MPRI Course 2-38-1: Algorithms and Combinatorics for Geometric Graphs
Exercise Sheet 1

1 Exercise 1

1. A simple bipartite cellularly embedded planar graph is bibipartite if its dual graph is simple and also bipartite. Give a complete list of all bibipartite planar graphs and prove that it is complete. *Hint: it is non-empty!*.

Solution: In a bipartite graph, faces have even degree. So in a bibipartite graph, vertices have even degree. Any vertex of degree two yields a multiple dual edge between the two adjacent faces, contradicting bibipartiteness. If there is a vertex of degree 0, it is unique because the graph is cellularly embedded. Otherwise, all the vertices have degree at least four. Dually, all the faces have degree at least four. As we have seen in class, this is not possible due to the Euler formula.

2. Let G be a simple planar graph, and suppose we arbitrarily color each edge of G either blue or red. Prove that for any embedding of G in the plane, there exists a vertex around which the incident red edges are consecutive.

Solution: Let us assume otherwise, then for each red edge incident to a vertex v, there exists another red edge incident to v so that there are blue edges on both sides of the path formed by these two edges. So we can find a red cycle C with the property that for every vertex on C except maybe one, there is at least one blue edge incident to that vertex that is inside the disk D bounded by C (which exists by the Jordan-Schoenflies theorem). The vertex where this might not happen is the one where we close the cycle, where we have no control, we call it the *closing vertex*. We call such a cycle a *balanced* red cycle. Let us pick such a balanced red cycle C so that the disk D it bounds is minimal with respect to containment: there is no other balanced red cycle within D.

We claim that this implies that there is no red edge in the interior of D, as otherwise we could follow it and either find a smaller disk if it closes into a disk, or find a red path within the red disk. If this red path closes itself, we have a smaller balanced cycle. Otherwise, one of the two endpoints of this red path is incident to a vertex that is not a closing vertex, and then it bounds a balanced cycle with one of the two red paths in C, as the four possible cases in Figure 1 picture, contradicting minimality. So all the edges within D are blue. If there is a vertex within D, all its incident edges are blue, a contradiction. Otherwise, we get a polygon with some edges inside, but any such polygon has at least two vertices of degree two. This contradicts the fact that in a balanced cycle, all the vertices except maybe one have an incident blue edge within the disk D, and finishes the proof.

3. Find universal constants α, β and γ (not depending on n or g) such that the following holds: For all integers n and g such that $n \geq \gamma g$, every simple n-vertex graph embedded on a surface of genus g has an independent set[1] of size n/α, in which every vertex has degree at most β.

[1] An independent set in a graph G is a subset of the vertices of G, no two of which are connected by an edge in G.

