Chapter 16

Planar separator

By Sariel Har-Peled, March 30, 2022[①] Version: 0.1

> "You see, dogs aren't enough any more. People feel so damned lonely, they need company, they need something bigger, stronger, to lean on, something that can really stand up to it all. Dogs aren't enough, what we need is elephants..."
>
> The roots of heaven, Romain Gary

The ***planar separator theorem*** is a fundamental result about planar graphs [Ung51, LT79]. Informally, it states that one can remove $O\left(\sqrt{n}\right)$ vertices from a planar graph with n vertices and break it into "significantly" smaller parts. It is widely used in algorithms to facilitate efficient divide and conquer schemes on planar graphs.

16.1. Planar separator from the circle packing theorem

Given a planar graph $G = (V, E)$, it can be drawn in the plane as a kissing graph, see Theorem 16.8.2. It turns out that the planar separator theorem is an easy consequence of this deep result.

Let $\mathcal{D}$ be the set of disks realizing G as a kissing graph, and let P be the set of centers of these disks. Let d be the *smallest* radius disk containing $n/8$ of the points of P, where $n = |P| = |V|$ (this disk is unique assuming general position assumption on P). To simplify the exposition, we assume that d is of radius 1 and it is centered in the origin. Randomly pick a number $x \in [1, 2]$ and consider the circle C_x of radius x centered at the origin. Let S be the set of all disks in $\mathcal{D}$ that intersect C_x. We claim that, in expectation, S is a good separator.

Lemma 16.1.1. *The separator S breaks G into two subgraphs with at most $(7/8)n$ vertices in each connected component.*

Proof: The circle C_x breaks the graph into two components: (i) the disks with centers inside C_x, and (ii) the disks with centers outside C_x.

Clearly, the corresponding vertices in G are disconnected once we remove S. Furthermore, a disk of radius 2 can be covered by 7 disks of radius 1, as depicted in Figure 16.1. As such, the disk of radius 2 at the origin can contain at most $7n/8$ points of P inside it, as a disk of radius 1 can contain at most $n/8$ points of P. We conclude that there are at least $n/8$ disks of $\mathcal{D}$ with their centers outside C_x, and, by construction, there are at least $n/8$ disks of $\mathcal{D}$ with centers inside C_x. As such, once S is removed, no connected component of the graph $G \setminus S$ can be of size larger than $(7/8)n$. ∎

Lemma 16.1.2. *We have $\mathbb{E}\left[|S|\right] \leq 11\sqrt{n}$, where $n = |V|$.*

Proof: Let $\ell < 1$ be a parameter to be specified shortly. We split $\mathcal{D}$ into two sets: $\mathcal{D}_{\leq \ell}$ and $\mathcal{D}_{> \ell}$ of all disks of diameter $\leq \ell$ and $> \ell$, respectively.

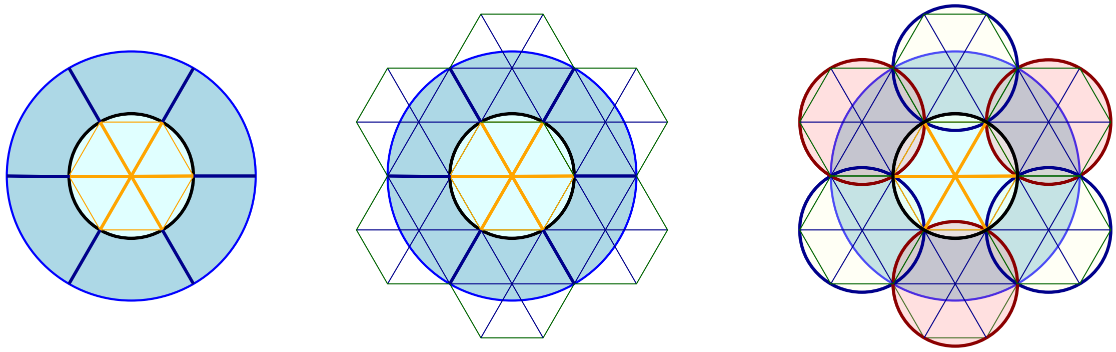

(A) (B) (C)

Figure 16.1: How to cover a disk of radius 2 by 7 disks of radius 1.

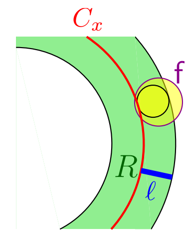

Consider the ring $R = \mathrm{disk}(0, x + \ell) \setminus \mathrm{disk}(0, x - \ell)$, and observe that any disk f of $\mathcal{D}_{>\ell}$ that intersects C_x, must contain inside it a disk of radius $\ell/2$ that is fully contained in R. As such, f covers an area of size at least $\alpha = \pi(\ell/2)^2$ of this ring. The area of R is $\beta = \pi\Big((x + \ell)^2 - (x - \ell)^2\Big) = 4\pi x \ell$. As such, the number of disks of $\mathcal{D}_{>\ell}$ that intersect C_x is $\le \beta/\alpha = 4\pi x\ell/(\pi\ell^2/4) = 16x/\ell$. As $\mathbb{E}[x] = 3/2$, we have $\mathbb{E}[\beta/\alpha] = 24/\ell$.

Consider a disk $\mathsf{u}_i \in \mathcal{D}_{\le\ell}$ of radius $r_i \le \ell/2$ centered at p_i. The circle C_x intersects u_i if and only if $x \in \Big[\|p_i\| - r_i, \|p_i\| + r_i\Big]$, and as x is being picked uniformly from $[1, 2]$, the probability for that is at most $2r_i/|2 - 1| = 2r_i \le \ell$. As such, since $|\mathcal{D}_{\le\ell}| \le n$, we have that the expected number of disks of $\mathcal{D}_{\le\ell}$ that intersect C_x is at most $n\ell$. Adding the two quantities together, we have that the expected number of disks intersecting C_x is bounded by $n\ell + 24/\ell$, which is $\le 2\sqrt{24n}$, for $\ell = 1/\sqrt{24n}$. ■

Now, putting Lemma 16.1.1 and Lemma 16.1.2 together implies the following.

Theorem 16.1.3. *Let $G = (V, E)$ be a planar graph with n vertices. There exists a set S of $11\sqrt{n}$ vertices of G, such that removing S from G breaks it into several connected components, each one of them contains at most $(7/8)n$ vertices.*

16.2. Planar separator in linear time

The circle packing theorem is non-constructive and algorithmically one can only approximate the circle-packing it defines. Fortunately, one can get the planar separator via a direct algorithm (if with somewhat more work). In particular, we present here a linear time algorithm for computing the planar separator.

The input is a planar graph G with n vertices together with its embedding.

Sketch of algorithm. The algorithm starts by computing a BFS tree of G. If there is any (light) BFS layer that is of size $O(\sqrt{n})$ and separates the graph, the algorithm just returns it as the separator. Otherwise, it must be that all the good separating layers are "heavy". Fortunately, all these heavy layers must be sandwiched between two light layers that are at most $O(\sqrt{n})$ apart. The idea is to compute a separator for the induced subgraph on these middle heavy layers. This new graph has only $O(\sqrt{n})$ layers in its BFS tree $\mathcal{T}$, and in particular, one can interpret this tree as a boundary of a polygon – triangulate this polygon, and find a separating edge e that breaks this polygon in a balanced way as far as the number of triangles. Adding e to the BFS tree $\mathcal{T}$ creates a cycle C of length $O(\sqrt{n})$, and together with the top/bottom light layers of the sandwich, it forms the desired separator in the original graph.

