MA431 Spectral Graph Theory: Lecture 9

Ahmad Abdi Neil Olver

19 Conductance and Cheeger's inequality, continued

Recall the statement of Cheeger's inequality.

Theorem 19.1 (Cheeger's inequality)**.** *With ν_2 denoting the second eigenvalue of $\mathcal{L}$, we have*

$$\tfrac{1}{2}\nu_2 \leq \phi(G) \leq \sqrt{2\nu_2}.$$

Last time we proved the easy direction under the assumption that G is d-regular. Today we prove the hard direction under the same assumption.

19.2 The hard direction

We now want to show that $\phi(G) \leq \sqrt{2\nu_2}$. Let x be an eigenvector of $\mathcal{L}$ corresponding to ν_2. Also assume that $|\{i : x_i > 0\}| \leq |V|/2$; otherwise simply replace x with $-x$. Note that x does have some positive component, since $x \perp \mathbf{1}$ and $x \neq 0$.

Now define $y \in \mathbb{R}^V$ by setting $y_v = \max\{x_v, 0\}$ for all $v \in V$. We can assume that $\max_v y_v = 1$, simply by scaling.

Lemma 19.2. $R(y) \leq R(x) = \nu_2$.

Proof. Let $Q = \{v : y_v > 0\}$. Then for $v \in Q$,

$$(\mathcal{L}y)_v = y_v - \tfrac{1}{d} \sum_{w: vw \in E} y_w \leq x_v - \tfrac{1}{d} \sum_{w: vw \in E} x_v = (\mathcal{L}x)_v = \nu_2 x_v = \nu_2 y_v.$$

Thus

$$y^\top \mathcal{L} y = \sum_{v \in Q} y_v (\mathcal{L}y)_v \leq \sum_{v \in Q} \nu_2 y_v^2 = \nu_2 \sum_{v \in V} y_v^2,$$

as required. $\square$

For any $\tau \in (0,1)$, let

$$S_\tau = \{v : y_v^2 \geq \tau\}.$$

Note that $S_\tau \neq \emptyset$ and $|S_\tau| \leq |V|/2$ for all $\tau \in (0,1)$. We will now choose τ uniformly at random from $[0,1]$ and see that "on average" this choice works.

With this choice of τ, we have

$$\mathbb{E}[\mathrm{vol}(S_\tau)] = d\sum_{v\in V}\mathbb{P}(y_v^2 \geq \tau) = d\sum_{v\in V} y_v^2,$$

from linearity of expectation. We also have

$$\begin{aligned}
\mathbb{E}[|\delta(S_\tau)|] &= \sum_{vw\in E}\mathbb{P}(vw \in \delta(S_\tau)) \\
&= \sum_{vw\in E}\mathbb{P}\Big((y_v^2 < \tau \wedge y_w^2 \geq \tau) \vee (y_v^2 < \tau \wedge y_w^2 \geq \tau)\Big) \\
&= \sum_{vw\in E}|y_v^2 - y_w^2| \\
&= \sum_{vw\in E}|y_v - y_w|(y_v + y_w) \\
&\leq \sqrt{\sum_{vw\in E}(y_v - y_w)^2}\cdot\sqrt{\sum_{vw\in E}(y_v + y_w)^2} \\
&\leq \sqrt{R(y)\cdot d\sum_{v\in V} y_v^2}\cdot\sqrt{2\sum_{vw\in E}(y_v^2 + y_w^2)} \\
&= \sqrt{2R(y)}\cdot d\sum_{v\in V} y_v^2.
\end{aligned}$$

(We have used Cauchy-Schwartz, and the simple inequality $2y_vy_w \leq y_v^2 + y_w^2$). Since $R(y) \leq \nu_2$, we can now deduce that

$$\mathbb{E}[|\delta(S_\tau)|] \leq \sqrt{2\nu_2}\,\mathbb{E}[\mathrm{vol}(S_\tau)].$$

Thus there must exist a choice t for which $|\delta(S_t)| \leq \sqrt{2\nu_2}\,\mathrm{vol}(S_t)$, and we're done.