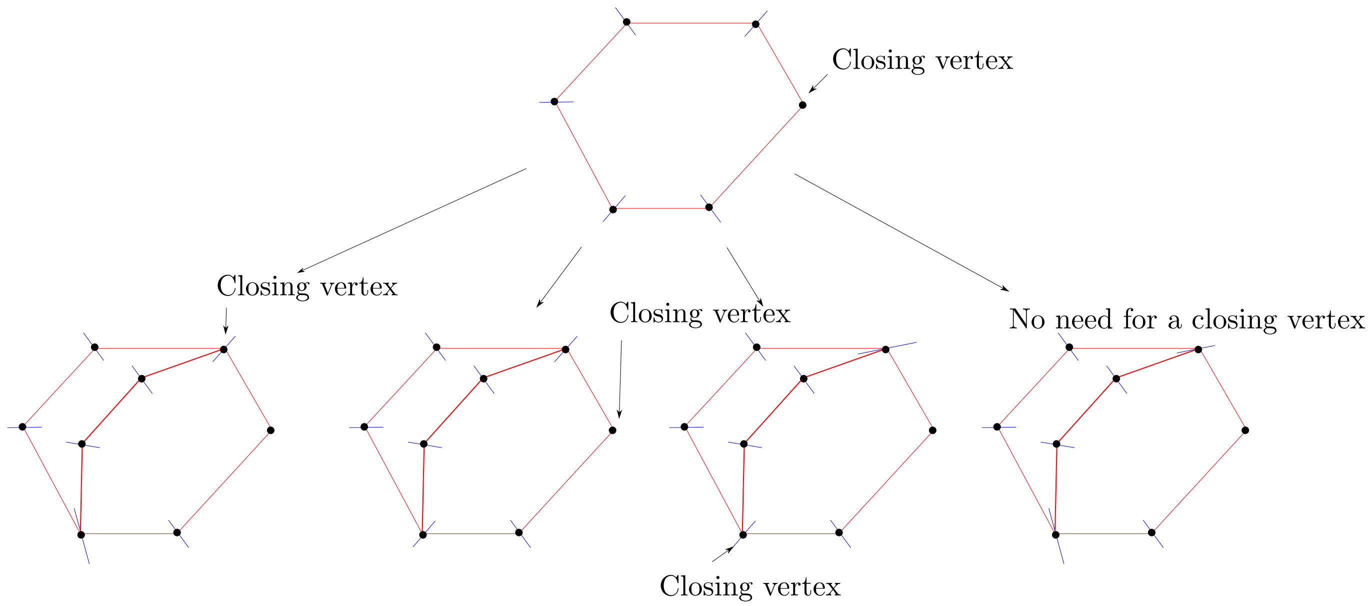

Figure 1: In cases 1, 3 and 4 the left cycle is balanced, and in case 2 the right cycle is balanced.

Solution: Let us pick $\beta = 24$, $\gamma = 2$ and $\alpha = 49$, which have not been optimized at all. Let us denote by n, e and f the number of vertices, edges and faces of a graph embedded on a surface of genus g. Without loss of generality, by reducing g if necessary, we can assume that the graph is cellularly embedded. Then the Euler formula gives $v - e + f = 2 - 2g$. We claim that for $n \geq \gamma g$, there are at most $n/2$ vertices of degree higher than β. Otherwise we would have $\beta n/2 \leq 2e$, and since the graph is simple, faces have degree at least 3, so $3f \leq 2e$. So we get $2 - 2g \leq n - e/3 \leq n(1 - \beta/12)$ which is a contradiction for our choice of β and γ.

So if we remove all the vertices of degree higher than β, we still have at least $n/2$ vertices. Now we can pick any of those, remove it and all its neighbors, and do it again. This will give us an independent set of size $\frac{n}{2(\beta+1)}$.

4. Describe an algorithm to find such an independent set in $O(n)$ time.

Solution: We first search through the graph in linear time to remove all vertices of degree bigger than β. Then we pick any vertex, and remove it and its neighbors, and we induct. Each of these removal steps takes constant time (because all the vertices have constant degree), so the whole procedure takes linear time.

2 Exercise 2

A cycle C on a graph G is *nonseparating* if $G \setminus C$ is connected.

1. Prove that any n-vertex triangulation of an orientable surface S of positive genus contains a non-separating cycle C of length at most $2\sqrt{n}$. *Hint: cut S along C, yielding two copies C_1 and C_2 of C on the boundary. How many independent paths are there from C_1 to C_2 and how long are they?*

Solution: This solution assumes that the reader is familiar with the setup for non-separating cycles that is introduced at the end of the lecture notes, but that we skipped in class. One can argue without it but it makes the arguments cleaner.

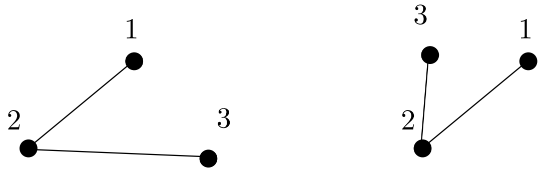

Figure 2: In a single morphing step, the vertex 3 would cross the edge (12) that does not move.