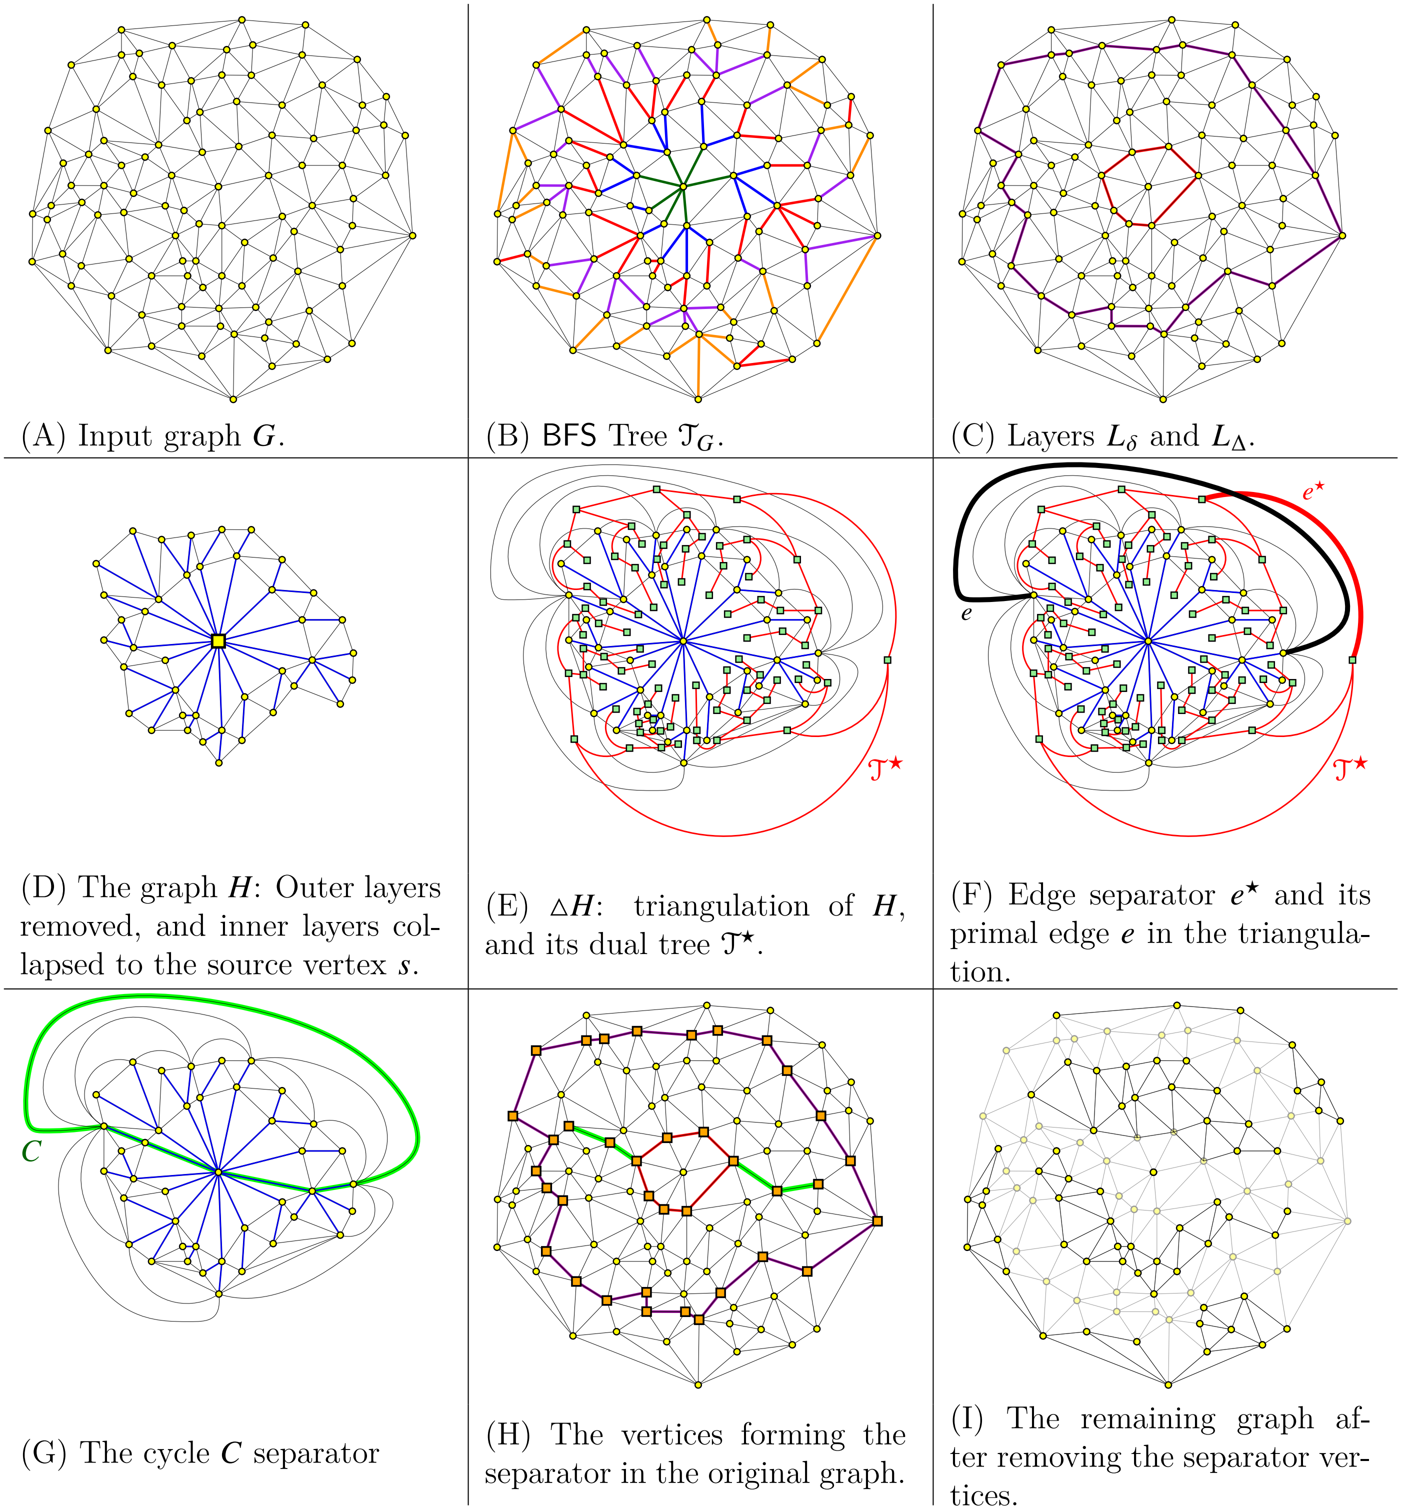

(A) Input graph G.

(B) BFS Tree $\mathcal{T}_G$.

(C) Layers L_δ and L_Δ.

(D) The graph H: Outer layers removed, and inner layers collapsed to the source vertex s.

(E) $\triangle H$: triangulation of H, and its dual tree $\mathcal{T}^\star$.

(F) Edge separator $e^\star$ and its primal edge e in the triangulation.

(G) The cycle C separator

(H) The vertices forming the separator in the original graph.

(I) The remaining graph after removing the separator vertices.

Figure 16.2: An illustration of the algorithm for computing the separator in a graph.

16.2.1. The planar separator algorithm

The following algorithm is illustrated in Figure 16.2. The algorithm works as follows:

(A) Pick an arbitrary vertex $s \in V(G)$, and do a BFS in G starting at s, and let $L_0, L_1, L_2, \ldots, L_t$ be the layers of the BFS tree $\mathcal{T}_G$, with $L_0 = \{s\}$. Let $L_{t+1} = \{\}$ be a fake "top" layer.

(B) Let k, K be the max/min indices, respectively, such that $\sum_{i=0}^{k-1} |L_i| < (1/3)n$ and $\sum_{i=K+1}^{t} |L_i| < (1/3)n$.

(C) If for any i, such that $k < i < K$, we have that $|L_i| \leq 2\sqrt{n}$, then return L_i as the separator.

(D) δ: maximum index such that $\delta < k$ and $|L_\delta| \leq 2\sqrt{n}$.

Δ: minimum index such that $\Delta > K$ and $|L_\Delta| \leq 2\sqrt{n}$.

(E) $U = \bigcup_{i=\delta+1}^{\Delta-1} L_i$: Set of middle layer vertices.

(F) H: Graph resulting from deleting all the top layers $L_\Delta, L_{\Delta+1}, \ldots, L_{t+1}$ from G, and collapsing all the bottom layers $B = L_0 \cup L_1 \cup \cdots L_\delta$ into the vertex s.

(G) $\triangle H$: graph resulting from triangulating H.

(H) $\mathcal{T}$: BFS tree of $\triangle H$ rooted at s.

(I) $(\triangle H)^\star$: Dual graph to Q.

(J) $\mathcal{T}^\star$: Spanning tree of $(\triangle H)^\star$ where two vertices are connected by an edge if they are adjacent in $(\triangle H)^\star$, and the edge separating the respective triangles in the primal is not in $\mathcal{T}$. Formally, the edges of $\mathcal{T}^\star$ are $\bigl(E(\triangle H) \setminus E(\mathcal{T})\bigr)^\star$.

(K) $e^\star$: Edge separator of $\mathcal{T}^\star$ computed using the algorithm of Lemma 16.2.3.

Removing the edge $e^\star$ breaks $\mathcal{T}^\star$ into two trees with at most $2/3$ fraction of the edges.

(L) $e = \bigl(e^\star\bigr)^\star = xy$: An edge of $\triangle H$ that is not in $\mathcal{T}$.

(M) v: LCA in $\mathcal{T}$ of x and y. Let π_x and π_y be the paths from x and y to v in $\mathcal{T}$, respectively.

(N) Return $\bigl(V\bigl(\pi_x \cup \pi_y\bigr) - s\bigr) \cup L_\delta \cup L_\Delta$ as the separator.

16.2.1.1. More details

Step (F): The graph H is planar since we took a connected components of G and collapsed it into a vertex (which preserve planarity, of course), and removed vertices/edges (which also preserve planarity). It particular, with some care, one can compute the embedding of H from the given embedding of G in linear time.

Step (G): Since we have the embedding of H, it is straightforward to triangulate it in linear time to get $\triangle H$.

Step (I): The graph $(\triangle H)^\star$ is the dual to a triangulation and it is as such 3-regular. This implies that $\mathcal{T}^\star$ has maximum degree 3, and edge separator with each part being at most $2/3$ fraction of the tree, see Lemma 16.2.3.

Step (J): A spanning tree of a planar graph has only one outer face – which implies that once we "cut" the plane along the edges of $\mathcal{T}$, what remains is a polygon. The tree $\mathcal{T}^\star$ is the spanning tree of the dual of the triangulation of this polygon. Thinking about $\mathcal{T}$ as a boundary of a polygon is somewhat bizarre, because it bounds the polygon from the inside (instead of the outside), but inside/outside is all the same for planar graphs.

It is now easy to verify that the above algorithm has linear running time.

16.2.1.2. Correctness

Separator size. If a separator is returned by step (C) then it is a balanced separator of size $\sqrt{n}$, and we are done.

Otherwise, if $K - k + 1 > \sqrt{n}/2$, then $\left|\bigcup_{i=k}^{K} L_i\right| > 2\sqrt{n}\sqrt{n}/2 = n$, which is impossible. As such $|K - k| < \sqrt{n}/2$. The same analysis implies that $|\Delta - \delta| < \sqrt{n}/2$. This implies that depth of the BFS tree $\mathcal{T}$ is $\leq \sqrt{n}/2$, which in turn implies that $|C| \leq 1 + \sqrt{n}$. The final separator size is $|C| + |L_\delta| + |L_\Delta| \leq 5\sqrt{n} + 1$.

Balanced separation. We have to consider only the scenario where the algorithm did not stop at step (C). So, let $N \leq n$ be the number of vertices of $\triangle H$, and let φ be the number of faces in $\triangle H$. We have $\varphi = 2N - 4$, by Lemma 16.8.1. Since e is an $[1/3, 2/3]$ edge separator, it follows that $\varphi_{\text{in}}, \varphi_{\text{out}} \leq (2/3)\varphi$, where φ_{in} and φ_{out} are the number of triangles inside and outside C in $\triangle H$, respectively.

