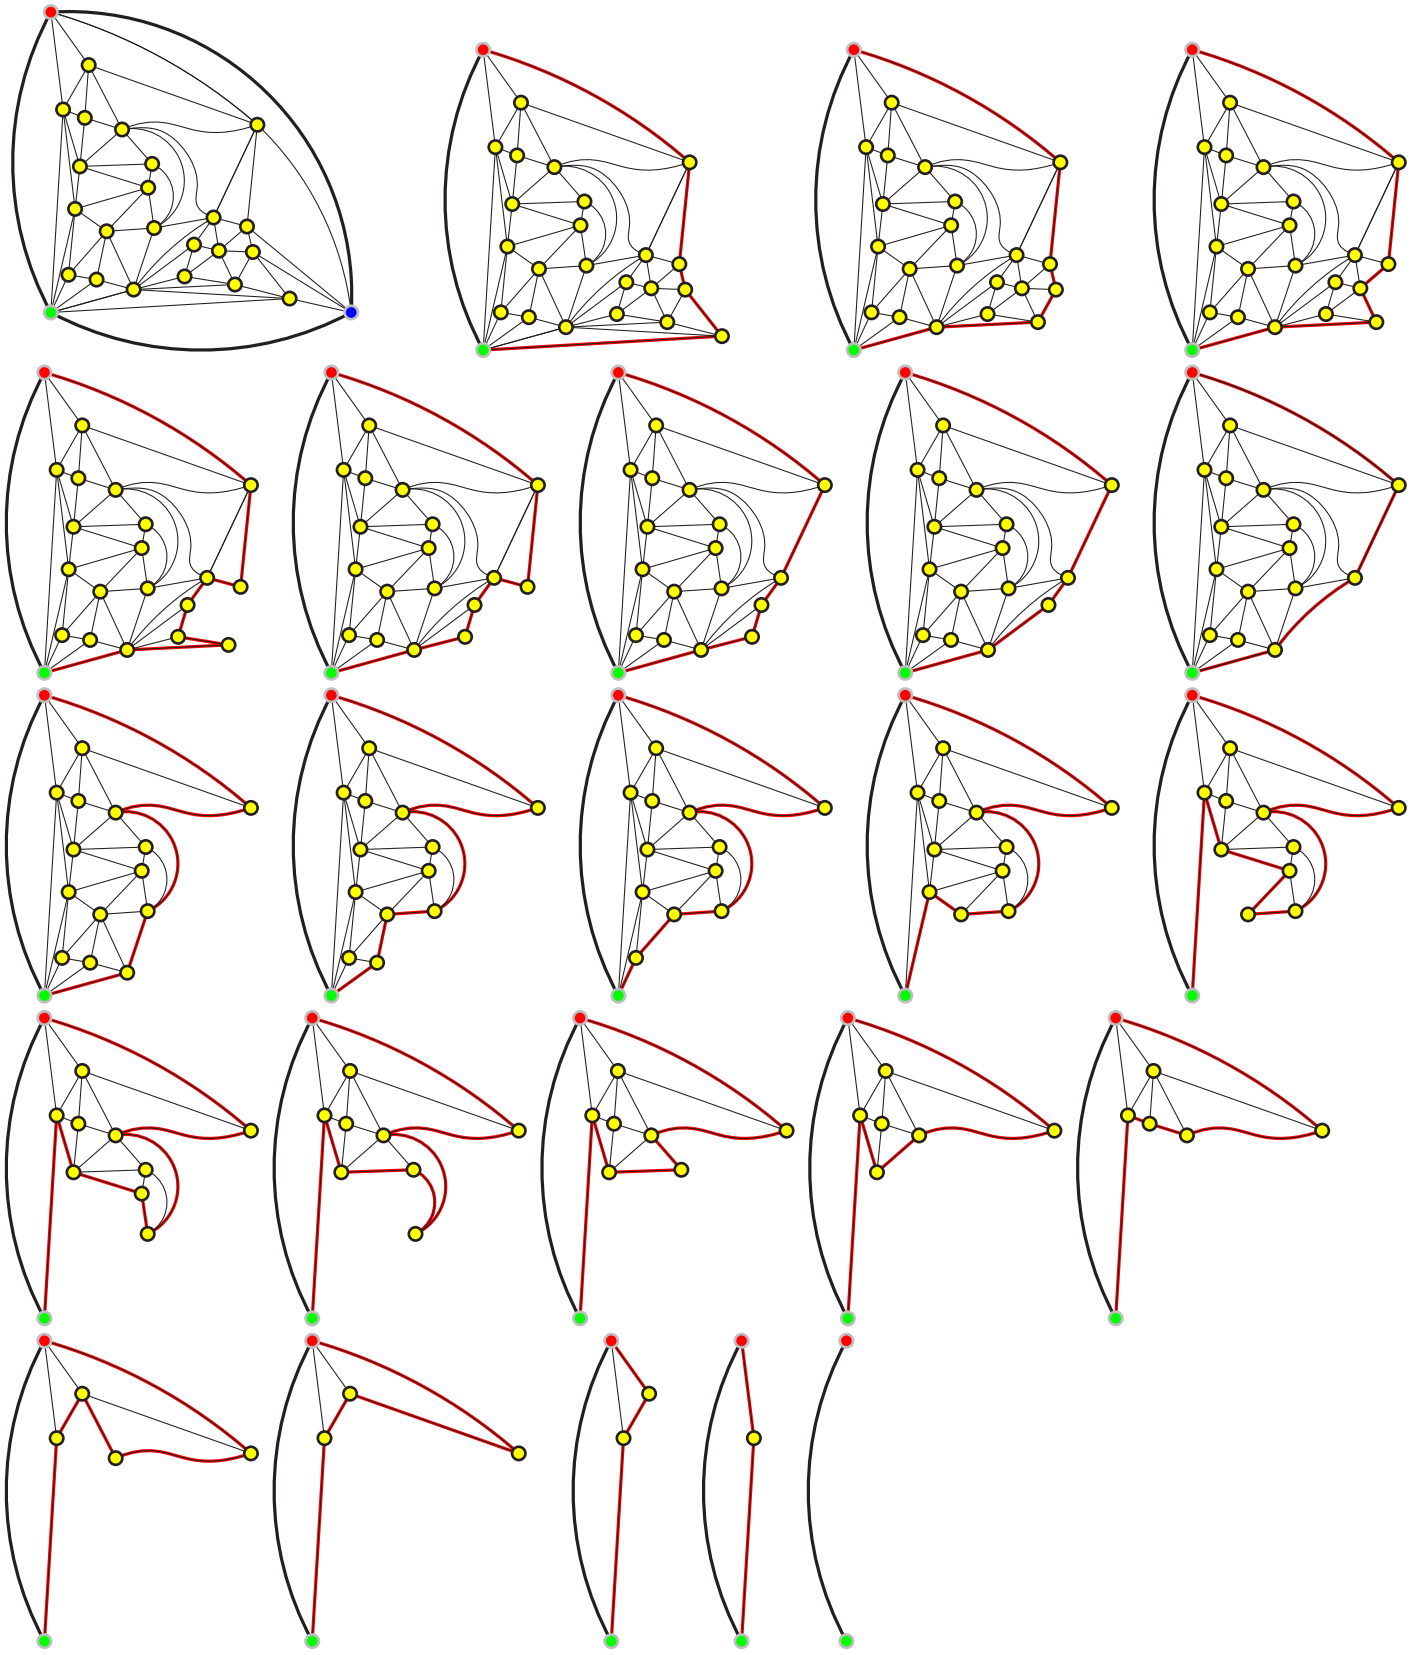

Figure 14.1: An example of canonical ordering for a triangulation.

Lemma 14.1.3. *Given a triangulation G with n vertices (together with its embedding), a canonical ordering of its vertices can be computed in linear time.*

Proof: The ordering is computed in reverse ordering. Initially, $v_n = \mathfrak{b}$ and clearly by the maximality of G the claim holds for it. Next, assume we have computed the suffix of the ordering $S_{k+1} = \{v_{k+1}, \ldots, v_n\}$, and it complies with the required conditions. In particular, let $G_k = G_{V(G)\setminus S_{k+1}}$. By the reverse induction, G_k is 2-connected, and has the edge $\mathfrak{rg}$ on its outer face. Let C_k be the boundary of the outer face of G_k. Since G_k is 2-connected, this cycle is simple (i.e., no repeated vertices), and it has no outer chords in the original graph G. Thus, applying Lemma 14.1.1 implies that there exists a vertex v_k, that is not adjacent to any inner chord of C_k, and v_k is not $\mathfrak{r}$ or $\mathfrak{g}$.