Note that the proof gives an algorithm to find a cut S for which $\phi(S) \leq \sqrt{2\nu_2}$. Simply try all distinct choices of S_τ (of which there are less than n) and choose one of minimum conductance. This will generally not return a cut of minimum conductance, but chaining both the easy and hard directions of Cheeger together we do have that this algorithm returns a cut S for which $\phi(S) \leq 2\sqrt{\phi(G)}$.

20 Expander graphs

We will now really restrict our attention to regular graphs (not just for streamlining a proof). In what follows, unless otherwise stated, $G = (V, E)$ will always refer to a connected, d-regular graph on n vertices, with Laplacian L and adjacency matrix A. The spectrum of L is $0 = \lambda_1 < \lambda_2 \leq \cdots \leq \lambda_n$. The spectrum of A is then $d > d - \lambda_2 \geq \cdots \geq d - \lambda_n$ (since $L = dI - A$) . Note that $\lambda_n \leq 2d$, since the spectrum of A is contained in $[-d, d]$.

We will say that G has *edge expansion* α if

$$\frac{|\delta(S)|}{|S|} \geq \alpha \qquad \text{for all } \emptyset \subsetneq S \subseteq V, |S| \leq n/2.$$

In other words, if $\phi(G) \geq \alpha/d$. Since the easy direction of Cheeger tells us that $\lambda_2/(2d) \leq \phi(G)$, we can deduce that the edge expansion of G is at least $\lambda_2/2$. This motivates the following definition.

Definition 20.1. A d-regular graph G with Laplacian spectrum $\lambda_1 < \lambda_2 \leq \cdots \leq \lambda_n$ is called a (spectral) *one-sided β-expander* for some $\beta \in (0, 1]$ if $\lambda_2 \geq \beta d$. If in addition, $\lambda_n \leq 2d - \beta d$, then G is called a *two-sided β-expander*.

So a one-sided β-expander has edge expansion $\beta d/2$. What's the motivation for the two-sided definition?

Firstly, simple random walk on a two-sided expander will mix fast. Lazy random walk (or a continuous time random walk) will mix fast on a one-sided expander; this we saw already. But if G is bipartite, or close to bipartite, simple random walk may not mix at all, or very slowly. It's the largest eigenvalue of A in *absolute value* aside from the d eigenvalue that matters.

Secondly, two-sided expanders can be thought of as good sparse approximations to the complete graph, in the following sense. Recall that K_n denotes the complete graph on n vertices.

Lemma 20.2. *If G is a two-sided $(1-\epsilon)$-expander (for some $0 < \epsilon < 1$) then*

$$(1-\epsilon)L(K_n) \preceq \tfrac{n}{d}L \preceq (1+\epsilon)L(K_n).$$

In other words, for any $x \in \mathbb{R}^V$,

$$(1-\epsilon)x^\top L(K_n)x^\top \leq \tfrac{n}{d}x^\top Lx \leq (1+\epsilon)x^\top L(K_n)x.$$

This lemma tells us that for G a two-sided $(1-\epsilon)$-expander if we take the weighted graph (G, w), where $w_e = n/d$ for all $e \in E$, then the corresponding weighted Laplacian $L_w = \frac{n}{d}L$ is "close" to the Laplacian of the complete graph. Take, for example, $x = \mathbf{1}_S$ for some set S; then we can deduce from this that

$$(1-\epsilon)|\delta_{K_n}(S)| \leq \tfrac{n}{d}|\delta_G(S)| \leq (1+\epsilon)|\delta_{K_n}(S)|.$$

That is, the wight across any cut is roughly the same in (G, w) as in K_n. There is a whole theory on finding sparse graphs that "approximate" some given graph; expanders are sparse graphs that approximate the complete graph.