Let n_{in} and n_{out} be the number of vertices of $\triangle H$ inside and outside C, respectively. By Lemma 16.2.4, we have that

$$n_{\text{in}} = \frac{\varphi_{\text{in}} - |C|}{2} + 1 \leq \frac{2\varphi}{3 \cdot 2} = \frac{2N - 4}{3} \leq \frac{2}{3}N \leq \frac{2}{3}n.$$

The same argument applies to n_{out}. We conclude that after removing the separator of step (N), every connected component of the graph has at most $(2/3)n$ vertices.

16.2.1.3. The result

Theorem 16.2.1. *Given a planar graph G in the plane with n vertices together with its embedding in the plane, one can compute, in linear time, a set $S \subseteq V(G)$ of size $\leq 5\sqrt{n} + 1$, such that each of the connected components of $G - S$ contains at most $(2/3)n$ vertices.*

Remark 16.2.2. One can derive a similar algorithm for the case that the vertices and edges have weight. There, the separation breaks the graph into components, each of weight at most $(2/3)W$, where W is the total weight. While the resulting separator has only $O(\sqrt{n})$ vertices also in this case, its weight might be arbitrarily large.

16.2.2. Some helper lemmas used above

We used the following two easy lemmas – we provide the proofs for the sake of completeness.

Lemma 16.2.3. *Let $\mathcal{T}$ be a tree with n vertices, with maximum degree $d \geq 2$. Then, there exists an edge whose removal break $\mathcal{T}$ into two trees, each with at most $\lceil (1 - 1/d)n \rceil$ vertices. This edge can be computed in linear time.*

Proof: Let v_1 be an arbitrary vertex of $\mathcal{T}$, and root $\mathcal{T}$ at v_1. For a vertex v of $\mathcal{T}$ let $n(v)$ denote the number of nodes in its subtree – this quantity can be precomputed, in linear time, for all the vertices in the tree using DFS.

In the ith step, v_{i+1} be the child of v_i with maximum number of vertices in its subtree. If $n(v_{i+1}) \leq \lceil (1 - 1/d)n \rceil$, then the algorithm outputs the edge xy as the desired edge separator, where $x = v_i$ and $y = v_{i+1}$. Otherwise, the algorithm continues the walk down to v_{i+1}. Since the tree is finite, the algorithm stops and output an edge.

Assume, for the sake of contradiction, that $n(y) < n/d$. But then, x has at most $d(x) - 1 \leq d - 1$ children (in the rooted tree), each one of them has at most $n(y)$ nodes (since y was the "heaviest" child). As such, we have $n(x) \leq 1 + (d-1)n(y) < 1 + (d-1)n/d \leq \lceil (1 - 1/d)n \rceil$ if d does not divides n. If d divides n then $n(x) \leq 1 + (d-1)n(y) \leq 1 + (d-1)(n/d - 1) = ((d-1)/d)n + 2 - d \leq \lceil (1 - 1/d)n \rceil$.

Namely, the algorithm would have stopped at x, and not continue to y, a contradiction.

As such, $n/d \leq n(y) \leq \lceil (1 - 1/d)n \rceil$. But this implies that xy is the desired edge separator. ∎

Lemma 16.2.4. *Let G be a planar graph with an outer face f, such that the boundary of f is a cycle C and no edge appears twice in C. Assume that G has φ inner faces, and furthermore all these faces are triangles. Then G has $(\varphi - |C|)/2 + 1$ internal vertices (i.e., vertices that do not appear in C).*

Proof: Let n be the number of vertices of G. Add an a vertex v to G in the outer face of G, and connect it to all the vertices lying on the boundary of the outer face of G. The resulting graph is a triangulation with $n + 1$ vertices, and $2(n + 1) - 4 = 2n - 2$ triangles, by Lemma 16.8.1. This counts $|C|$ triangles that were created by the addition of v. As such, $\varphi = 2n - 2 - |C| \implies n = \varphi/2 + 1 + |C|/2$. The number of inner vertices is $n - |C| = (\varphi - |C|)/2 + 1$. ■

16.3. Extensions

Here, we show various extensions of the above existential proof of the planar separator theorem, The reader might want to only skim this part on first reading.

16.3.1. Weighted version

Lemma 16.3.1. *Let $G = (V, E)$ be a planar graph with n vertices, and assume that the vertices have non-negative weights assigned to them, with total weight W. There exists a set S of $4\sqrt{n}$ vertices of G, such that removing S from G breaks it into several connected components, each one of them contains a set of vertices of total weight at most $(9/10)W$.*

Proof: The proof of Theorem 16.1.3 goes through, with the minor modification that that d is picked to be the smallest disk, such that the total weight of the centers of the disks it covers is $\geq W/10$. ■

Note, that if there is a vertex in the graph with weight $\geq W/10$, then the returned separator could be this single vertex, which is a legal answer (as the weight of the remaining graph is sufficiently small).

16.3.2. Cycle separators

A planar graph G is ***maximal*** if one can not add edges to it without violating its planarity. Any drawing of a maximal planar graph is a triangulation; that is, every face is a triangle. But then, in the realization of the graph as a kissing graph of disks, a face of the complement of the union of the disks has three touching disks as its boundary.

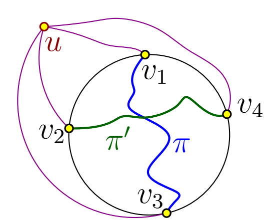

In particular, consider the separating cycle C_k, and two disks f and f' that intersect it consecutively along C_x. Let I be interval on C_x between $\mathsf{f} \cap C_x$ and $\mathsf{f}' \cap C_x$. The interval I belong to a single face of the complement of the union of disks, and in particular, this face has both f and f' on its boundary. As such, the vertices of G that corresponds to f and f' are connected by an edge. That is, the resulting separator is a cycle in G. Since C_x intersects a disk along an interval (or not at all), it follows that this cycle is simple. Thus, we get the following.

Theorem 16.3.2 ([Mil86])**.** *Let $G = (V, E)$ be a maximal planar graph with n vertices. There exists a set S of $4\sqrt{n}$ vertices of G, such that removing S from G breaks it into several connected components, each one of them contains at most $(9/10)n$ vertices. Furthermore S is a simple cycle in G.*

Cycle separator if the graph is not triangulated.

Lemma 16.3.3 ([Mil86]). *Let $G = (V, E)$ be a connected planar graph with n vertices, where the ith face has d_i vertices on its boundary, and let $N = \sum_i d_i^2$. Then, there exists a set S of $4\sqrt{N}$ vertices of G, such that removing S from G breaks it into several connected components, each one of them contains at most $(9/10)n$ vertices. Furthermore S is a cycle in G.*

In particular, if the maximum face degree in G is d, then the separator size is $O\left(\sqrt{nd}\right)$.

Proof: The idea to fill in the faces of G so that they are all triangulated.

So, consider a cycle C (not necessarily simple – an edge might be traversed twice) with k vertices that forms the boundary of a single face in the given embedding of G. Next, we build a graph having $C_1 = C$ as its outer boundary, as follows – it has k copies of C one inside the other, where the ith copy C_i is connected to the $i-1$ and $i+1$ copies, in the natural way, where a vertex is connect to its copies. Drawn in the plane, this results in a grid like construction. We also triangulate the inner most copy C_k in an arbitrary fashion, and every quadrilateral face is triangulated in an arbitrary fashion. The resulting graph G_C has k^2 vertices, and has the property that the any path between any two vertices of C in G_C, the corresponding shortest path in C is shorter (or of the same length). See Figure 16.3 for an example.

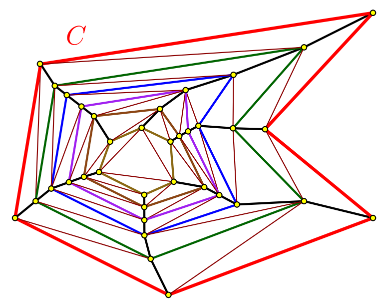

Figure 16.3

We repeat this fill-in process for all the faces of G, and let G' be the resulting graph. G' is still planar, and clearly the number of resulting vertices in the new graph is $N = \sum_i d_i^2$. Observe that $\sum_i d_i \leq 6n$, as every vertex v incident on a face r, can be charged to an edge adjacent to both v and f. Clearly, if done in a consistent fashion, an edge would be charged at most twice, and the maximum number of edges in a planar graph is $3n - 6$ by Euler's formula.

In particular, if the maximum value of d_i is d, then maximum of $N = \sum_i d_i^2$ is $O(nd)$, as can be easily verified.

Now, we assign weight zero to all the newly introduced vertices in G', and assign weight one for the original vertices (that appear in G). The graph G' is a fully triangulated planar graph and it has N vertices. By Lemma 16.3.1, there is separator providing the desired partition, and the number of vertices on this separator is $\leq 4\sqrt{N}$. Since G' is triangulated, the separator is a simple cycle in G'. We now replace portions of it that uses the face grids by the appropriate paths along the original boundary of the faces. Clearly, the resulting cycle in G has the same number of vertices, provide the same quality of separation (or better, since some vertices migrated to the separator), as desired. ■

Miller's result is somewhat stronger than Lemma 16.3.3, as he assumes the graph is 2-connected, and can ensure that in this case the separator is a *simple* cycle.

16.3.3. Ball systems that are k-ply

A set of balls $\mathcal{B}$ in $\mathbb{R}^d$ is ***k-ply***, if no point of $\mathbb{R}^d$ is contained in more than k balls of $\mathcal{B}$.

Definition 16.3.4. The ***doubling constant*** of a metric space is the smallest number of balls of the same radius needed to cover a ball of twice the radius (formally, we take the maximum such number over all possible balls to be covered). The doubling constant of $\mathbb{R}^d$ is $\mathfrak{c} \leq 2^{O(d)}$ [Ver05].