It is easy to verify that the desired properties hold. The only maybe non-trivial property is that the graph $G_{k-1} = G_{V(G)\setminus S_k}$ is 2-connected. Assume this is false, and let u be a vertex that its removal disconnects G_{k-1}. The vertex u must be on the outer face of G_{k-1}, and it must be connected to v_k in G_k. Furthermore, u must also appear in C_k, as otherwise its removal can not disconnect G_{k-1}. But then, the edge uv_k is an inner chord of C_k, contradicting the choice of v_k.

As for the algorithm implementation, each vertex maintains a counter with the number of its neighbors on the outer face as the vertices are being removed. A vertex that is on the outer face and with counter value of two, which is not $\mathfrak{r}$ or $\mathfrak{g}$, is a *candidate* to be removed. In each iteration one of the active candidate is being removed. When a vertex gets exposed and appear on the outer face, one scans its adjacency lists and increase the counter for all its neighbors, and furthermore, it marks itself as being on the outer face. Finally, the algorithm maintains a set with all the active candidates. This set can be easily maintained in constant time per operation – for example, by maintaining a list of active candidates, and for each vertex there is flag that marks if it is active, and if so it has a pointer to its location in the list of candidates. As such, the algorithm can retrieve a candidate to be removed in each iteration in constant time. The time to maintain the candidate list can be charged to the new exposed vertices and their adjacent edges in each iteration, which takes linear time overall. ■

Definition 14.1.4. Given a triangulation G and its canonical ordering, $v_1, v_2, \ldots, v_n$, with the outer boundary being $v_1 = u_1, \ldots, u_m = v_2$. Now, consider the vertex v_k, and its neighbors $u_i, \ldots, u_j$ in G_k. Then (i) u_i is the ***right child*** of v_k, (ii) $u_{i+1}, \ldots, u_{j-1}$ are the ***inner children***, and (iii) u_j is the ***left child*** of v_k.

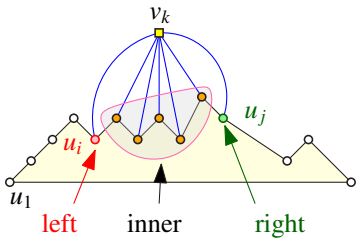

14.1.2. Canonical embedding

A nice implication of the canonical ordering is yet another proof that planar graphs have straight line embeddings. To this end, let G be a given triangulation, and let $v_1, v_2, \ldots, v_n$ be a canonical ordering of its vertices. Initially, we map v_1 to $p_1 = (0, 0)$ and v_2 to some arbitrary point p_2 on the x-axis to the right of p_1. In the ith iteration, the outer face of G_k, is bounded by a cycle C_k, which is embedded as an x-monotone polygon $\mathcal{P}_k$, having the segment p_1p_2 as its base. When adding the next vertex v_{i+1} in the canonical ordering, its neighbors are a sequence of vertices $u_i, \ldots, u_j$, with locations $q_i, \ldots, q_j$, respectively. Here, we have that $x(q_i) < x(q_{i+1}) < \cdots < x(q_j)$.

Think about the polygon $\mathcal{P}_k$ as an obstacle, and observe that the $q_i, \ldots, q_j$ are visible from above by a point that is sufficiently far away. In particular, pick an arbitrary coordinate x_k in the range $[x(q_i), x(q_j)]$, and let y_k be a number such that the point $p_k = (x_k, y_k)$ sees all the vertices of $q_i, \ldots, q_j$, but $(x_k, y_k - \varepsilon)$ does not, where ε is some arbitrary parameter.

Repeating the above for $k = 3, \ldots, n$, yields a straight line embedding. We thus get the following.

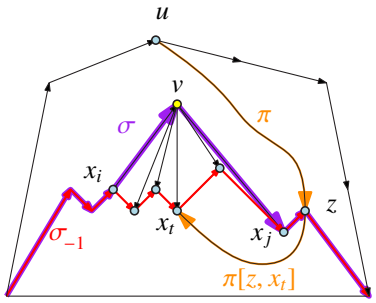

Figure 14.2

Claim 14.1.5. *The above algorithm computes, in linear time, to a straight line embedding of G. which is a* **canonical embedding** *of G.*

Remark. Note, that above canonical embedding is not uniquely defined, and furthermore, naively, the numbers needed to represent the embedded graph can be potentially quite large.

14.1.2.1. How to stretch an embedding

Given a partial (straight line) canonical embedding, we want to stretch it by moving some vertices to the right. Clearly, if we do it without care, it is easy to end up with an embedding that has crossing segments. Here, we develop mechanism to do such a stretching operation.

To this end, we define a DAG (directed acyclic graph) H_k that is a directed version of G_k, for all k. In particular, the invariant kept is that all the edges on the boundary of the embedding (i.e., edges of C_k) are directed from left to right.

Initially, we direct the first edge $v_1 \to v_2$. In the kth iteration, assume that $C_k \equiv \mathfrak{r} = u_1, \ldots, u_m = \mathfrak{g}$ form the boundary of the outer face of G_k, ordered in the given canonical embedding from left to right, and assume that v_{k+1}, the next vertex in the canonical ordering, is connected to $u_i, \ldots, u_j$. Then, in the new graph H_{k+1}, the new directed edges are $u_i \to v_{k+1}$ (i.e., left child goes into v_{k+1} and $v_{k+1} \to u_{i+1}, \ldots, v_{k+1} \to u_{j-1}, v_{k+1} \to u_j$ (i.e., edges go from v_{k+1} to the inner children and to the right child).

It is easy to verify that this is indeed a DAG, as any path from two vertices x and y in H_{k+1} that uses the new vertex v_{k+1} as an internal vertex, can be realized as a path from x to y in the earlier graph H_k.

For two vertices, $x, y \in V(H_k)$, the vertex x ***dominates*** y, if there is a path in H_k from x to y. In particular, let $D_k(x)$ be the set of vertices of G_k that x dominates (note, that $x \in D_k(x)$. It is easy to verify that the following properties hold (using the above notations):

(I) $u_j \in D_k(u_i)$ if $j \geq i$.
(II) $D_k(u_1) \supset D_k(u_2) \supset \cdots \supset D_k(u_m)$.

Observation 14.1.6. *For a vertex $x \in H_k$, let $R_k(x)$ be the set of vertices of C_k that dominates x, where $C_k \equiv \mathfrak{r} = u_1, \ldots, u_m = \mathfrak{g}$ is the boundary of the outer face of G_k. Then $R_k(x)$ is a prefix $u_1, u_2, \ldots u_{m'}$ of C_k. This follows readily as the edges of C_k are directed from left to right in the embedding of H_k.*

Lemma 14.1.7. *For a vertex $v \in H_k$, let $x_i, \ldots, x_j$ be its children when it was added by the canonical ordering. Then, we have that $R_k(x_i) \subseteq R_k(v) = R_k(x_{i+1}) = \cdots = R_k(x_{j-1}) \subseteq R_k(x_j)$.*

Proof: The proof is illustrated in Figure 14.2. Since $x_i \to v \in E(H_k)$, we have $R_k(x_i) \subseteq R_k(v)$. Similarly, $R_k(v) \subseteq R_k(x_\ell)$, for $\ell = i + 1, \ldots, j$.

We are left with the task of proving the equality. Assume, for the sake of contradiction, that there is a vertex $u \in R_k(x_t) \setminus R_k(v)$, for some t in the range $i+1, \ldots, j-1$. Let π be the path in H_k from u to x_t, and let σ (resp. σ_{-1}) be the directed chain from v_1 to v_2 that forms the outer face of H_α (resp.$H_{\alpha-1}$), where v is the αth vertex in the canonical ordering. If π intersects σ before v appears in σ, then there is a path in H_{k+1} from u to v, and $u \in R_k(v)$, a contradiction. As such, it must be that π intersects σ at a vertex z after v. But then the path $\pi[z, x_t]$ (i.e., the portion of the path π from z to v) together with $\sigma_{-1}[x_t, z]$ form a cycle in H_{k+1}, contradicting that H_{k+1} is a DAG. ■

So, consider a canonical embedding of G_k, with the outer cycle $C_k \equiv v_1 = u_1, \ldots, u_m = v_2$. A ***chain translation*** is a sequence α of numbers $\alpha_1, \ldots, \alpha_m \geq 0$ assigned to the vertices of C_k.