Proof. Note that all eigenvalues of $L(K_n)$ aside from the 0-eigenvalue have value n, by observing that $L(K_n) = nI - J$, where J is the all-ones matrix. We clearly have $\mathbf{1}^\top L(K_n)\mathbf{1} = 0 = \mathbf{1}^\top L\mathbf{1}$. So consider any $x \perp \mathbf{1}$. Then $x^\top L(K_n)x = nx^\top x$, whereas

$$d(1-\epsilon)x^\top x \leq x^\top Lx \leq d(1+\epsilon)x^\top x$$

by the CFT. This shows the claim. □

Trivially, the complete graph K_n is an $(n-1)$-regular two-sided $(1 - \frac{1}{n-1})$-expander. So in that sense, expanders certainly exist. But what's much more interesting, and much less obvious, is the existence of *sparse* expanders. Fix the degree d; for n very large, do β-expanders exist for any fixed positive $\beta > 0$? This is already not obvious. Happily, these amazing objects do exist. They have many useful properties (and are excellent candidates to reach for when looking for counterexamples or extremal examples to many problems). We won't say much (some further facts will appear as exercises).

20.1 A randomized expander construction

During the lecture, we saw a random construction not of an expander, but a bipartite graph with certain vertex expansion properties. In these notes, instead, we give the construction of a d-regular edge expander that follows along very similar lines (in particular, the same approach of using union bound over a large collection of "bad events" that some subset S has all its neighbours within some smallish set $S \cup T$). We will omit one part of the proof.

Fix d to be a constant (later, we will need to choose it large enough). Let n be large and even, and let $V = \{1, 2, \ldots, n\}$. For $i \in \{1, 2, \ldots, d\}$, let M_i be an independent, uniformly chosen perfect matching on V. In other words, we start with $M_i = \emptyset$; then repeatedly pick any node v not yet covered by M_i, match it with a uniformly chosen second node w not yet covered by M_i, and add $\{v, w\}$ to M_i; and repeat this process until all nodes are covered by M_i. Then let $G = (V, E)$ be the union of all these perfect matchings, taking multiple copies of an edge if it appears in multiple matchings.

The potential overlap of the matchings is an annoyance, since it means that G may be a multigraph (multiple copies of the same edge allowed), whereas we have been working always with simple graphs. If one doesn't mind this, it's straightforward to extend the definition of an edge expander to allow multigraphs, and one can also extend the proof of Cheeger's inequality to allow for this. The other option is to simply repeat this construction from scratch, until we get lucky and none of the matchings overlap. Fortunately, the probability of an overlap is not too large, meaning one would only need to repeat a constant number of times (the constant depends on d) on average:

Lemma 20.3. *The probability that G is d-regular is bounded from below by a positive constant depending only on d.*

We won't prove this, however; the proof is somewhat involved.

Let $N(S)$ denote the set of neighbours of S, for any $S \subseteq V$.

Lemma 20.4. *Consider some $k \leq n/2$, a set $S \subseteq V$ of size k, and a set $T \subseteq V$ of size $\lfloor k/6 \rfloor$ that is disjoint from S. Then*

$$\mathbb{P}[N(S) \subseteq S \cup T] \leq \left(\frac{k}{n}\right)^{dk/6}.$$

Proof. We will omit all floors for notational convenience.

It suffices to show that for the matching M_1, the probability that all nodes in S are matched to nodes in $S \cup T$ is bounded by $(k/n)^{k/6}$. Since all the matchings are independent, the claim then follows immediately.