Theorem 16.3.5 ([MTTV97]). *Let $\mathcal{B}$ be a set of n balls that is k-ply in $\mathbb{R}^d$. Then, there exists a sphere $\mathbb{S}^{(d)}$ that intersects $4k^{1/d}n^{1-1/d}$ balls of $\mathcal{B}$. Furthermore, the number of balls of $\mathcal{B}$ that are completely inside (resp. outside) $\mathbb{S}^{(d)}$ is $\geq n/(\mathfrak{c}+1)$.*

Proof: Let P be the set of centers of the balls of $\mathcal{B}$. As above, let b be the smallest ball containing $n/(1+\mathfrak{c})$ points of P. As above, assume that b is centered at the origin and has radius 1. Let $\mathbb{S}^{(d)}$ be a random sphere centered at the origin with radius x picked randomly from the range $[1,2]$.

Now, arguing as above, there are at most $(\mathfrak{c}/(\mathfrak{c}+1))n$ points of P inside $\mathbb{S}^{(d)}$, and as such, at least $(1-\mathfrak{c}/(\mathfrak{c}+1)) = n/(\mathfrak{c}+1)$ points of P outside $\mathbb{S}^{(d)}$. As such $\mathbb{S}^{(d)}$ is a good separator for the balls.

As for the expected number of balls intersecting $\mathbb{S}^{(d)}$, let $v_d r^d$ be the volume of a ball of radius r in $\mathbb{R}^d$, where v_d is a constant that depends on the dimension. As above, we clip the balls of $\mathcal{B}$ to the ball of radius 2 centered at the origin, replacing every lens, by a an appropriate ball of the same volume. Let r_i denote the radius of the ith such ball f_i, for $i = 1, \ldots, n$. By the k-ply property, we have that

$$\sum_i r_i^d = \frac{1}{v_d}\left(\sum_i v_d r_i^d\right) \leq \frac{k}{v_d}\,\mathrm{Vol}\big(\mathrm{ball}(2)\big) \leq k2^d,$$

where ball(2) denotes a ball of radius 2 in $\mathbb{R}^d$. As before, the probability of the ith ball to intersect $\mathbb{S}^{(d)}$ is bounded by $2r_i$. Let S be the set of balls of $\mathcal{B}$ that intersects $\mathbb{S}^{(d)}$. We have, by Hölder's inequality, that

$$\begin{aligned}\mathbb{E}\Big[|S|\Big] &= \sum_i \mathbb{P}\Big[\mathsf{f}_i \cap \mathbb{S}^{(d)} \neq \emptyset\Big] \leq \sum_i 2r_i = 2\sum_i 1 \cdot r_i \leq 2\left(\sum_{i=1}^n 1^{d/(d-1)}\right)^{(d-1)/d}\left(\sum_{i=1}^n r_i^d\right)^{1/d}\\ &\leq 2n^{1-1/d}\Big(k2^d\Big)^{1/d} \leq 4n^{1-1/d}k^{1/d},\end{aligned}$$

as desired. ∎

16.3.4. Separators for the kth nearest neighbor graph

Let P be a set of n points in $\mathbb{R}^d$, and let k be a parameter. The ***kth nearest neighbor graph*** $G_k = (P, E)$ is the graph, where two points $p, q \in P$ are connected by an edge $pq \in E$, if q is the ith nearest neighbor of p in P (or p is the ith nearest neighbor of q), for $i \leq k$.

Theorem 16.3.6 ([MTTV97]). *Let P be a set of n points in $\mathbb{R}^d$, and let k be a parameter. The kth nearest neighbor graph $G_k = (P, E)$ has a separator of size $O(k^{1/d}n^{1-1/d})$, such that each connected component has at most $(\mathfrak{c}/(\mathfrak{c}+1))n$ vertices, where $\mathfrak{c}$ is the doubling constant of $\mathbb{R}^d$, see Definition 16.3.4.*

Proof: We follow the proof of Miller et al. [MTTV97]. A point $q \in P$ is an ***i-client*** of $p \in P$, if p is the ith nearest neighbor of q, for $i \leq k$. If q is a k-client of p, then create a ball of radius $\|p - q\|$ centered at q. Let $\mathcal{B}$ be the resulting set of n balls. The key observation is that this set of balls is $O(k)$-ply – which we reprove here using a standard argument.

We claim that every point $p \in P$ can serve at most $O(k)$ clients. To this end, cover the sphere of directions around p with cones with angular diameter at most 30°. It is easy to verify that at most $c = 2^{O(d-1)}$ such cones are needed.

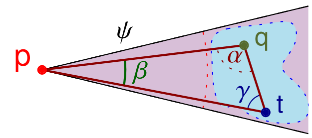

The key observation is now that for any two points $q, s \in P$ that belong to the same cone ψ of p, it must be that $\|q - s\| \le \|p - s\|$, assuming that q is closer to p than s, as an easy geometric argument shows. That is, if $q_1, \ldots, q_k$ are the k closest points to p in $P \cap \psi$, then these are the only points of $P \cap \psi$ that might be k-clients of p. It follows that p can have at most ck k-clients, and as such its degree in G_k is $\le ck + k$. That is, the maximum degree of a vertex in G_k is $O(k)$.

To see why this implies that the set of balls $\mathcal{B}$ is k-ply, consider any point $p \in \mathbb{R}^d$, insert it into P, and observe that the degree of p in the graph G_{k+1} bounds the number of balls of $\mathcal{B}$ that cover it. By the above, this is $O(k)$, as desired.

By Theorem 16.3.5, there are $4k^{1/d}n^{1-1/d}$ balls of $\mathcal{B}$, such that their removal breaks the intersection graph of $\mathcal{B}$ into connected components each of size at most $(\mathfrak{c}/(\mathfrak{c}+1))n$. Clearly, the corresponding set of points of P is the desired separator of G_k. ■

16.3.5. Separator for r vertices in a planar graph

Our purpose here is to show that in a triangulated planar graph, there is always a cycle of size $O(\sqrt{r})$ that its removal separates (roughly) r vertices from remainder of the graph. To this end, we need the following.

Lemma 16.3.7. *Let $\mathcal{B}$ be a set of n balls in $\mathbb{R}^d$ that are interior disjoint, and let $r > 0$ be some prespecified integer number. Let* b *be the smallest ball that contains r centers of the balls of $\mathcal{B}$. Then* b *intersects at most $(\mathfrak{c})^2(r+1)$ balls of $\mathcal{B}$. Furthermore,* $2\mathsf{b}$ *intersects at most $(\mathfrak{c})^3(r+1)$ balls of $\mathcal{B}$, where $\mathfrak{c}$ is the doubling constant of $\mathbb{R}^d$, see Definition 16.3.4.*

Proof: Assume b is of radius one and it is centered at the origin. Consider the ball $4\mathsf{b}$, and observe that it can be covered by $(\mathfrak{c})^2$ balls of radius one, and let $\mathcal{B}'$ be this set of balls. As such, $4\mathsf{b}$ contains at most $(\mathfrak{c})^2 r$ centers of balls of $\mathcal{B}$. Any other ball of $\mathcal{B}$ that intersect b must be radius at least 3, as its center is at distance at least 4 from the origin.

It is easy to verify that such a ball b' must contain fully at least one ball of $\mathcal{B}'$. Indeed, consider the segment connecting the center of b' with the origin, and consider the point on this segment on $\partial 4\mathsf{b}$. Clearly, this point must be covered by one of the balls of $\mathcal{B}'$, and this ball is fully contained in b'. ■

Lemma 16.3.8. *Let G be a planar graph with n vertices, and let $r > 0$ be an integer number which is sufficiently large. There exists a set of vertices S of size $\le 4\ell_2\sqrt{r}$, such that $G \setminus S$ is disconnected into two sets of vertices, X and Y, such that $r/2\ell_2 \le |X| \le r$, where ℓ_2 is a constant (see Definition 16.3.4). Furthermore, if G is triangulated then S is a cycle in the graph.*

Proof: Let $\mathcal{B}$ be the realization of G as a kissing graph of interior disjoint disks. Let d be the smallest disk containing r/ℓ_2 centers of $\mathcal{B}$, and assume that it is of radius one and centered at the origin. Lemma 16.3.7 implies that $2\mathsf{d}$ intersects at most $r(\ell_2)^2$ disks of $\mathcal{B}$, and let $\mathcal{B}'$ be this set of balls. Now consider the circle C_x centered at the origin of radius x, where x is picked randomly and uniformly from the range $[1, 2]$. Let S be the set of disks of $\mathcal{B}'$ that intersects C_x.

Now, by the analysis of Lemma 16.1.2, the expected number of disks of $\mathcal{B}'$, and thus of $\mathcal{B}$ that intersects C_x is $\le 4\sqrt{|\mathcal{B}'|} \le 4\ell_2\sqrt{r}$. This implies that the number of disks strictly inside C_x is at least $r/\ell_2 - 4\ell_2\sqrt{r} \ge r/2\ell_2$, if $r \ge 64(\ell_2)^4$. Similarly, it is easy to argue that C_x contains at most r disks of $\mathcal{B}$. ■

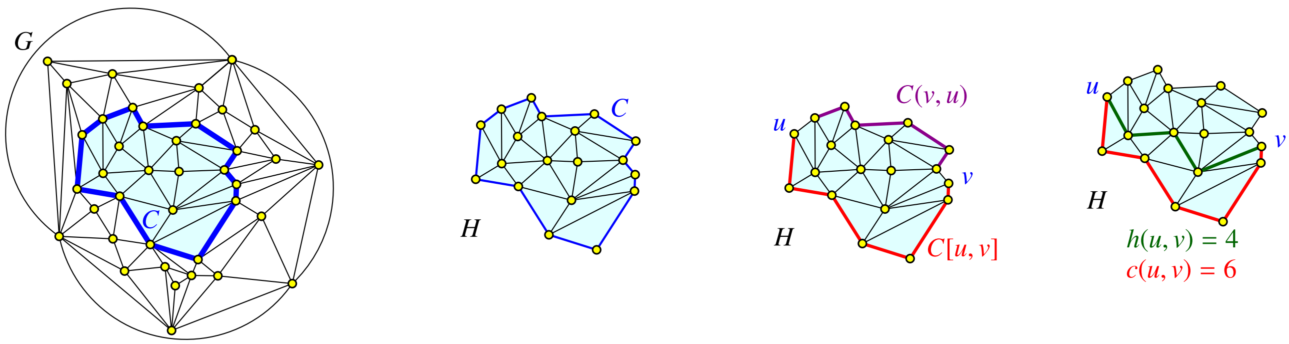

Figure 16.4: Some definitions.