The size of the induced ***translation*** for a vertex $x \in V(G_k)$ is

$$\Delta(x) = \sum_{i : x \in D_k(u_i)} \alpha_i.$$

In words, we translate the location of x to the right, where the translation is the total sum of the translations of all the top chain vertices that dominates it. In particular, for the ith vertex in the top chain u_i, we have that $\Delta(u_i) = \sum_{j=1}^{i} \alpha_j$.

Lemma 14.1.8. *Given a canonical straight line embedding of G_k, with a top chain $C_k \equiv v_1 = u_1, \ldots, u_m = v_2$, and a chain translation $\alpha_1, \ldots, \alpha_m \geq 0$, the embedding resulting from translating every vertex $x \in V(G_k)$ by distance $\Delta(x)$ to the right, is a valid straight line canonical embedding of G_k.*

Proof: Consider a triangle $\triangle$ in the embedding of G_k, and let v_i be the highest index vertex (in the canonical ordering) that is a vertex of $\triangle$. As such, $\triangle$ was created when v_i was added. Let z_1, z_2 be the two other vertices of $\triangle$. If z_1 and z_2 are inner children of v, then Lemma 14.1.7 implies that $R_k(v_i) = R_k(z_1) = R_k(z_2)$, and as such $\Delta(z_1) = \Delta(z_2) = \Delta(v_i)$. Namely, the new embedding just translates the triangle to the right rigidly, and it is still a valid triangle.

Similarly, if z_1 is the left child of v_i, then Lemma 14.1.7 implies that $R_k(z_1) \subseteq R_k(v_i) \subseteq R_k(z_2)$, which implies that $\Delta(z_1) \leq \Delta(v_i) \leq \Delta(z_2)$. In addition, we have that in the current embedding $x(z_1) < x(v_i)$ and $y(z_1), y(z_2) < y(v_i)$ (this are all properties of canonical embedding). Which implies that this triangle had not flipped after the new translation is applied, and the same properties hold for the new locations of the vertices. The case that z_2 is the right child of v_i follows by a similar argument.

As such, all the triangles in the embeddings of G_k are preserved under this translation. As such, the result is a valid straight line embedding, and it is easy to verify that it is a canonical embedding. ■

14.1.3. Embedding on the grid using canonical ordering

Let G be a triangulation with n vertices (its embedding is also provided), with the other face vertices being $\mathfrak{r}, \mathfrak{g}$ and $\mathfrak{b}$. Let $v_1 = \mathfrak{r}, v_2 = \mathfrak{g}, \ldots, v_{n-1}, v_n = \mathfrak{b}$ be a canonical ordering of G.

Here, we are going to show how to embed the G on a small grid. The idea is to use canonical embedding together with the above stretching mechanism. So, we add the vertices of the canonical ordering. In the end of the kth iteration, we have an embedding g_k of the graph G_k with the following properties:

(A) The embedding g_k maps all the vertices of G_k to integer points in the rectangles $[0, 2k-4] \times [0, k]$.

(B) The base segment v_1v_2 (i.e., $\mathfrak{rg}$) is mapped to the interval $[0, 2k-4]$ on the x-axis. That is $g_k(v_1) = (0, 0)$ and $g_k(v_2) = (2k-4, 0)$.

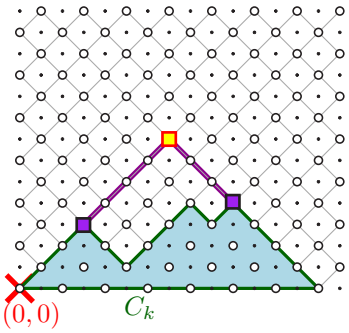

Figure 14.3: The vertices of the outer face of G_k lies on the even grid G^0.

(C) The boundary C_k of the exterior face of G_k is an x-monotone polygonal curve above $g_k(v_1v_2)$. Formally, if $C_k \equiv \mathfrak{r} = u_1, u_2, \ldots, u_{m-1}, u_m = \mathfrak{g}$, then in the embedding g_k we have

$$x\big(g_k(u_1)\big) < x\big(g_k(u_2)\big) < \cdots < x\big(g_k(u_m)\big). \tag{14.1}$$

In particular, the embedding of G_k is contained in the polygon $g_k(C_k)$.

(D) The segments $g_k(u_iu_{i+1})$ have slope either $+1$ or -1, for all i.

Let G^0 (resp. G^1) be the set of all integer grid points such that the sum of their coordinates is even (resp. odd). Observe that lines with slope ± 1 that passes through points of G^0 do not contain any point of G^1 (and vice versa). As such, starting at $g_k(u_1) = (0,0)$ and tracing $g_k(C_k)$, it follows that $g_k(u_1), \ldots, g_k(u_m) \in \mathsf{G}^0$, by condition (D) above. This also implies that the Manhattan distance[②] between $g_k(u_i)$ and $g_k(u_j)$ is even (that is, $\|g_k(u_i) - g_k(u_j)\|_1$ is even) for any i, j.

Let v_{k+1} be the next vertex in the canonical ordering, and assume it connected to $u_i, \ldots, u_j$. Consider the ℓ^+ of slope 1 passing through $g_K(u_i)$, and the line ℓ^- of slope -1 that passes through $g_k(u_j)$. The intersection point $p = \widehat{g_k}(u_i, u_j) = \ell^+ \cap \ell^- \in \mathsf{G}^0$ is the natural location for the v_{k+1}. Clearly, the segments $pg_k(u_{i+1}), \ldots, pg_k(u_{j-1})$ are all valid. Furthermore, if the segments $pg_k(u_i)$ and $pg_k(u_j)$ do not intersect the interior of the segments $g_k(u_iu_{i+1})$ and $g_k(u_{j-1}u_{j-1})$, respectively, then this is a valid location for v_{k+1} preserving all the desired properties. The algorithm updates the embedding g_{k+1} to map v_{k+1} to p and continues to the next iteration. See Figure 14.3.

Otherwise, we apply stretching to $g_k(G_k)$ as follows. We set the numbers $\alpha_1, \ldots, \alpha_m$ all to zero except for $\alpha_{i+1} = \alpha_j = 1$. Let g'_k be the resulting embedding after applying the induced translations, as described in Lemma 14.1.8. Note, that u_i is in the same location as before, and u_j has moved to the right by distance 2. As such, $g'_k(u_i), g'_k(u_j) \in \mathsf{G}^0$, but importantly, the segment $g'_k(u_iu_{i+1})$ has slope strictly smaller than 1, and the segment $g'_k(u_iu_{i+1})$ has slope strictly larger than -1. As such, the new location $p' = \widehat{g'_k}(u_i, u_j) \in \mathsf{G}^0$ introduces two new segments to the left and right child, respectively, that no longer intersects the top chain of $g'_k(G_k)$. As such, we set $g_{k+1}(v_{k+1}) = p'$, and otherwise it equal to g'_k. Clearly, the resulting embedding is a canonical embedding with all the top chain vertices belonging to G^0.

Inserting v_{k+1} had increase the width of the embedding by at most two, and it is not hard to verify that $y(p') \le k + 1$. As such, we get the desired properties.

Running time. Naively, the running time is quadratic, but a more careful implementation works in linear time, see the bibliographical notes.

[②] Also known as the Shampoo-Banana distance.

14.1.3.1. The result

Theorem 14.1.9. *Given a maximal planar graph G with n vertices, one can embed its vertices on the integer grid $[\![2n]\!] \times [\![n]\!]$, where $[\![n]\!] = \{0, 1, \ldots, n-1\}$. This embedding can be computed in linear time.*

Remark 14.1.10. Interestingly, the above algorithm does not require having the embedding of G. It is enough to designate the outer face, and the rest of the embedding is computed implicitly from the above. This implies, for example, that maximal planar graphs, up to designating what is the outer face, are uniquely defined.

14.2. Breaking a triangulation into oriented trees

In the following, a ***triangulation*** is a maximal simple planar graph given together with its embedding.

Consider a graph G and an edge uv in it. Consider the graph resulting from merging u and v together, removing any extra parallel edges that might have been created or self loops. The resulting graph, denoted by G/uv, is the result of the ***edge contraction*** of uv. If we do the contraction over a planar embedding, of a triangulation, then not all edges can be contracted safely. This is illustrated in Figure 14.5.