Consider constructing M_1 edge by edge. We begin by picking any $v_1 \in S$, and choosing (uniformly at random) its partner w_1. If $w_1 \notin S \cup T$, we already fail. Otherwise, we pick any $v_2 \in S$ not yet covered by M_1, choose its partner w_2, etc. Think of each successive choice of an edge as a trial; we will need to succeed at least $k/2$ times (each trial can use up at most two nodes of S, depending on whether w_i is in S or in T). The success probability of the i'th trial, given that all the previous trials succeeded, is precisely $\frac{|S \cup T| - 2(i-1) - 1}{n - 2(i-1) - 1}$. So the probability of overall success is at most

$$\prod_{i=1}^{k/2} \frac{7k/6 - 2i + 1}{n - 2i + 1} \leq \prod_{i=k/3+1}^{k/2} \frac{7k/6 - 2i + 1}{n - 2i + 1} \leq \prod_{i=k/3+1}^{k/2} \frac{k/2}{n - k} \leq (k/n)^{k/6}.$$

□

Now we apply a union bound over choices of S and T to show the expansion property we desire. For $k \leq n/2$, let

$$\mathcal{R}_k = \{(S,T) : S, T \subseteq V, S \cap T = \emptyset, |S| = k \text{ and } |T| = \lfloor k/6 \rfloor\}.$$

Let $\mathcal{R} = \bigcup_{k=1}^{n/2} \mathcal{R}_k$. For $(S,T) \in \mathcal{R}$, let $Y_{S,T}$ be the indicator variable for the event that $N(S) \subseteq S \cup T$. Our goal is to show that the event that $Y_{S,T} = 0$ for *all* $(S,T) \in \mathcal{R}$ is very likely.

We have

$$\begin{aligned}
\mathbb{P}(\exists (S,T) \in \mathcal{R} : Y_{S,T} = 1) &\leq \sum_{(S,T) \in \mathcal{R}} \mathbb{P}(Y_{S,T} = 1) \\
&= \sum_{k=1}^{n/2} \sum_{(S,T) \in \mathcal{R}_k} \mathbb{P}(Y_{S,T} = 1) \\
&\leq \sum_{k=1}^{n/2} \sum_{(S,T) \in \mathcal{R}_k} \left(\frac{k}{n}\right)^{kd/6} \\
&\leq \sum_{k=1}^{n/2} \binom{n}{k} \binom{n-k}{k/6} \left(\frac{k}{n}\right)^{kd/6} \\
&\leq \sum_{k=1}^{n/2} \left(\frac{en}{k}\right)^{k} \left(\frac{k}{n}\right)^{kd/6}.
\end{aligned}$$

In the last step, we used the estimates $\binom{n-k}{k/6} \leq \binom{n}{k}$ and $\binom{n}{k} \leq (en/k)^k$. Now we require that d be large enough; $d \geq 18$ suffices. Then

$$\begin{aligned}
\mathbb{P}(\exists (S,T) \in \mathcal{R} : Y_{S,T} = 1) &\leq \sum_{k=1}^{n/2} \left(e(k/n)^2\right)^k \\
&= \sum_{k=1}^{\sqrt{n}} \left(e(k/n)^2\right)^k + \sum_{k=\sqrt{n}+1}^{n/2} \left(e(k/n)^2\right)^k \\
&\leq \sum_{k=1}^{\sqrt{n}} (e/n)^k + \sum_{k=\sqrt{n}+1}^{n/2} (e/4)^k.
\end{aligned}$$

It's not too hard to see now that the first term is $O(1/n)$ and the second term is much smaller. So this probability tends to 0 as $n \to \infty$.

So we have shown that G has edge expansion $1/6$, with probability going to 1 as $n \to \infty$; taking on faith Lemma 20.3, we can deduce that with positive constant probability, G is a simple d-regular graph with edge expansion $1/6$.

21 Tutte drawings of planar graphs

In the last few lectures, we saw how lower bounding the second least eigenvalue of the Laplacian of a graph, i.e. the spectral gap, relates to the mixing time, conductance, and also expansion properties of the graph.