16.4. A short balanced cycle and the planar separator theorem

Here, we present a third proof of the planar separator theorem, which relies on the idea that there should be a short cycle that separates the graph in a balanced way. Intuitively, the shortest such cycle, that is a balanced separator, can not be made shorter, implying that there must be many vertices on both sides of it.

We provide this additional proof because it is technically interesting and different than the previous versions. It is probably less useful than the previous versions and reading it is optional.

16.4.1. The separator definition

Let G be a triangulation with n vertices. A simple cycle C in the embedding of G forms a disk, and the vertices of G are partitioned into three sets $\mathrm{in}(C), V(C)$, and $\mathrm{out}(C)$ – the vertices inside C, on C, and outside C, respectively. Let $n_{\mathrm{in}}(C) = |\mathrm{in}(C)|$, $n(C) = |V(C)|$, and $n_{\mathrm{out}}(C) = |\mathrm{out}(C)|$.

Let $k = \lfloor \sqrt{2n} \rfloor$, and let $\mathcal{C}$ be the set of simple cycles C' in G, such that:

(A) $n(C') \le 2k$, and

(B) $n_{\mathrm{out}}(C') \le 2n/3$.

Clearly, this family is not empty, as it contains all the boundary of the faces of the graph as members (here, the cycles are reversed, and the interior of the face is their "outer" side). Let C be the most "balanced" separator in this family. Formally,

$$C = \arg\min_{C' \in \mathcal{C}} \Bigl(n_{\mathrm{in}}(C') - n_{\mathrm{out}}(C') \Bigr). \tag{16.1}$$

The claim is that C is the desired separator. That is, $n_{\mathrm{in}}(C), n_{\mathrm{out}}(C) \le 2n/3$.

16.4.2. Correctness

Consider the planar graph H induced by G on the set of vertices $\mathrm{in}(C) \cup V(C)$. The graph H is a triangulated planar graph having C as the boundary of its outer face.

For two vertices $u, v \in V(C)$, let $c(u,v) = \mathrm{d}_C(u,v)$ and $h(u,v) = \mathrm{d}_H(u,v)$ be the number of edges in the shortest path between u and v in C and H, respectively.

In the following, for two vertices $u, v \in V(C)$, let $C[u,v]$ be the path formed by tracing C from u to v, in a counterclockwise direction. Similarly, let $C(v,u)$ be he path resulting from removing $C[u,v]$ from C. See Figure 16.4.

Lemma 16.4.1. *If $n_{\mathrm{in}}(C) \ge (2/3)n$ then, for all $u, v \in V(C)$, we have $c(u,v) = h(u,v)$.*

Proof: As $C \subseteq H$, we readily have that $c(u,v) \geq h(u,v)$. To prove that $c(u,v) \leq h(u,v)$, we assume for the sake of contradiction that there is pair $u, v \in V(C)$, such that

(i) $c(u,v) > h(u,v)$, and

(ii) $h(u,v)$ is minimal among all such pairs.

Let π be a path in H between u and v with $h(u,v)$ edges. If the interior of π visits C, then there is a vertex $w \in \mathrm{int}(\pi) \cap V(C)$, and then

$$h(u,w) + h(w,v) = h(u,v) < c(u,v) \leq c(u,w) + c(w,v).$$

As such, either $h(u,w) < c(u,w)$ or $h(w,v) < c(w,v)$. But both options contradicts (ii) – the minimality of u, v. As such, the interior of π is disjoint from C.

The vertices u, v of π partitions C into two paths, and gluing π into the two paths, results into two new cycles C_1 and C_2. Here $\mathrm{in}(C_1) \cup V(\mathrm{int}(\pi)) \cup \mathrm{in}(C_2) = \mathrm{in}(C)$, and assume that $n_{\mathrm{in}}(C_1) \geq n_{\mathrm{in}}(C_2)$. Since $\pi \subseteq C_1$, we have

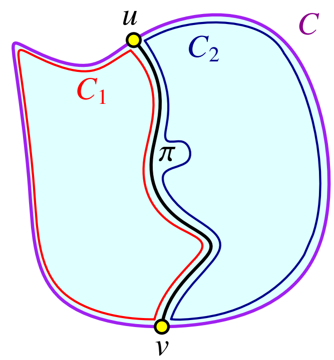

$$\begin{aligned} n - n_{\mathrm{out}}(C_1) &= n_{\mathrm{in}}(C_1) + n(C_1) > n_{\mathrm{in}}(C_1) + n(\mathrm{int}(\pi)) \\ &= \frac{2n_{\mathrm{in}}(C_1) + 2n\big(\mathrm{int}(\pi)\big)}{2} \geq \frac{n_{\mathrm{in}}(C_1) + n_{\mathrm{in}}(C_2) + n(\mathrm{int}(\pi))}{2} \\ &= \frac{n_{\mathrm{in}}(C)}{2} \geq \frac{1}{3}n, \end{aligned}$$

It follows that $n_{\mathrm{out}}(C_1) < (2/3)n$ – namely, condition (B) holds for C_1.

Assume that $C(v,u)$ is the portion of C that does not appear in C_1. By definition, we have that $c(u,v) - 1 = \min\Big(n(C(v,u)), n(C(u,v))\Big)$. Thus, we have

$$2k \geq n(C) = n(C_1) - \big(n(\sigma) - 2\big) + n\big(C(v,u)\big) \geq n(C_1) - (h(u,v) - 1) + (c(u,v) - 1) > n(C_1),$$

as $c(u,v) > h(u,v)$. Namely, condition (A) holds for C_1, and $C_1 \in \mathcal{C}$. Since C minimizes the imbalance between the inside and the outside, see Eq. (16.1), we have

$$n_{\mathrm{in}}(C_1) - n_{\mathrm{out}}(C_1) \geq n_{\mathrm{in}}(C) - n_{\mathrm{out}}(C) \iff 0 \geq n_{\mathrm{in}}(C_1) - n_{\mathrm{in}}(C) \geq n_{\mathrm{out}}(C_1) - n_{\mathrm{out}}(C).$$

We conclude that $n_{\mathrm{out}}(C) \geq n_{\mathrm{out}}(C_1)$, which implies that $n_{\mathrm{out}}(C) = n_{\mathrm{out}}(C_1)$. That implies that $C(v,u)$ does not contain any vertex. Namely, $vu \in E(C)$ and $c(u,v) = 1$. But this implies that $1 = c(u,v) > h(u,v) \implies h(u,v) = 0$, which is impossible. ■

Lemma 16.4.2. *If $n_{\mathrm{in}}(C) \geq (2/3)n$ then $|V(C)| = 2k$.*

Proof: Assume for the sake of argument that $|V(C)| < 2k$, and consider a triangular face $f = \triangle uvz$ of H such that uv is an edge of C.

If $z \in V(C)$, then since (by assumption) $n_{\mathrm{in}}(C) > 0$, it follows that $C \neq uvz$. It must be that either uz or vz are not edges of C. If $uz \in E(H) \setminus E(C)$ then $1 = h(u,z) < c(u,z)$, which contradicts Lemma 16.4.1 (alternatively apply the argument to vz). As such z is not in $V(C)$ (i.e., it is an interior vertex to the region bounded by C).

Next consider the exchange of removing the edge uv from C, and replacing it by the path uz, zv. Let C' be the resulting cycle. Clearly, $n_{\mathrm{out}}(C') = n_{\mathrm{out}}(C)$, $n_{\mathrm{in}}(C') = n_{\mathrm{in}}(C) - 1$, and $|V(C')| \leq 2k$. Namely, $C' \in \mathcal{C}$ and $n_{\mathrm{in}}(C') - n_{\mathrm{out}}(C') < n_{\mathrm{in}}(C') - n_{\mathrm{out}}(C')$, which is a contradiction to the choice of C, see Eq. (16.1). ■

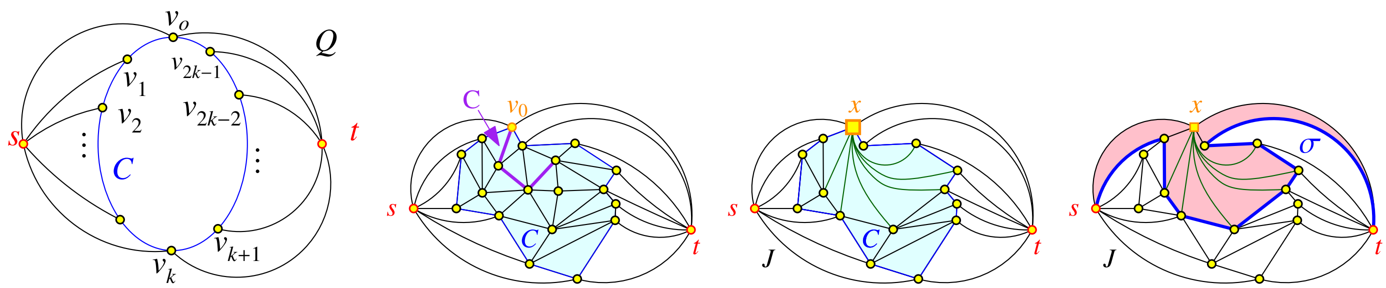

Figure 16.5

We need the following classical result.

Theorem 16.4.3 (Menger's Theorem). *Let G be an undirected graph with two vertices s and t. The size of the minimum vertex cut between s and t (i.e., the minimum number of vertices one has to remove from G to disconnect s from t) is equal to the maximum number of (interior) vertex-disjoint paths between s and t.*

Lemma 16.4.4. *If $n_{\text{in}}(C) \geq (2/3)n$ then $|V(H)| \geq (k+1)^2/2$.*

Proof: Let $C = \langle v_0, v_1, \ldots, v_{2k-1} \rangle$. Add a source vertex s and connect it by $k+1$ edges to $v_0, \ldots, v_k$. Similarly, create a target vertex t and connect $v_k, \ldots, v_{2k}$ to t. Let Q be the resulting graph. See Figure 16.5.