Definition 14.2.1. An edge uv in a triangulation G is ***contractible*** if u and v have exactly two common neighbors.

Lemma 14.2.2. *Let G be a triangulation with $n > 3$ vertices, and let y be any of the three outer vertices of G. Then, there exists an internal vertex u in G such that yu is contractible.*

Proof: Let $\pi \equiv u_1, \ldots, u_k$ be the neighbors of y in G ordered in counterclockwise order around y (i.e., u_1 and u_k are the two other outer vertices of G) – see Figure 14.4. The proof is now by induction on k. If $k = 3$, then the claim holds immediately for yu_2, as y is of degree three (indeed, any edge adjacent to a vertex of degree three can be contracted). Otherwise, pick an arbitrary vertex u_i, for some i such that $1 < i < k$. If u_i has only its two immediate neighbors as neighbors in the chain π, then it is the desired vertex.

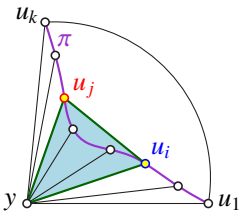

Figure 14.4

Otherwise, there is an index j (say $j > i$ the other case is handled similarly), such that $u_iu_j \in E(G)$, and $j - i > 1$. The claim now follows by applying induction on y and the subchain $u_i, \ldots, u_j$. See Figure 14.4. ■

In the following, given a triangulation G, with n vertices, we use $\mathfrak{r}, \mathfrak{g}$ and $\mathfrak{b}$ to denote the three outer vertices in G, in their counterclockwise order along the outer triangle. Let $H_1 = G$. For $i \geq 1$, let v_i be a vertex that has an edge to one of the outer vertices $y_i \in \{\mathfrak{r}, \mathfrak{g}, \mathfrak{b}\}$, such that v_iy_i can be contracted (such a vertex exist by Lemma 14.2.2 above). Let $H_{i+1} = H_i/v_iy_i$. The reverse sequence $v_{n-3}, v_{n-2}, \ldots, v_1$ is an ***expansion sequence*** for G. Such an expansion sequence starts with a triangle $\triangle\mathfrak{r}\mathfrak{g}\mathfrak{b}$, and in each iteration it adds a new vertex and three edges to the graph by expanding an edge from one of the outer vertices. In the end of this expansion process we end up with the original graph G. Such an expansion process is depicted in Figure 14.7.

Lemma 14.2.3. *Given a triangulation G, one can color and orient the edges of G, such that the internal edges of G decomposes into three oriented trees, $\mathcal{T}_\mathfrak{r}, \mathcal{T}_\mathfrak{g}, \mathcal{T}_\mathfrak{b}$, their edges colored by **red**, **green**, and blue, respectively, such that:*

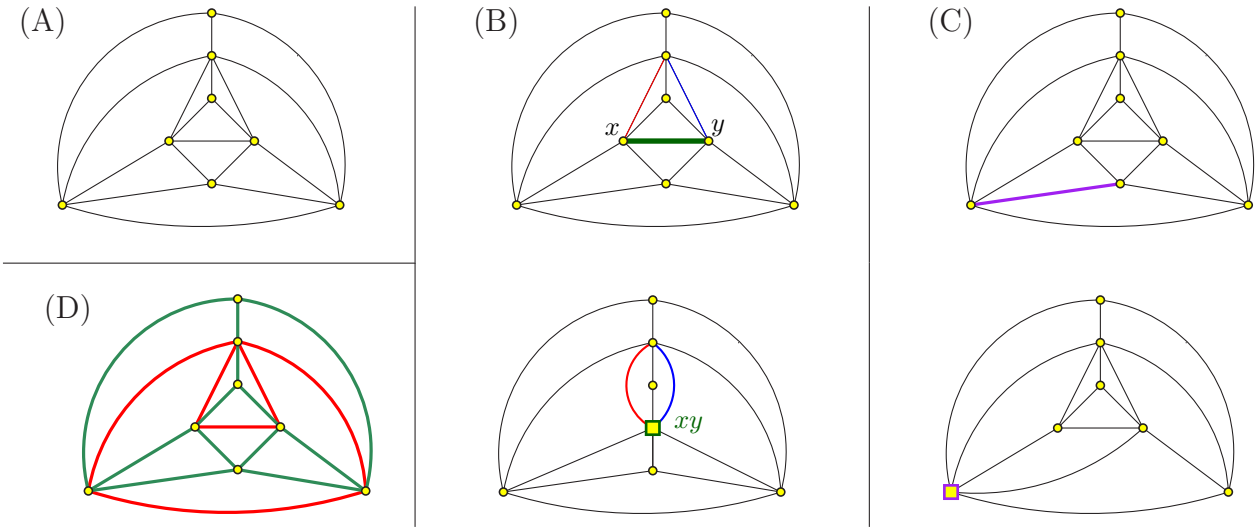

Figure 14.5: (A) A graph. (B) An unsafe edge contraction may lead to parallel edges that are different in the planar embedding, and cannot just be simply removed. (C) A safe contraction. (D) Classification of edges into safe (green) and unsafe (red) edges.

(A) Every internal vertex has one outgoing edge in each of the three colors.
(B) The trees are rooted at their respective three outer vertices.
(C) The colors of the three outgoing edges around an internal vertex, in counterclockwise circular order, are ***red****, green, and blue.*
(D) All the edges adjacent to the outer vertex $\mathfrak{r}$ *are red in-edges (similar claims holds for the* $\mathfrak{g}$ *and* $\mathfrak{b}$ *outer vertices).*
(E) Colors of incoming edges into an internal vertex v *comply with (see Figure 14.6):*

(a) All ***red*** *in-edges of into* v *are between the* ***green*** *and* ***blue*** *out-edges (in counterclockwise order). (b) All* ***green*** *in-edges into* v *are between the* ***blue*** *and* ***red*** *out-edges. (c) All* ***blue*** *in-edges into* v *are between the* ***red*** *and* ***green*** *out-edges.*

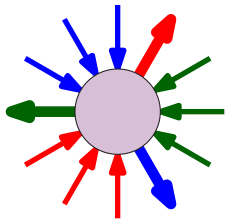

Figure 14.6: Legal pattern around a vertex. See Lemma 14.2.3.

Proof: Consider an expansion sequence $x_1, \ldots, x_{n-3}$. The trees are constructed iteratively from this sequence. Let G_0 be the triangle formed by $\mathfrak{r}, \mathfrak{g}, \mathfrak{b}$. In the ith iteration, consider the vertex x_i, and the outer vertex $u_i \in \{\mathfrak{r}, \mathfrak{g}, \mathfrak{b}\}$ it was contracted into. Expanding the edge $x_i u_i$ in G_{i-1} creates the graph G_i. By this end of this process we end up with the original graph G.

Each edge expansion creates three new edges in the graph, around the new vertex x_i. We orient these three edges away from x_i. The edge $x_i \to u_i$ is colored in the color of the outer vertex u_i. Next, color the two other new edges, so that the going around u_i in a (circular) counterclockwise order, these three outgoing edges, are colored $\mathfrak{r}, \mathfrak{g}, \mathfrak{b}$.

It is now easy to verify, by induction, that the edges of each color class form a tree rooted in the respective outer vertex. The other claims follow in a similar straightforward fashion.