In the final stretch of the course, we see how the Laplacian can be used to glean insights about structural properties of graphs. More specifically, we prove a remarkable theorem of Tutte on how to draw a 3-connected planar graph through a connection to the Laplacian. Then we prove that for such a graph, the second least eigenvalue of the Laplacian has multiplicity at most 3.

All the graphs in this section and the next, unless stated otherwise, are assumed simple. Let $G = (V, E)$ be a 3-connected planar graph. We exhibit an elegant way to embed G on the plane (so that no two edges cross), by solving a matrix equation involving the Laplacian matrix of G. This drawing is originally due to Tutte.

21.1 Barycentric extensions

Let $G = (V, E)$ be a connected graph, let $X \subseteq V$ be nonempty, and let $\phi' : X \to \mathbb{R}^2$ be an embedding of the vertices of X. A function $\phi : V \to \mathbb{R}^2$ is a *barycentric extension of* ϕ' if for every $v \in V$,

$$\phi(v) = \begin{cases} \phi'(v) & \text{if } v \in X, \\ \frac{1}{\deg(v)} \sum_{u \in N(v)} \phi(u) & \text{otherwise.} \end{cases}$$

Lemma 21.1. *Let $G = (V, E)$ be a connected graph, let $X \subseteq V$ be nonempty, and let $\phi' : X \to \mathbb{R}^2$ be an embedding of the vertices of X. Then a barycentric extension of ϕ' exists uniquely.*

Proof. Let L be the Laplacian of G. For subsets $S_1, S_2 \subseteq V$, denote by $L[S_1, S_2]$ the submatrix of L whose rows correspond to S_1 and whose columns correspond to S_2. Let $\phi : V \to \mathbb{R}^2$ be an arbitrary extension of ϕ', and let A be the $V \times \{x, y\}$ matrix whose rows are $\phi(v), v \in V$. Observe that ϕ is a barycentric extension of ϕ' if, and only if, $\deg(v) \cdot \phi(v) - \sum_{u \in N(v)} \phi(u) = \mathbf{0}$ for all $v \in V - X$. Writing the latter in matrix form, we see that ϕ is a barycentric extension if, and only if, $L[V - X, V] \cdot A = \mathbf{0}$, which can be rewritten as

$$L[V - X, V - X] \cdot A[V - X, \{x, y\}] = -L[V - X, X] \cdot A[X, \{x, y\}].$$

In the matrix equation above, the values of $A[X, \{x, y\}]$ are given by the coordinates of ϕ', whereas the values of $A[V - X, \{x, y\}]$ are determined by the unknown coordinates of ϕ on $V - X$. As $X \neq \emptyset$, $L[V - X, V - X]$

is a proper principal submatrix of L. Moreover, since G is a connected graph, $L[V-X, V-X]$ is nonsingular matrix (exercise). Thus, the values of ϕ on $V-X$ are determined uniquely as follows:

$$A[V-X,\{x,y\}] = -L[V-X,V-X]^{-1} \cdot L[V-X,X] \cdot A[X,\{x,y\}].$$

This finishes the proof. $\square$

We may therefore speak of *the* barycentric extension of a partial mapping. Now, given that $G=(V,E)$ is a 3-connected planar graph, we can describe a fairly intuitive method for a straight-line embedding of G in the plane:

Step 1 Identify a "facial" cycle $C \subseteq E$ of G.

Step 2 Let $\phi' : V(C) \to \mathbb{R}^2$ be the function that maps C to a strictly convex polygon.

Step 2 Let $\phi : V \to \mathbb{R}^2$ be the barycentric extension of ϕ'.

Remarkably, ϕ yields a *plane* embedding of G, one where no two edges cross – we prove this next time.

You may wonder where we use 3-connectedness of G. This is used in two places. We will use in Step 1 to identify a facial cycle. It turns out that 3-connected planar graphs have *essentially* a unique plane embedding, and so one can characterize (and also find) the facial cycles independently of the embedding. Secondly, 3-connectedness is used in a subtle way in Step 2 to guarantee that the barycentric extension indeed yields a plane embedding.