Let $S \subseteq V(Q)$ be a minimum vertex cut between s and t in Q, and let $\alpha = |S|$. Observe that the graph Q is a planar graph with all interior faces being triangles, and with s, t, v_0, v_k being on the outer face. Furthermore, consider the drawing where s (resp. t) is the leftmost (resp. rightmost) vertex in the drawing. Observe that $v_0, v_k \in S$, as otherwise S would not be an s–t vertex cut.

Consider the induced subgraph Q_S. If Q_S is disconnected, then let C be its connected component that contains v_0. We contract C to a single vertex in Q (removing parallel and self loop that are created by this process), and let J be the resulting graph. Let x be the vertex that corresponds to the contracted J – it lies on the top part of the outer face of the drawing of J. All its internal faces of J are triangles. The path sxt is in J. Consider the union of all triangular faces adjacent to x, that are between sx and xt. This union is a simply connected region (otherwise there would be parallel edges), and its boundary contains a simple path σ between s and t in J (which is also a path in the graph Q), see Figure 16.5. This path can not contain a vertex of S (since all its vertices are adjacent to vertices of C), which implies that S is not a vertex cut between s and t, which is a contradiction.

As such, Q_S is connected, and it contains a path π between v_0 and v_k, and by itself π forms a cut between s and t in Q. As such, it must be that $n(\pi) = |S|$, and the edge length of π is $\alpha - 1$. In particular, applying this argument to the shortest path between v_0 and v_k in H, we conclude that $h(v_0, v_k) = \alpha - 1$. However, by Lemma 16.4.1, we have $\alpha - 1 = h(v_0, v_k) = c(v_0, v_k) = k$. Thus, $\alpha = |S| = k + 1$.

By Menger's Theorem (T16.4.3), there are $k+1$ vertex disjoint paths between s and t. The ith path in this collection, π_i corresponds to a path between v_i and v_{2k-i} in H, for $i = 0, \ldots, k$. By Lemma 16.4.1, we have that $|V(\pi_i)| - 1 \geq h(v_i, v_{2k-i}) = c(v_i, v_{2k-i}) = 2\min(i, k-i)$. As such, we have that

$$|V(H)| \geq \sum_{i=0}^{k} |V(\pi_i)| \geq \sum_{i=0}^{k} \Bigl(1 + 2\min(i, k-i)\Bigr) \geq \frac{(k+1)^2}{2},$$

by silly calculations[2]. ■

[2] If $k = 2t+1$, then $\Delta = \sum_{i=0}^{k} \min(i, k-i) = 2\sum_{i=0}^{t} i = t(t+1) = (k^2-1)/4$, and $S = k+1+2\Delta = (k^2-1)/2+k+1 = (k+1)^2/2$. If $k = 2t$, then $\Delta = t + 2\sum_{i=0}^{t-1} i = t + (t-1)t = t^2 = k^2/4$, and $S = k+1+2\Delta = k^2/2+k+1 \geq (k+1)^2/2$.

Theorem 16.4.5. *Given a triangulated planar graph H with n vertices, the cycle C (defined in Section 16.4.1) is a separator in G with $\leq 2\sqrt{2n}$ vertices, and having $\leq (2/3)n$ vertices of G on each side of it.*

Proof: By (A) and (B) we have $|V(C)| \leq 2k$ and $n_{\mathrm{out}}(C) \leq 2n/3$, respectively, where $k = \lfloor\sqrt{2n}\rfloor$. As such, if $n_{\mathrm{in}}(C) \leq 2n/3$ then we are done. Otherwise, by Lemma 16.4.4, we have $n = |V(G)| > |V(H)| \geq (k+1)^2/2 \geq n$, which is impossible. ■

16.5. Cycle separator

Let $G = (V, E, F)$ be a triangulated planar graph embedded on the plane, and let $n = |V|$, and $\varphi = |F|$. In this section, we describe the linear time construction for cycle separators of G.

Our construction is composed of three phases. First, we find a possibly long cycle separator S, by finding a spanning tree $\mathcal{T}$ of G, and a balanced edge separator $(uv)^*$ in its dual tree. The unique cycle in $\mathcal{T} \cup \{uv\}$ is guaranteed to be a (possibly long) cycle separator (Section 16.5.1). This part of the construction is similar to Lemma 2 of Lipton and Tarjan [LT79], and we include the details for completeness. Next, we build a nested sequence of cycles $C_1 \preceq C_2 \preceq \ldots \preceq C_k$ (Section 16.5.2). The specific construction of these cycles, which is guided by S, is perhaps the central insight of this paper that results in our simple algorithms. Finally, we consider the collection of all cycles $C_1, \ldots, C_k$ and S to construct a set of short cycles one of which is guaranteed to be a balanced separator (Section 16.5.3).

16.5.1. A possibly long cycle separator

We start by computing a balanced separator that unfortunately can be too long.

For a BFS tree $\mathcal{T}$, we denote by $\pi(\mathcal{T}, u)$ the unique shortest path in $\mathcal{T}$ between the root of $\mathcal{T}$ and u.

Lemma 16.5.1 ([LT79]). *Given a triangulated planar graph G, one can compute, in linear time, a BFS tree $\mathcal{T}$ rooted at a vertex* root*, and an edge $uv \in E(G)$, such that:*
(A) the (shortest) paths $p_u = \pi(\mathcal{T}, u)$ and $p_v = \pi(\mathcal{T}, v)$ are edge disjoint,
(B) the cycle $S = p_u \cup p_v \cup uv$ is a 2/3-separator for G.

Proof: Our proof is a slight modification of the one provided by Lipton and Tarjan [LT79], and we include it for the sake of completeness. Let $r' \in V$ be any vertex, and let $\mathcal{T} = (V_{\mathcal{T}}, E_{\mathcal{T}})$ be a BFS tree rooted at r'. Also, let $D = E \setminus E_{\mathcal{T}}$, and note that the dual set of edges D^* is a spanning tree of the dual G^*. Since G is a triangulation, D^* has maximum degree at most three. Thus, it contains an edge $(uv)^*$ whose removal leaves two connected components, D^*_{in} and D^*_{out}, each with at most $\lceil(2/3)\varphi\rceil$ (dual) vertices, see Lemma 16.2.3, where $\varphi = |F|$ is the number of faces of G. Let D^*_{out} be the component that contains the dual of the outer face, and let D^*_{in} be the other one.

Let uv be the original edge that is dual of uv^*, and S the unique cycle in $\mathcal{T} \cup \{uv\}$. The sets of faces inside and outside S correspond to vertex sets of D^*_{in} and D^*_{out}, respectively. Thus, S is a 2/3-cycle separator.

Now, let root be the lowest common ancestor of u and v in $\mathcal{T}$. The cycle S is composed of $p_u = \mathcal{T}[r, u]$, $p_v = \mathcal{T}[r, v]$ and the edge uv. Since $\mathcal{T}$ is a BFS tree, and root is an ancestor of u and v, the paths p_u and p_v are shortest paths in G.

To get a BFS tree rooted at root, one simply recompute the BFS tree starting from root, where we include the edges of p_u and p_v in the newly computed BFS tree $\mathcal{T}$. ■

For the rest of the algorithm, let S, root, uv, p_u and p_v be given by Lemma 16.5.1. We emphasize that the graph is unweighted, p_u and p_v are shortest paths, and u and v are neighbors.

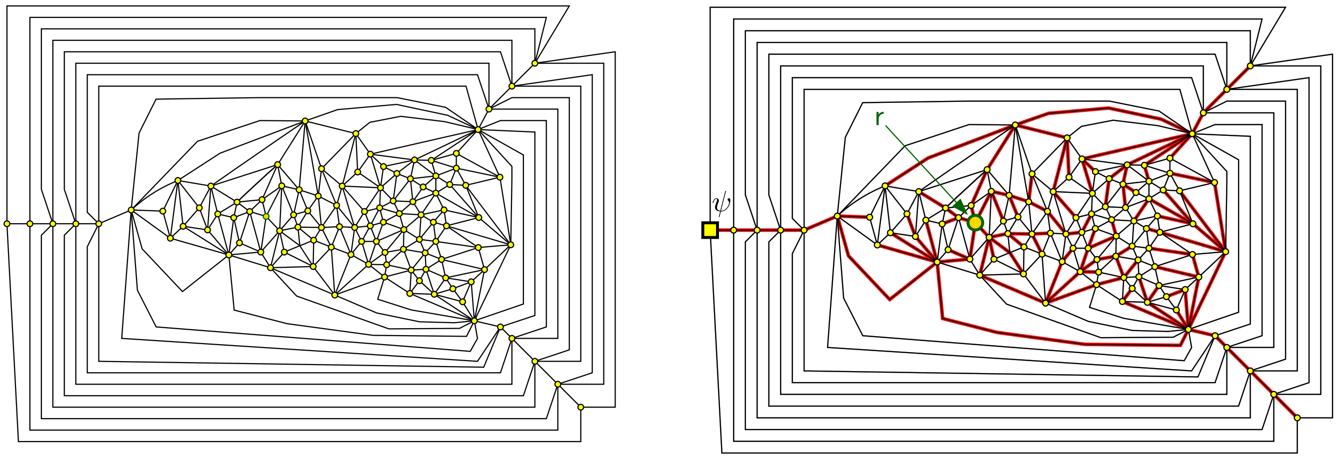

Figure 16.6: A graph and its BFS tree.