Claim (E) about the color patterns of the incoming edges can be similarly proved using induction. It obviously holds for a new vertex. Now, an existing vertex v gets a new incoming edge, during the expansion process, only if it is attached to one of the outer vertices, say $\mathfrak{r}$, and the new vertex u added is also attached to $\mathfrak{r}$. If the new edge $u \to \mathfrak{r}$ is counterclockwise to $v \to \mathfrak{r}$, then the new edge is colored

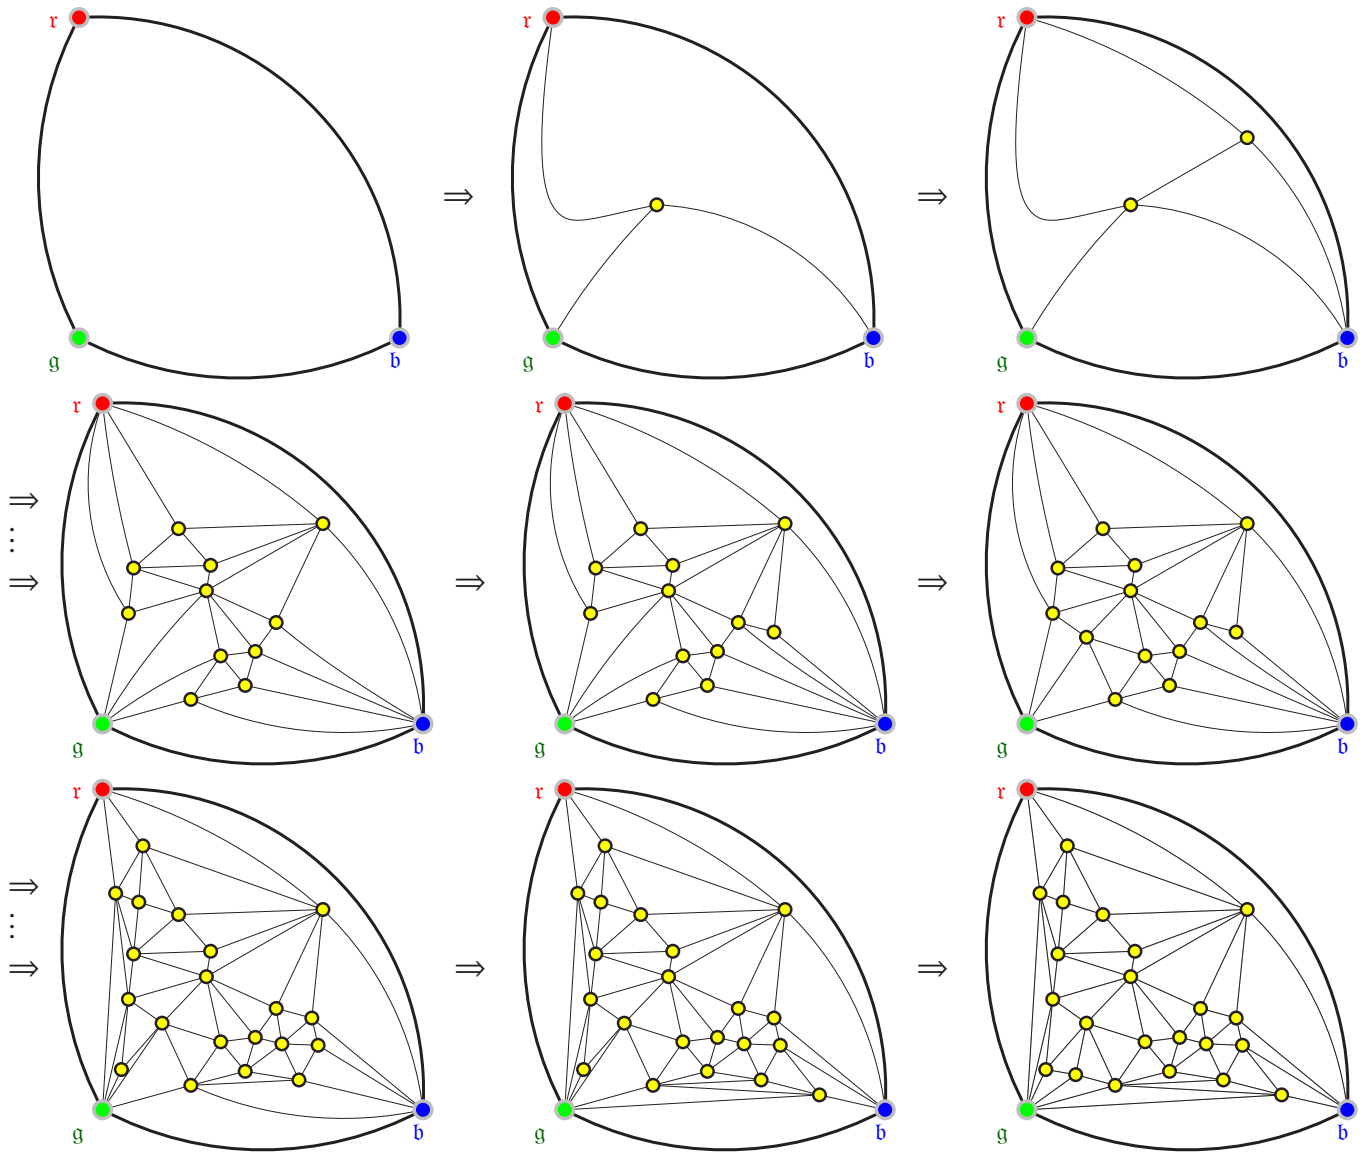

Figure 14.7: Snapshots of an expansion process of a triangulation.

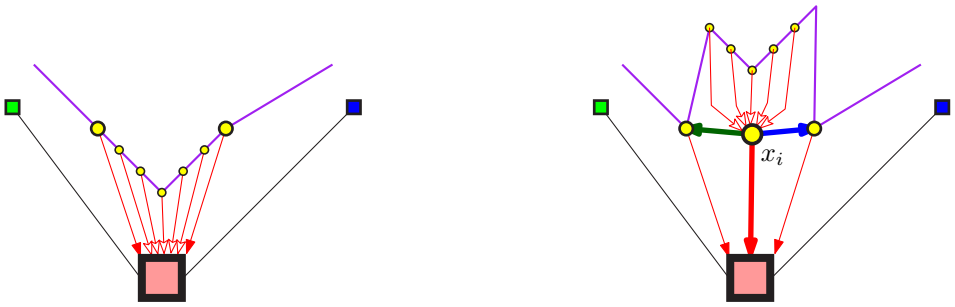

Figure 14.8: An expansion of an edge – a consecutive group of incoming red edges get reassigned to the new vertex, and three new out-edges are introduced from the new vertex x_i.

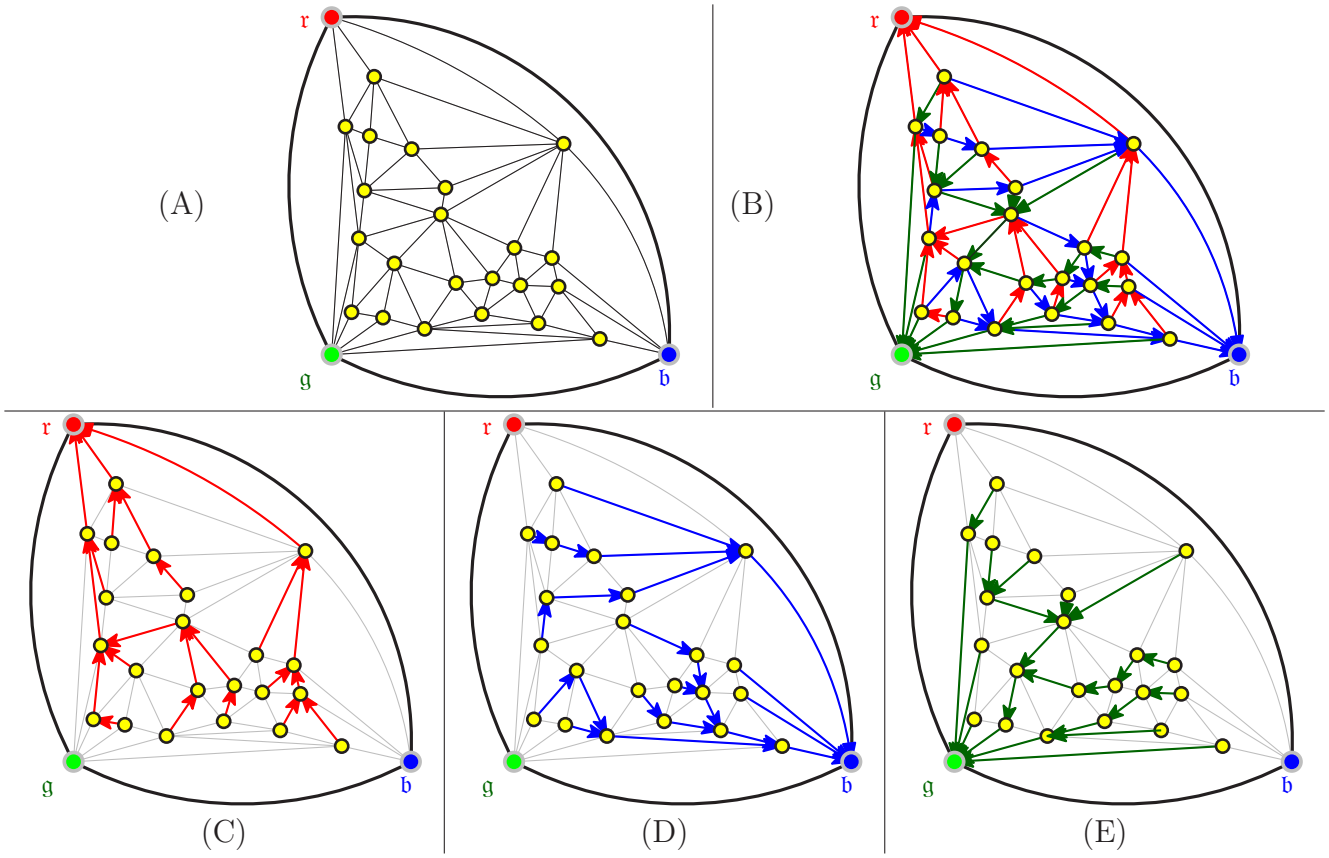

Figure 14.9: (A) A triangulation. (B) The decomposition into three directed trees. (C) Each tree separately.