Any orientable surface has a non-separating simple cycle (this is proved at the end of the lecture notes), and we pick a minimal one C, and cut along it. By Menger's theorem, the number of independent paths from C_1 to C_2 equals the size of the smallest vertex cut S separating C_1 from C_2. The induced subgraph $G[S]$ contains a cycle that is homologous to C and thus non-separating. Therefore it must have size at least that of C. Each of these independent paths must have length at least $|C|/2$, as otherwise one would find another non-separating cycle by taking it and taking the smaller of the subpaths of C, yielding a shorter non-separating cycle than C. Therefore there are at least $|C| \times |C|/2$ vertices in G, and thus there is a non-separating cycle of length as most $\sqrt{2n}$. (Yes there was a typo in the exercise and the 2 should have been under the square root (but it was still correct.))

2. Deduce that any n-vertex graph on an orientable surface of genus g has a 2/3-separator S of size $O(g\sqrt{n})$, and such that each component of $G \setminus S$ is planar.

Solution: Each time we remove a nonseparating cycle of the graph and all the adjacent edges, the graph can be embedded on a surface of lower genus. So after applying g times the inequality in the first question, we have removed $O(g\sqrt{n})$ vertices and the remaining graph is planar. We can now find a separator for planar graphs of size $O(\sqrt{n}$. Taking the union of the planar separator and the removed vertices yields the surface separator.

3 Exercise 3

Let G be a planar graph, and let G_1 and G_2 be two isomorphic[2] straight-line embeddings of G, where each face, including the outer face, is a triangle. A *morphing step* between G_1 and G_2 is a straight-line continuous transformation of one into the other, such that the graph stays planar at all times: for each vertex v of G, we denote by $S(v)$ the segment connecting v_1, the embedding of v in G_1 to v_2, the embedding of v in G_2, and we slide v from v_1 to v_2 at uniform speed along this segment. At a time $t \in [1,2]$, we denote by v_t the position of v, and for any edge (uv) in E, we connect v_t to u_t with a straight segment. This defines a family of drawings $(G_t)_{t\in[1,2]}$, and this is a morphing step if all these drawings are planar embeddings. A *morphing* from G_1 to G_2 is a sequence of morphing steps $G_1 \to G' \to G^{(2)} \ldots \to G^{(k)} = G_2$, where the graphs $G^{(i)}$ are all straight-line embeddings of G. The integer k is the *complexity* of the morphing.

1. Provide an example of a planar graph G and two straight-line embeddings (not necessarily triangulated) that are not connected by a single morphing step.

Solution: See Figure 2

The rest of the exercise aims at proving that for any two straight-line embeddings G_1 and G_2 with the above conditions, there always exists a morphing of finite complexity between G_1 and G_2. The proof is by induction.

[2]This means here that G_1 and G_2 have the same outer face and the same combinatorics *as an embedded graph*: same set of facial walks when turning clockwise.

2. Prove the base case of induction for $n = 4$.

Solution: Even for $n = 4$, one can not in general do it in a single step: think about two K_4, one of which has been rotated by a 180 compared to the first one. But one can first move everything so that the middle vertex coincide, then rotate, and then a single morphing step will work.

The *visibility kernel* of a polygon is the set of points inside or on the polygon that can be "seen" from any vertex of the polygon, i.e., the set of points p such that for any vertex v of the polygon the segment pv does not cross the polygon.

3. Prove that for any polygon with at most 5 vertices, one of the vertices is contained in its visibility kernel.

Solution: For convex polygons, and thus for triangles, the visibility kernel is the whole polygon. For a concave quadrilateral, the concave vertex is in the visibility kernel. For pentagons, either the two concave vertices are consecutive, or not. In the first case, the vertex opposite to the consecutive vertices sees everything. In the second case, the vertex inbetween the consecutive vertices sees everything.

The *link* $L(v)$ of a vertex v of G_1 or G_2 is the polygon defined by the neighbors of v.

4. Prove that there exists a vertex v of G so that both in G_1 and in G_2, the link $L(v)$ contains a vertex u that is in the visibility kernel of $L(v)$. (Note that u might be different in G_1 and G_2.

Solution: Since G is planar, at least one vertex has degree at most 5. Its link in G_1 is a polygon, which therefore has a vertex contained in the visibility kernel of the link by the previous question. Likewise in G_2.

We first assume that there are no edges in G connecting non-adjacent vertices of $L(v)$.