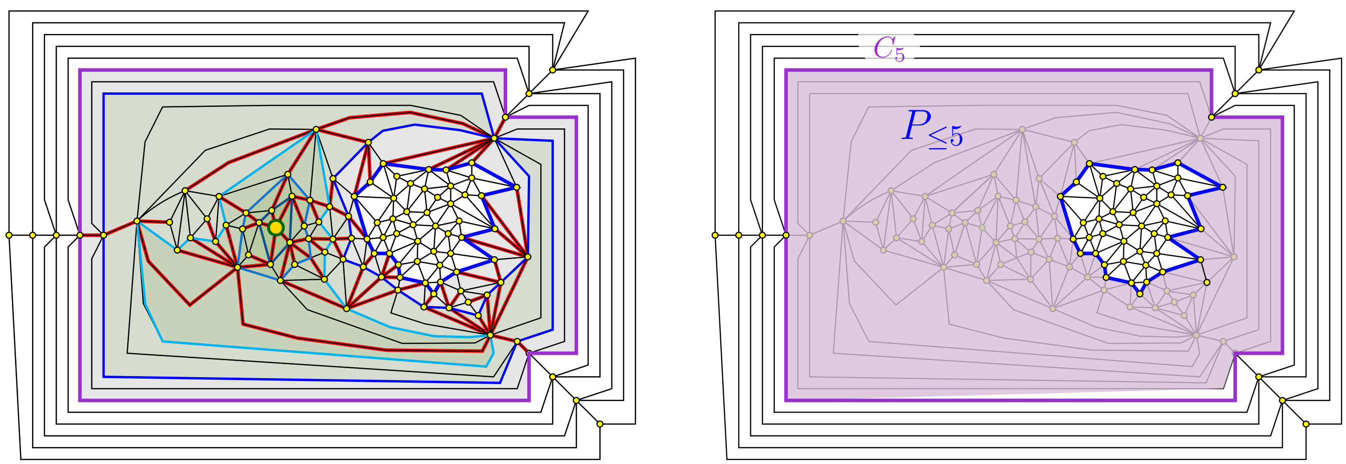

Figure 16.7: The region $P_{\le 5}$ and the associated outer cycle C_5.

16.5.2. A nested sequence of short cycles

Let root be the root node of the BFS tree $\mathcal{T}$ computed by Lemma 16.5.1. For $x \in V(G)$, let $\ell(x)$ be the distance in $\mathcal{T}$ of x from the root root. The ***level*** of a (triangular) face $\eta = xyz$ of G is $\ell(\eta) = \max(\ell(x), \ell(y), \ell(z))$. In particular, a face $\eta = uvz \in \mathrm{faces}(G)$ is ***i-close*** to root if $\ell(\eta) \le i$. The union of all i-close faces, form a region $P_{\le i}$ in the plane[3]. This region is simple, but it is not necessarily simply connected.

Let $\mathsf{h} = \max(\ell(u), \ell(v))$, and let $\psi \in \{u, v\}$ be the vertex realizing h. We assume, for the sake of simplicity of exposition, that ψ is one of the vertices of the outer face[4].

For $i < \mathsf{h}$, let ξ_i be the outer connected component of $\partial P_{\le i}$. This is a closed curve in the plane, with ψ being outside it (as long as $i < \mathsf{h}$), and let C_i be the corresponding cycle of edges in G that corresponds to ξ_i. The resulting set of cycles is $C_0, \ldots, C_{\mathsf{h}-1}$ (i.e., a cycle C_i is empty if $i \ge \mathsf{h}$).

Lemma 16.5.2. *We have the following:*

(A) For any $i < \mathsf{h}$, the vertices of C_i are all at distance i from root *in $\mathcal{T}$.*

---

[3] Here, conceptually, we consider the embedding of the edges of G to be explicitly known, so that $P_{\le i}$ is well defined. The algorithm does not need this explicit description.

[4] This can be ensured by applying inversion to the given embedding of G – but it is not necessary for our algorithm.

(B) For any $i < \mathsf{h}$, the cycle C_i is simple.
(C) For any $i < j < \mathsf{h}$, the cycles C_i and C_j are vertex disjoint.
(D) For $i < \mathsf{h}$, the cycle C_i intersects the cycle S.

Proof: (A) Consider a vertex x in G with $\ell(x) < i$. As $\mathcal{T}$ is a BFS tree, we have that all the neighbors y of x in G, have $\ell(y) \leq \ell(x) + 1 \leq i$. Namely, all the triangles adjacent to x are i-close, and the vertex x is internal to the region $P_{\leq i}$, which implies that it can not appear in C_i.

(B) Since ξ_i is the (closure) of the outer boundary of a connected set, the corresponding cycle of edges C_i is a cycle in the graph. The bad case here is that a vertex x is repeated in C_i more than once. But then, x is a cut vertex for $P_{\leq i}$ – removing it disconnects $P_{\leq i}$ – see Figure 16.8. Now, $\ell(x) < i$ as the BFS from root must have passed through x from one side of $P_{\leq i}$ to the other side. Arguing as in (A), implies that x is internal to $P_{\leq i}$, which is a contradiction.

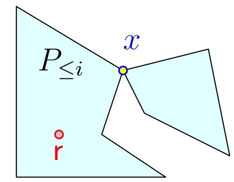

Figure 16.8

(C) is readily implied by (A).

(D) Indeed, C_i must intersect the shortest path from root to ψ, and as this path is part of S, the claim follows. ■

Computing the cycles C_i, for all i, can be done in linear time (without the explicit embedding of the edges of G). To this end, compute for all the (triangular) faces of G their level, mark all the edges between faces of level i and $i+1$ as boundary edges forming $\partial P_{\leq i}$ – this yields a collection of cycles. To identify the right cycle, consider the shortest p_ψ path between root and ψ. The cycle with a vertex that belongs to π, is the desired cycle C_i. Clearly, this can be done in linear time overall for all these cycles.

Lemma 16.5.3. *Let $\Delta > 0$ be an arbitrary parameter. If $\mathsf{h} = \ell(\psi) > \Delta$, then there exist an integer $i_0 \in [\![\Delta]\!]$, such that $\left|C_{i_0}\right| > 0$ and $\sum_{j\geq 0} \left|C_{i_0+j\Delta}\right| \leq n/\Delta$, where $\left|C_k\right|$ denotes the number of vertices of C_k.*

Proof: Setting $g(i) = \sum_{j\geq 0} \left|C_{i+j\Delta}\right|$. By Lemma 16.5.2 (D), $g(i) > 0$, for $i = 0, \ldots, \Delta - 1$. We have

$$\sum_{i=0}^{\Delta-1} g(i) \leq \sum_{i=0}^{\Delta-1} \sum_{j\geq 0} \left|C_{i+j\Delta}\right| = \sum_{k\geq 0}^{\mathsf{h}-1} \left|C_k\right| \leq |V(G)| \leq n,$$

as the cycles $C_0, C_1, \ldots, C_{\mathsf{h}-1}$ are disjoint. As such, there must be an index $i = i_0$ of the first summation that does not exceed the average. ■

16.5.3. Constructing cycle separators

Let $\Delta = \Theta(\sqrt{n})$ be a parameter to be specified shortly. Let S be a 2/3-cycle separator, and root, u, v, p_u, and p_v as given by Lemma 16.5.1. If $|S| \leq 2\Delta$ then this is the desired a short cycle separator. So, assume that $\mathsf{h} \geq |S|/2 > \Delta$.

For $j \geq 0$, let $\alpha_j = i_0 + (j-1)\Delta$ be the index of the jth cycle in the small "ladder" of Lemma 16.5.3. Since $\mathsf{h} > \Delta$ and by Lemma 16.5.2 (D), the cycles $C_{i_0} = C_{\alpha_0}$ of the ladder intersects S. In particular, let $D_j = C_{\alpha_j}$, for $j = 1, \ldots, k-1$, be the jth nested cycles of this light ladder that intersects S. Specifically, let k the minimum value such that $\alpha_k \geq \mathsf{h}$. Let D_0 be the trivial cycle formed by the root vertex root. Similarly, let D_k be the trivial cycle of the ψ, such that its interior contains the whole graph.

For $j = 0, \ldots, k$, let φ_j be the number of faces in the interior of D_j. If for some j, we have that $\lfloor \varphi/3 \rfloor \leq \varphi_j \leq \lceil (2/3)\varphi \rceil$, then D_j is the desired separator, as its length is at most n/Δ by Lemma 16.5.2, where φ is the number of faces of G.

Otherwise, there must be an index i, such that $\varphi_i < \varphi/3$, and $\varphi_{i+1} > (2/3)\varphi$. Assume, for the sake of simplicity of exposition that $0 < i < k-1$ (the cases that $i = 0$ or $i = k-1$ are degenerate and can be handled in a similar fashion to what follows).

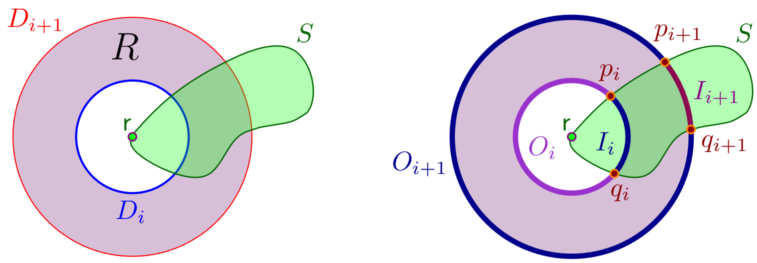

Figure 16.9

Consider the “heavy” ring R bounded by the two of the nested cycles D_{i+1} and D_i, see Figure 16.9.

Observation 16.5.4. *By Lemma 16.5.2, the cycles D_i and D_{i+1} intersects S in two vertices exactly. And D_i is nested inside D_{i+1}.*

Let I_i and O_i the portions of D_i inside and outside S, respectively (define I_{i+1} and O_{i+1} similarly). Let p_i and q_i (resp., p_{i+1} and q_{i+1}) be the end points of I_i (resp., I_{i+1}), such that p_i is adjacent to p_{i+1} along S.