g, otherwise the new edge is colored b, see Figure 14.8. This clearly complies with the color pattern restrictions. Now, if the claim on the color pattern is correct before adding this new edge, it is going to be correct after adding the new edge. The other cases where the outer vertices are g and b are handled in a similar fashion. ■

Definition 14.2.4. A decomposition of a triangulation G into three oriented trees $(\mathcal{T}_{\mathfrak{r}}, \mathcal{T}_{\mathfrak{g}}, \mathcal{T}_{\mathfrak{b}})$ that comply with the conditions of Lemma 14.2.3 is a ***realizer***.

Lemma 14.2.5. *Let G be a triangulation with a realizer $(\mathcal{T}_{\mathfrak{r}}, \mathcal{T}_{\mathfrak{g}}, \mathcal{T}_{\mathfrak{b}})$. For an internal vertex v of G, let $\pi(v, \mathfrak{r})$, $\pi(v, \mathfrak{g})$, and $\pi(v, \mathfrak{b})$ be the paths in $\mathcal{T}_{\mathfrak{r}}, \mathcal{T}_{\mathfrak{g}}, \mathcal{T}_{\mathfrak{b}}$ from v to the root nodes $\mathfrak{r}, \mathfrak{g}, \mathfrak{b}$, respectively. Then, these paths are interior disjoint.*

Proof: The proof is by contradiction. Let i be the first iteration of the expansion process such that the graph generated G_i no longer has the desired property. Let v be the vertex inserted in this iteration. Clearly, by minimality, it must be that v is the vertex for which the property fails, as otherwise, i would not be minimal.

Here v was inserted by expanding it (say from $\mathfrak{r}$), and the three paths $\pi(v,\mathfrak{r})$, $\pi(v,\mathfrak{g})$, $\pi(v,\mathfrak{b})$ are not interior disjoint. The path $\pi(v,\mathfrak{r})$ is the edge $v \to \mathfrak{r}$, so it can not participate in an interior intersection. As such, it must be that $\pi(v,\mathfrak{g})$ and $\pi(v,\mathfrak{b})$ share an interior vertex, say u.

By the color patterns allowed around a vertex, see Figure 14.6, it must be that the paths are locally crossing at u, which means that immediately after u, the path $\pi(v,\mathfrak{b})$ is inside the region bounded by the Jordan curve formed by $\pi(v,\mathfrak{r}) \cup \mathfrak{r}\mathfrak{g} \cup \pi(v,\mathfrak{g})$. This implies that $\pi(v,\mathfrak{b})$ must have a second intersection vertex, denoted by x, with $\pi(v,\mathfrak{g})$. See Figure 14.10.

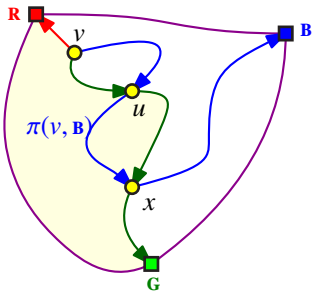

Figure 14.10

However, this implies that u violates the desired property as the paths $\pi(u,\mathfrak{g})$ and $\pi(u,\mathfrak{b})$ have an interior intersection (i.e., x), and this intersection is present in an earlier iteration than i. This contradicts the choice of i. ■

Lemma 14.2.5 implies that given a realizer $(\mathcal{T}_\mathfrak{r}, \mathcal{T}_\mathfrak{g}, \mathcal{T}_\mathfrak{b})$, and a vertex v, the (bounded) region bounded by $\triangle\mathfrak{r}\mathfrak{g}\mathfrak{b}$ is partitioned into three regions by the paths $\pi(v,\mathfrak{r})$, $\pi(v,\mathfrak{g})$, and $\pi(v,\mathfrak{b})$. The region not adjacent to $\mathfrak{r}$ (resp. $\mathfrak{g}$ and $\mathfrak{b}$) is denoted by $\Delta\mathfrak{r}(v)$ (resp. $\Delta\mathfrak{g}(v)$ and $\Delta\mathfrak{b}(v)$). See Figure 14.11 (A). For our purposes, boundary vertices (e.g., vertices that belong to $\pi(v,\mathfrak{r})$) are in both regions that they bound (e.g., $\Delta\mathfrak{g}(v)$, $\Delta\mathfrak{b}(v)$).

Lemma 14.2.6. *Consider the region $\Delta\mathfrak{r}(v)$, and a vertex $u \in V(G)$, such that $u \neq v$, and u lies in $\Delta\mathfrak{r}(v)$. Then $\Delta\mathfrak{r}(u) \subsetneq \Delta\mathfrak{r}(v)$. Similar claims holds for $\Delta\mathfrak{g}(\cdot)$ and $\Delta\mathfrak{b}(\cdot)$.*

Proof: Assume u is in the interior of $\Delta\mathfrak{r}(v)$ (if it is on the boundary a similar argument works). Consider the path $\pi(u,\mathfrak{g})$ and assume, for the sake of contradiction, that it intersects the path $\pi(v,\mathfrak{b})$ at a vertex z, see Figure 14.11 (B). However, at z, the green and blue oriented paths cross in an illegal pattern of colors. Indeed, by Figure 14.6, as a green path traverses a vertex, the blue path comes from the right and exists on the left, but here the situation is in reverse, which is not possible.

As such, $\pi(u,\mathfrak{g})$ intersects only $\pi(v,\mathfrak{g})$ (potentially at the outer vertex $\mathfrak{g}$), and $\pi(u,\mathfrak{b})$ intersects only $\pi(v,\mathfrak{b})$, which readily implies the claim. See Figure 14.11 (C). ■

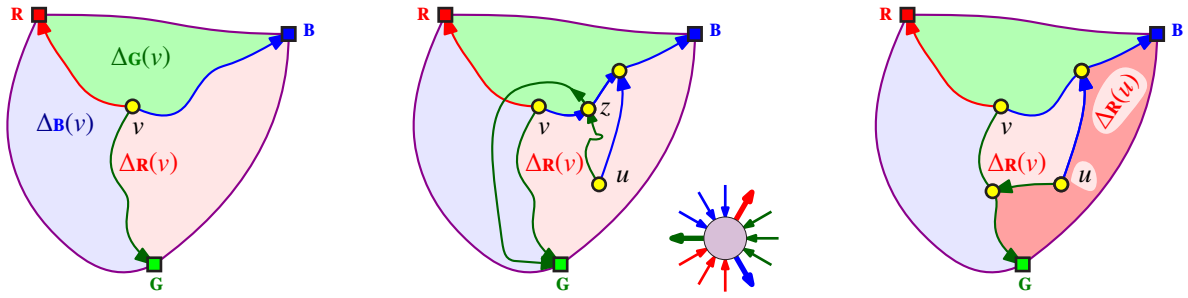

Figure 14.11: (A) The three regions. (B) A crossing has an invalid color pattern. (C) Monotonicity of the regions.

14.3. Grid embedding via realizers

14.3.1. Embedding via barycentric coordinates

We had sowed the trees of the realizers of a triangulation, and we shell now reap the straight line embedding on a grid from them. The idea is somewhat surprising – given a realizer $(\mathcal{T}_\mathfrak{r}, \mathcal{T}_\mathfrak{g}, \mathcal{T}_\mathfrak{b})$, it induces a natural mapping of the vertices of a graph to barycentric coordinates.