5. (*) Prove that there exists a straight-line embedding G' of $G \setminus v$ so that $L(v)$ is convex.

Solution: One can do it by hand (see reference at the end of the document). But we take a simpler way and prove that G' is 3-connected, and the result will follow by using Tutte's theorem on 3-connected graphs. Note that G' is triangulated except for possibly the face $L(v)$, and $L(v)$ is a cycle and in this face there are not edges connecting non-adjacent vertices of $L(v)$. Let us call such a graph *almost triangulated.* In an almost-triangulated graph, if e is an edge belonging to exactly two triangles, or to exactly one triangle and the face $L(v)$, then contracting that edge e yields another almost triangulated graph (note that this is not true without the assumption). Let us assume that there are two vertices u and w disconnecting G', into at least two components X and Y. We claim that there is always an edge different from uw that belongs to two triangles or to exactly one triangle and the face $L(v)$: pick any edge, and if it belongs to too many triangles, one of them is separating. Then we pick another edge inside that separating triangle. The separating triangles between nested, this process stops and we have found our edge, proving the claim. Now, we repeatdely contract one edge different from uw that belongs to two triangles or to exactly one triangle and the face $L(v)$, keeping at least one vertex in X and in Y. At this stage, there are exactly four vertices in the graph: u, w and one from X and one from Y. The graph is still almost-triangulated, thus and if u and w are on the face $L(v)$, and it has degree more than 3 (thus degree four), then u and w are adjacent. Therefore, the graph is either K_4, or K_4 minus an edge with u and w adjacent. Furthermore, u and w disconnect that graph, which is not the case. Hence we have a contradiction, finishing the proof.

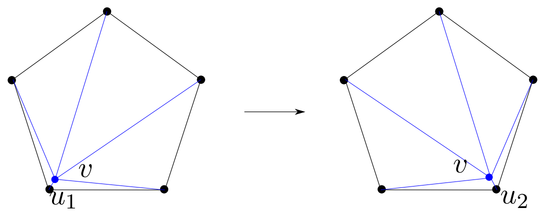

Figure 3: Moving the vertex v between u_1 and u_2.

6. What is the visibility kernel of $L(v)$ in G'? Assuming the induction hypothesis (every two straight-line triangulations with $n-1$ vertices can be morphed one into the other), prove that one can morph G_1 into G_2. *Hint: contract an edge, and use the induction to morph into G'.*

Solution: Since $L(v)$ is convex, its visibility kernel is $L(v)$. By Question 4., there exists a vertex v so that there is a vertex u_1 is in the visibility kernel of $L(v)$ in G_1, and a vertex u_2 in the visibility kernel of $L(v)$ in G_2. In G_1, we send the vertex v to u_1 using a single morphing step. Actually, we send it very close to u_1, with the idea that any move of u_1 will tell us how to move v. This is a tad painful to justify accurately (the keyword is pseudomorphs, you can look it up online). Likewise, in G_2 we send v to u_2 using a single morphing step. Now G_1 and G_2 have become G_1' and G_2' which are supergraphs of G' (they have more edges in the face $L(v)$). They both have one less vertex than G. By the induction hypothesis and question 5., G_1' can be morphed into a straight-line embedding of G' where the face $L(v)$ is convex. Likewise, G_2' can be morphed into the same (except for the edge within $L(v)$) straight-line embedding. Remembering that we put v just next to u_1 (respectily u_2), now there just remains to move the vertex v between those two embeddings, which can be done using a single morphing step since $L(v)$ is convex, see Figure 3.

We now remove the additional assumption.

7. (*) Prove the induction step in the general case. *Hint: without the assumption, there is no hope of finding a straight-line embedding where $L(v)$ is convex, but we can still find an embedding G' where all the vertices of $L(v)$ except the non-adjacent ones which are joined by an edge of G are in the visibility kernel of $L(v)$.*

Solution: Since nobody even attempted this question, I will just refer to the original article that inspired this exercise, and is quite readable: Cairns, S.: Deformations of plane rectilinear complexes. The American Mathematical Monthly 51(5), 247252 (1944).