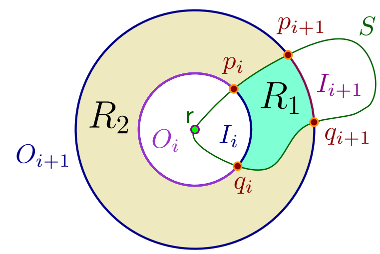

Figure 16.10

We can now partition R into two cycles R_1 and R_2. The region R_1 is bounded by the cycle formed by $\mathsf{D}_1 = S[q_i, q_{i+1}] \circ I_{i+1} \circ S[p_{i+1}, p_i] \circ I_i$. The region R_2 is bounded by the cycle formed by $\mathsf{D}_2 = S[q_i, q_{i+1}] \circ O_{i+1} \circ S[p_{i+1}, p_i] \circ O_i$, see Figure 16.10.

We have that $|\mathsf{D}_1| \leq |D_i| + |D_{i+1}| + 2\Delta \leq n/\Delta + 2\Delta$, by Lemma 16.5.3. In particular, if $\varphi(R_1) \geq \varphi/3$, then D_1 is the desired cycle separator, since $\varphi(R_1) \leq \varphi(S) \leq \lceil (2/3)\varphi \rceil$.

Similarly, if $\varphi(R_2) \geq \varphi/3$, then D_2 is the desired cycle separator, since $\varphi(R_2) \leq \varphi - \varphi(S) \leq \lceil (2/3)\varphi \rceil$.

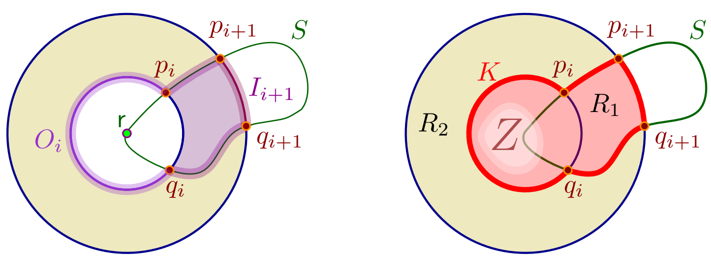

Figure 16.11

Lemma 16.5.5. *Assume that $\varphi(R_1) < \varphi/3$ and $\varphi(R_2) < \varphi/3$. Consider the region Z, formed by the union of the interior of D_i, together with the interior of R_1. Its boundary, is the cycle K formed by $O_i \circ S[q_i, q_{i+1}] \circ I_{i+1} \circ S[p_{i+1}, p_i]$, see Figure 16.11. The cycle K is a 2/3-cycle separator with $n/\Delta + 2\Delta$ edges.*

Proof: We have the following: (i) $\varphi_i < \varphi/3$, (ii) $\varphi_i + \varphi(R_1) + \varphi(R_2) = \varphi_{i+1} > (2/3)\varphi$, (iii) $\varphi(R_1) < \varphi/3$, and (iv) $\varphi(R_2) < \varphi/3$. Assume that $\varphi_i + \varphi(R_1) < \varphi/3$. But then $\varphi_{i+1} = \varphi_i + \varphi(R_1) + \varphi(R_2) < (2/3)\varphi$, which is impossible. The region Z bounded by K contains $\varphi_i + \varphi(R_1)$ faces, and we have $\varphi/3 < \varphi_i + \varphi(R_1) < (2/3)\varphi$, which implies the separator property.

As for the length of K, observe that $|K| \leq |D_i| + |D_{i+1}| + |S[p_i, p_{i+1}]| + |S[q_i, q_{i+1}]| \leq n/\Delta + 2\Delta$, by Lemma 16.5.3. ■

Theorem 16.5.6. *Given an embedded triangulated planar graph G with n vertices and φ faces, one can compute in linear time a simple cycle K that is a 2/3-separator of G. The cycle K has at most $O(1) + \sqrt{8n}$ edges.*

This cycle K also 2/3-separates the vertices of G – namely, there are at most $(2/3)n$ vertices of G on each side of it.

Proof: The construction is described above. As for the length of K, set $\Delta = \left\lceil \sqrt{n/2} \right\rceil$, and by Lemma 16.5.5 we have $|K| \leq 2\Delta + n/\Delta \leq O(1) + \sqrt{2n} + \sqrt{2n} \leq O(1) + \sqrt{8n}$. (The separator cycle is even shorter if one of the other cases described above happens.)

As for the running time, observe that the algorithm runs BFS on the graph several times, identify the edges that form the relevant cycles. Count the number of faces inside these cycles, and finally counts the number of edges in R_1 and R_2. Clearly, all this work (with a careful implementation) can be done in linear time.

The second claim follows from a standard argument, see Lemma 16.5.7 (III) below for details. ■

16.5.4. From faces separation to vertices separation

Lemma 16.5.7. *(I) A simple planar graph G with n vertices has at most $3n-6$ edges and at most $2n-4$ faces. A triangulation has exactly $3n-6$ edges and $2n-4$ faces. (II) Let G be a triangulated planar graph and let C be a simple cycle in it. Then, there are exactly $(\varphi(C) - |C|)/2 + 1$ vertices in the interior of C, where $\varphi(C)$ denotes the number of faces of G in the interior of C. (III) A simple cycle C in a triangulated graph G that has at most $\lceil (2/3)\varphi \rceil$ faces in ts interior, contains at most $(2/3)n$ vertices in its interior, where n and φ are the number of vertex and faces of G, respectively.*

Proof: (A) is an immediate consequence of Euler's formula.

(B) Let n be the number of vertices of G in or on C – delete the portion of G outside C, and add a vertex v to G outside C, and connect it to all the vertices of C. The resulting graph is a triangulation with $n+1$ vertices, and $2(n+1) - 4 = 2n - 2$ triangles, by part (A). This counts $|C|$ triangles that were created by the addition of v. As such, $\varphi(C) = 2n - 2 - |C| \implies n = \varphi(C)/2 + 1 + |C|/2$. The number of inner vertices is $n - |C| = (\varphi(C) - |C|)/2 + 1$.

(C) Part (B) implies that number of vertices internal to the cycle C is at most

$$
\begin{aligned}
(\varphi(C) - |C|)/2 + 1 &\leq (\lceil (2/3)\varphi \rceil - |C|)/2 + 1 = (\lceil (2/3)(2n-4) \rceil - |C|)/2 + 1 \\
&\leq \frac{(2/3)(2n-4) + 1 - |C|}{2} + 1 \leq \frac{2}{3}n,
\end{aligned}
$$

as claimed. ■

16.6. Bibliographical notes

History. The planar separator theorem was proved by Ungar [Ung51] which provided a bound $O\big(\sqrt{n} \log n\big)$ on the separator size. Lipton and Tarjan [LT79] presented an optimal separator of size $O(\sqrt{n})$. For further details on planar separators and their applications, see Wikipedia (http://en.wikipedia.org/wiki/Planar_separator_theorem).

Presentation. The presentation in Section 16.1 follows [Har13], and demonstrates that the planar separator theorem is an easy consequence of the circle packing theorem, originally proved by Paul Koebe in 1936 [Koe36]. The circle packing theorem is thus the "true" magic – converting a topological property (a graph being planar) into a packing property (i.e., disks touching each other). Most[5] of the main ingredients of the proof of [Har13] are present in earlier work on this problem. See Miller et al. [MTTV97], Smith and Wormald [SW98], and Chan [Cha03]. Furthermore, the constants in the separator are inferior to known constructions [AST94].

The presentation in Section 16.2 follows roughly the work of Lipton and Tarjan [LT79].

The presentation of Section 16.3 follows [Har13]. For all the results in this part there are known linear time algorithms that work directly on the graph, but they tend to be significantly more tedious.

Section 16.4 is from the work by Alon et al. [AST94].

16.7. Exercises

16.8. From previous lectures

Lemma 16.8.1. *A simple planar graph G with n vertices has at most $3n - 6$ edges and at most $2n - 4$ faces. A triangulation has exactly $3n - 6$ edges and $2n - 4$ faces.*

Theorem 16.8.2 (Circle packing theorem). *Let $H = (V, E)$ be a finite simple planar graph. Then H can be realized by a set of interior-disjoint disks, where every disk corresponds to a vertex, and two disks touch, if and only if the corresponding vertices has an edge between them in the original graph.*

Definition 16.8.3. A maximal planar graph is a ***triangulation***. In any embedding of a triangulation, all its faces, including the outer face, are triangles (i.e., the boundary of a face is a cycle with three edges).

References

[AST94] N. Alon, P. Seymour, and R. Thomas. *Planar separators*. *SIAM J. Discrete Math.*, 2(7): 184–193, 1994.

[Cha03] T. M. Chan. *Polynomial-time approximation schemes for packing and piercing fat objects*. *J. Algorithms*, 46(2): 178–189, 2003.

[Har13] S. Har-Peled. *A simple proof of the existence of a planar separator*. *ArXiv e-prints*, 2013. arXiv: `1105.0103 [cs.CG]`.

[Koe36] P. Koebe. *Kontaktprobleme der konformen Abbildung*. *Ber. Verh. Sächs. Akademie der Wissenschaften Leipzig, Math.-Phys. Klasse*, 88: 141–164, 1936.

[LT79] R. J. Lipton and R. E. Tarjan. *A separator theorem for planar graphs*. *SIAM J. Appl. Math.*, 36: 177–189, 1979.

[Mil86] G. L. Miller. *Finding small simple cycle separators for 2-connected planar graphs*. *J. Comput. Sys. Sci.*, 32(3): 265–279, 1986.

[MTTV97] G. L. Miller, S. H. Teng, W. P. Thurston, and S. A. Vavasis. *Separators for sphere-packings and nearest neighbor graphs*. *J. Assoc. Comput. Mach.*, 44(1): 1–29, 1997.

[5] And by most we mean all.

[SW98] W. D. Smith and N. C. Wormald. *Geometric separator theorems and applications*. *Proc. 39th Annu. IEEE Sympos. Found. Comput. Sci.* (FOCS), 232–243, 1998.

[Ung51] P. Ungar. *A theorem on planar graphs*. *J. London Math. Soc.*, 26: 256–262, 1951.

[Ver05] J.-L. Verger-Gaugry. *Covering a ball with smaller equal balls in* $\mathbb{R}^n$. *Discrete Comput. Geom.*, 33(1): 143–155, 2005.