Definition 14.3.1. A point $p = (p_1, \ldots, p_d) \in \mathbb{R}^d$ is in ***barycentric coordinates*** if $p_1, \ldots, p_d \geq 0$ and $\sum_i p_i = 1$. Namely, p encodes a convex combination of d objects. Let Δ_d denotes the set of all points with barycentric coordinates in $\mathbb{R}^d$

Observe that $\Delta_3 = \mathcal{CH}\Big((1,0,0).(0,1,0),(0,0,1)\Big)$. For a mapping $f : V(G) \to \Delta_3$, we use $f_i(v)$ to denote the ith coordinate of $f(v)$. Given three points p_1, p_2, p_3 in the plane, barycentric coordinates yield a natural parameterization of the triangle $\triangle p_1 p_2 p_3$, where a point $(x_1, x_2, x_3) \in \Delta_3$, is interpreted as the point $\sum_i x_i p_i$. Under this interpretation, the segment in Δ_3 that is $(x_1 = \alpha) \cap \Delta_3$, corresponds to a segment in $\triangle p_1, p_2, p_3$ parallel to $p_2 p_3$, such that its two endpoints are $\alpha p_1 + (1-\alpha) p_2$ and $\alpha p_1 + (1-\alpha) p_3$. See Figure 14.12.

In the following, we use $f(x) <_i f(y)$ (resp $=_i$) to express that the ith coordinate of $f(x)$ is smaller than (resp. equal to) the ith coordinate of $f(y)$. Clearly, for x, y, z if $f(x) <_i f(z)$ and $f(y) <_i f(z)$ then, for any point p on the segment $f(x)f(y)$, we have that $p <_i f(z)$. We denote this by $f(xy) <_i f(z)$.

Definition 14.3.2. A ***barycentric representation*** of a graph G is an injective function $f : V(G) \to \Delta_3$, such that for any edge $xy \in E(G)$, and any vertex $z \in V(G) \setminus \{x, y\}$, there is an index i, such that $f(xy) <_i f(z)$.

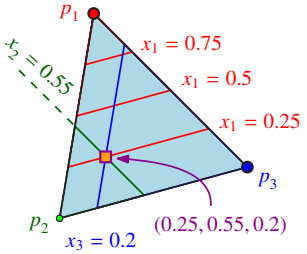

Figure 14.12

Surprisingly, once one has a barycentric coordinates, it leads immediately to a straight-line embedding. In particular, we have the following (we prove a more general version of this lemma below – see Lemma 14.3.6).

Lemma 14.3.3. *Let f be a barycentric representation of G. For any three points p_1, p_2, p_3 in the plane that are not collinear, the mapping $g(v) = \sum_i p_i f_i(v)$ induces a straight line embedding of G in the plane.*

14.3.2. From realizers to barycentric coordinates

Consider a vertex $v \in V(G)$, and its associated regions $\Delta\mathfrak{r}(v)$, $\Delta\mathfrak{g}(v)$, and $\Delta\mathfrak{b}(v)$. We assign a vertex u of $V(G)$ to $\Delta\mathfrak{r}(v)$ (resp. $\Delta\mathfrak{g}(v)$ or $\Delta\mathfrak{b}(v)$) if it lies in the interior of $\Delta\mathfrak{r}(v)$ (resp. $\Delta\mathfrak{g}(v)$ or $\Delta\mathfrak{b}(v)$). If $u \in \pi(v, \mathfrak{g})$ (resp. $u \in \pi(v, \mathfrak{b})$ or $u \in \pi(v, \mathfrak{r})$), then we assign it to $\Delta\mathfrak{r}(v)$ (resp. $\Delta\mathfrak{g}(v)$ or $\Delta\mathfrak{b}(v)$). Note, that this also assigns the three outer vertices. To avoid v from being abused, we do not assign it to any of these three groups. Let $n_\mathfrak{r}(v), n_\mathfrak{g}(v), n_\mathfrak{b}(v)$ be the number of vertices of $V(G) - v$ assigned to the $\Delta\mathfrak{r}(v)$, $\Delta\mathfrak{g}(v)$, $\Delta\mathfrak{b}(v)$, respectively. For $v \in V(G) \setminus \{\mathfrak{r}, \mathfrak{g}, \mathfrak{b}\}$, let

$$f(v) = \frac{1}{n-1}\Big(n_\mathfrak{r}(v), n_\mathfrak{g}(v), n_\mathfrak{b}(v)\Big)$$

be the embedding of the vertices of G into barycentric coordinates. The embedding for the three outer vertices is

$$f(\mathfrak{r}) = \frac{(n-2,1,0)}{n-1}, \quad f(\mathfrak{g}) = \frac{(0,n-2,1)}{n-1}, \quad \text{and} \quad f(\mathfrak{b}) = \frac{(1,0,n-2)}{n-1}. \tag{14.2}$$

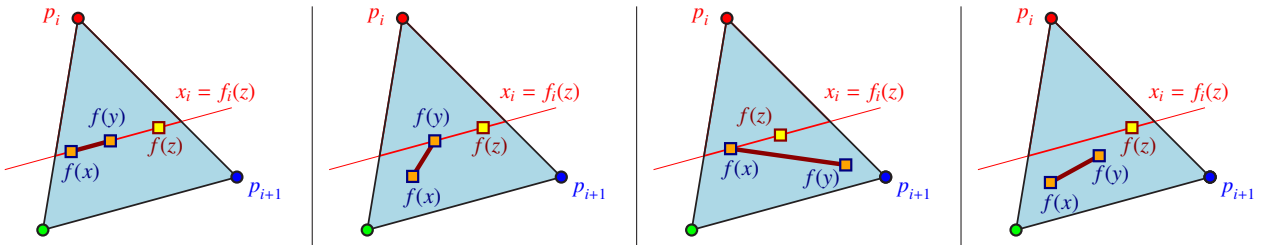

Figure 14.13: Case analysis for the proof of Lemma 14.3.6 for the case $f(xy) \prec_i f(z)$.

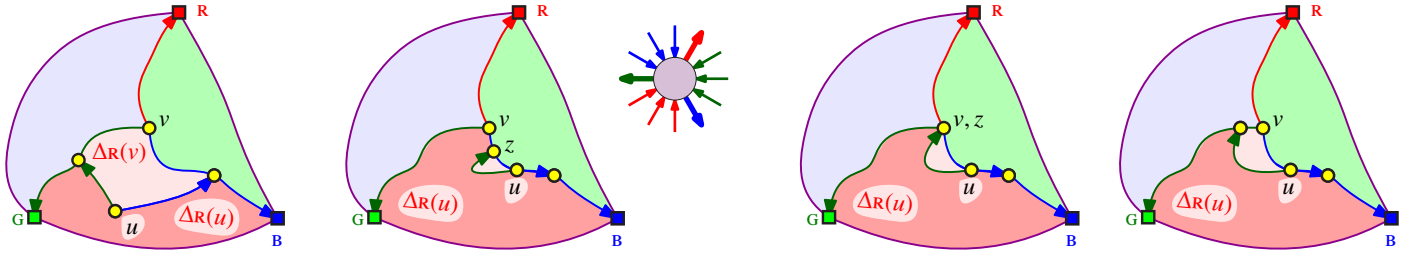

Figure 14.14

14.3.2.1. Embedding via *weak* barycentric representation

We somewhat weaken the demands on the barycentric representation, by using a variant of lexicographical ordering to allow equality in some coordinates.

Definition 14.3.4. For $f : V(G) \to \Delta_3$, let $f(u) \prec_i f(v)$ denote that $f(u) <_i f(v)$, or $f(u) =_i f(v)$ and $f(u) <_{i+1} f(v)$, where $<_4$ is $<_1$.

Definition 14.3.5. A ***weak barycentric representation*** of a graph G is an injective mapping $f : V(G) \to \Delta_3$, such that for each edge $xy \in E(G)$ and a vertex $z \in V(G) \setminus \{x, y\}$, there is an index i, such that $f(xy) \prec_i f(z)$.

Lemma 14.3.6. *Let f be a weak barycentric representation of G. For any three points p_1, p_2, p_3 in the plane that are not collinear, the mapping $g(v) = \sum_i p_i f_i(v)$ induces a straight-line embedding of G in the plane. Specifically, we map an edge $xy \in E(G)$ to the segment $g(x)g(y)$.*

Proof: For an edge xy and a vertex $z \neq x, y$, there is an index i, such that $f(z) \succ_i f(xy)$. This implies by an easy case analysis that $f(z)$ is not on the segment $f(x)f(y)$, see Figure 14.13. This readily implies that two adjacent edges do not intersect in their interior in the embedding.