21.2 Plane embeddings and peripheral cycles

The reason we work with planar graphs that are 3-connected is because such graphs essentially have a unique embedding, and for such graphs, we can characterise the facial cycles independently of the embedding. To elaborate, let $G=(V,E)$ be a connected graph. A cycle $C \subseteq E$ is *peripheral* if G/C has no cut-vertex. We need the following classic result on planarity:

Theorem 21.2. *Let $G=(V,E)$ be a 3-connected planar graph. Then the following statements hold:*

1. *Every edge belongs to exactly two peripheral cycles, and the two cycles are internally vertex-disjoint.*
2. *Given a plane embedding of G, the facial cycles are precisely the peripheral cycles of G.*

Given disjoint edge sets $I, J \subseteq E$, the *minor* $G \setminus I/J$ is the (not necessarily simple) graph obtained from deleting the edges in I and contracting the edges in J. Observe that every minor of a planar graph is also planar, so because $K_{3,3}$, the complete bipartite graph with 3 vertices on each side, is not planar, we get the following useful, widely known fact:

Remark 21.3. *A planar graph has no $K_{3,3}$ minor.*

We also need the following lemma:

Lemma 21.4 (Key Lemma)**.** *Let $G = (V, E)$ be a 3-connected planar graph, let $e = uv \in E$, and let C_1, C_2 be the two peripheral cycles of G containing e. Then every path connecting $V(C_1) - \{u, v\}$ to $V(C_2) - \{u, v\}$ and every uv-path different from $\{e\}$ have a vertex in common.*

Proof. Let us embed G on the unit sphere $\mathbb{S}^2 = \{x \in \mathbb{R}^3 : x^\top x = 1\}$.[1] By Theorem 21.2, C_1, C_2 are the boundaries of the two faces containing edge e; call the two faces $F_1, F_2 \subseteq \mathbb{R}^2$, respectively. Let $F := \mathbb{S}^2 \setminus (F_1 \cup F_2)$. Observe that F is homeomorphic to a disc, and its boundary is the cycle $C := C_1 \triangle C_2$.

Let $P \subseteq E$ be a path connecting a vertex $a \in V(C_1) - \{u, v\}$ to a vertex $b \in V(C_2) - \{u, v\}$, and let Q be a uv-path different from $\{e\}$. We need to prove that P, Q have a vertex in common. Since F_1, F_2 are faces of the embedding, the two paths P, Q are disjoint from them and therefore contained in F. We know that P joins the two boundary points a, b of F, and Q joins the two boundary points u, v. Since the four points appear on the boundary of F, which is C, as a, u, b, v clockwise or anticlockwise, the two paths P, Q must have a point, which must be a vertex, in common. $\square$

21.3 Locally convex embeddings

Let $G = (V, E)$ be a 3-connected graph, and let $C \subseteq E$ be a peripheral cycle. A *locally convex embedding* of the pair (G, C) is a straight-line plane drawing of G as specified by $\phi : V \to \mathbb{R}^2$, where

1. ϕ maps C to a strictly convex polygon,
2. for every vertex $v \in V - V(C)$, $\phi(v) \in \operatorname{relint}(\operatorname{conv}(N(v)))$, that is, v is mapped to the relative interior of the convex hull of the images of the neighbours of v (here, we assume that the relative interior of a single point is the point itself).

We see in the next subsection why such an embedding exists. For now, let us prove the following:

Theorem 21.5. *Let $G = (V, E)$ be a 3-connected planar graph, and let C be a peripheral cycle. Then every locally convex embedding of (G, C) is a plane embedding of G.*

Proof. Let $\phi : V \to \mathbb{R}^2$ be a locally convex embedding of (G, C).

Claim 1. $\phi(v) \in \operatorname{int}(\operatorname{conv}(\phi(V(C))))$ for every vertex $v \in V - V(C)$.