Next, consider two edges that are disjoint $xy, uv \in E(G)$. By definition, there are indices $i, j, h, k \in \{1, 2, 3\}$, such that (i) $f(x) \succ_i f(uv)$, (ii) $f(y) \succ_j f(uv)$, (iii) $f(u) \succ_h f(xy)$, and (iv) $f(v) \succ_k f(xy)$. If $i = h$ then $f(x) \succ_i f(u) \succ_i f(x)$, which is impossible. Arguing similarly, we get that $\{i, j\} \cap \{h, k\} = \emptyset$. As, such $i = j$ or $h = k$ (since there are only three possible values for these four indices). If $i = j$ then $f(xy) \succ_i f(uv)$, which implies that the segments are disjoint, by a similar case analysis to the one above. Similarly, if $h = k$ then $f(uv) \succ_h f(xy)$, which implies disjointness again. ■

14.3.3. The realizer embedding is a *weak* barycentric representation

Lemma 14.3.7. *Let u, v be two distinct vertices in G. Then*

(i) *if* $u \in \Delta\mathfrak{r}(v)$ *then* $f(u) \prec_1 f(v)$,
(ii) *if* $u \in \Delta\mathfrak{g}(v)$ *then* $f(u) \prec_2 f(v)$, *and*
(iii) *if* $u \in \Delta\mathfrak{g}(v)$ *then* $f(u) \prec_3 f(v)$.

Proof: We prove only case (i) as the other cases are similar. If $v = \mathfrak{r}$ then the claim obviously holds. We remind the reader that $n_{\mathfrak{r}}(v) = |\Delta\mathfrak{r}(v)| - |\pi(v, \mathfrak{b})|$. In particular, if $u \in \Delta\mathfrak{r}(v) \setminus \pi(v, \mathfrak{b})$ then $\Delta\mathfrak{r}(u) \setminus \pi(u, \mathfrak{b}) \subsetneq \Delta\mathfrak{r}(v) \setminus \pi(v, \mathfrak{b})$ by Lemma 14.2.6. See Figure 14.14 (A). This in turn implies that $f(u) \prec_1 f(v)$, and we are done.

So the problematic case is when $u \in \pi(v, \mathfrak{b})$. Observe that $\pi(u, \mathfrak{g})$, except for u, can not contain a vertex z that lies on $\pi(v, \mathfrak{b})$ as this would be a color pattern violation. See Figure 14.14 (B) and Figure 14.11 (B). As a concrete example, the case that $z = v$ is impossible as the path from v to $\mathfrak{b}$, fails to be to the left of the green path $\pi(u, \mathfrak{g})$ going through v. This is depicted in Figure 14.14 (C).

This implies that $\Delta\mathfrak{r}(v) \setminus \pi(v, \mathfrak{b}) \subseteq \Delta\mathfrak{r}(u) \setminus \pi(u, \mathfrak{b})$, see Figure 14.14 (D). If the inclusion is proper then we are done as $f(u) \prec_1 f(v)$. As such, the remaining case is that $\Delta\mathfrak{r}(v) \setminus \pi(v, \mathfrak{b}) = \Delta\mathfrak{r}(u) \setminus \pi(u, \mathfrak{b})$ – that is $f(v) =_1 f(u)$.

Now, as $u \in \pi(v, \mathfrak{b})$, we have that $u \notin \pi(v, \mathfrak{r})$, since by construction the paths $\pi(v, \mathfrak{r})$, $\pi(v, \mathfrak{g})$ and $\pi(v, \mathfrak{b})$ are interior disjoint. As such, $u \in \Delta\mathfrak{g}(v) \setminus \pi(v, \mathfrak{r})$, and it follows that $f(u) \prec_2 f(v)$, by Lemma 14.2.6. This implies that $f(u) \prec_1 f(v)$. ■

Lemma 14.3.8. *The mapping defined in Section 14.3.2 is a weak barycentric representation of the triangulation* G.

Proof: The mapping f is injective by Lemma 14.3.7. So consider any edge $xy \in E(G)$ and a vertex $z \in V(G) \setminus \{x, y\}$.

If $x, y \in \Delta\mathfrak{r}(z)$ then $f(xy) \prec_1 f(z)$, by Lemma 14.3.7, and the weak barycentric conditions are met for f. The claim similarly follows if $x, y \in \Delta\mathfrak{g}(z)$ or $x, y \in \Delta\mathfrak{b}(z)$.

As such, the edge xy is not on the boundary, or in the interior, of one of the regions $\Delta\mathfrak{r}(z)$, $\Delta\mathfrak{g}(z)$, or $\Delta\mathfrak{b}(z)$. These three regions cover the interior of G (excluding the three outer edges).

So xy is one of the three outer edges of G. Say $xy = \mathfrak{r}\mathfrak{g}$. If $f_3(z) > 1$, then $f(z) \succ_3 f(xy)$ and we are done, as $f_3(\mathfrak{r}) = 0$ and $f_3(\mathfrak{g}) = 1$, see Eq. (14.2). Observe that for a vertex $u \in V(G)$, we have $f_3(u) = 0$ only for $u = \mathfrak{r}$. As such, the remaining case is that $f_3(z) = 1$. But then $f(z) \succ_3 f(\mathfrak{r}) = 0$. As for the other vertex $\mathfrak{g}$, we have $f(z) =_3 f(\mathfrak{g}) = 1$, but $f_1(\mathfrak{g}) = 0$, and $f_1(z) \geq 1$, which implies that $f(z) \succ_1 f(\mathfrak{g})$. This implies that $f(z) \succ_3 f(xy)$, as desired. ■

14.3.4. The result

The algorithms. Given a triangulation G with n vertices, one can compute its realizer in linear time. To this end, one maintains for each vertex three counters with the number of its common neighbors with $\mathfrak{r}$,$\mathfrak{g}$, and $\mathfrak{b}$, respectively. It is easy to verify that these counters can be maintained in linear time overall through a contraction sequence, and then it is straightforward to maintain a list of all vertices that currently have only two common neighbors with the outer vertices and are connected to them. As such, one can maintain the contractible edges, and do the contractions, thus getting the desired contraction/expansion sequence, which then can be turned into the realizer in a straightforward fashion.

Computing the barycentric mapping of Section 14.3.2 can be easily done in linear time per vertex, and quadratic time overall. If one works harder, one can get linear time - the details are somewhat tedious, and we leave them as an exercise to the interested reader.

Theorem 14.3.9. *Given a triangulated planar graph G with an embedding (i.e., a triangulation) with n vertices, one can compute in linear time a equivalent straight-linear embedding of G, where the vertices are integer points in the grid $\{0, 1, \ldots, n-1\}^2$.*

Proof: We compute the barycentric mapping f of Section 14.3.2 for G. Observe, that by construction we have $f_i(v) \in \left\{ \frac{i}{n-1} \;\middle|\; i = 0, \ldots, n-1 \right\}$ for $v \in V(G)$, and $i \in \{1, 2, 3\}$. By Lemma 14.3.8 f is a weak barycentric representation of G, and by Lemma 14.3.6 it yields a straight-line embedding in the plane. As such, picking $p_1 = (0, 0)$, $p_2 = (0, n-1)$, and $p_3 = (n-1, 0)$, the mapping $g(v) = \sum_i p_i f_i(v)$ is the desired straight-linear embedding. ■

14.4. Bibliographical notes

The canonical ordering of Section 14.1 and the grid embedding is from de Fraysseix et al. [**fpp-hdpgg-90**]. The linear time implementation is described by Chrobak and Payne [**cp-ltadp-95**].

The oriented trees decomposition is from the work by Walter Schnyder [**s-epgg-90**]. The resulting embedding algorithm on the grid, described in Section 14.3, is somewhat more complicated than the previous algorithm, but the realizer idea seems useful and worthy of presentation. This decomposition of planar graphs into three trees is sometime referred to as *Schnyder Woods*. Schnyder work was motivated by his work on characterizing planar graphs in terms of the dimension of its poset, showing that this dimension is at most three [**s-pgpd-89**].

14.5. Exercises

14.6. From previous lectures