Proof of Claim. Exercise. $\Diamond$

Claim 2. Let $a \in \mathbb{R}^2$ and $\beta \in \mathbb{R}$, and let $X := \{v \in V : a^\top \phi(v) > \beta\}$. If $X \neq \emptyset$, then the induced subgraph $G[X]$ is connected.

Proof of Claim. Suppose $X \neq \emptyset$. By Claim 1, $X \cap V(C) \neq \emptyset$, and clearly, every pair of vertices in $X \cap V(C)$ are connected via a subpath of C that appears in $G[X]$. Let $K \subseteq X$ be the connected component of $G[X]$

[1] It is a well-known fact that such an embedding can be obtained by an appropriate "lifting" of the embedding of G in $\mathbb{R}^2$.

containing the vertices of C. Let $\beta^\star = \max\{a^\top \phi(v) : v \notin K\}$, and let L collect the vertices that attain this maximum:

$$L = \{v \notin K : a^\top \phi(v) = \beta^\star\}.$$

Suppose for a contradiction $K \neq X$. Then $\beta^\star > \beta$ and so $L \subseteq X - K$. Since G is connected, $\delta(L) \neq \emptyset$, so there exists an edge uv where $u \in L$ and $v \notin L$. Since there is no edge between L and K, we have $v \notin K$, so our choice of $\beta^\star$ implies that $a^\top \phi(v) < \beta^\star$. By local convexity, there must exist another $w \in N(u)$ for which $a^\top \phi(w) > \beta^\star$. Our choice of $\beta^\star$ would then imply that $w \in K$, which is a contradiction as there is no edge between L and K. Thus, $K = X$, thereby proving the claim. $\Diamond$

Claim 3 (Nondegeneracy)**.** For every vertex $u \in V$, the points $\{\phi(v) : v \in N(u)\}$ are not collinear.

Proof of Claim. Suppose otherwise. That is, there exist $a \in \mathbb{R}^2$ and $\beta \in \mathbb{R}$ such that $a^\top \phi(v) = \beta$ for every $v \in N(u)$. Clearly, $u \in V - V(C)$. Thus $\phi(u) \in \mathrm{relint}(\mathrm{conv}(N(u)))$, so we also have $a^\top \phi(u) = \beta$. Let

$$A_1 := \{w \in V : a^\top \phi(w) > \beta\}$$
$$A_2 := \{w \in V : a^\top \phi(w) < \beta\}.$$

By Claim 1, these sets are nonempty, and by Claim 2, $G[A_1], G[A_2]$ are connected subgraphs. Let $A_3' \subseteq V - (A_1 \cup A_2)$ be the set of all vertices w such that $a^\top \phi(w') = \beta$ for all $w' \in N(w)$. That is, A_3' collects all the vertices w such that neither w nor its neighbours belong to $A_1 \cup A_2$. By definition, $u \in A_3'$. Let

$$A_3 := \text{the connected component of } G[A_3'] \text{ containing } u \subseteq A_3'.$$

Next, let $B := V - (A_1 \cup A_2 \cup A_3')$. That is, B collects vertices w such that $a^\top \phi(w) = \beta$ but $a^\top \phi(w') \neq \beta$ for some $w' \in N(w)$. In fact, local convexity implies that each vertex in B has a neighbour in A_1 as well as a neighbour in A_2. Moreover, B *cuts* the vertex set A_3 from the rest of the graph, that is, $G \setminus B$ has A_3 as a connected component. Subsequently, since G is 3-connected, B must contain at least three vertices b_1, b_2, b_3 each of which has a neighbour in A_3.

To summarise, we have found three vertex-disjoint connected subgraphs, $G[A_1], G[A_2], G[A_3]$, and three vertices outside, b_1, b_2, b_3, such that each b_i has a neighbour in each A_j. Consequently, G has a $K_{3,3}$ minor, a contradiction to Remark 21.3 as G is planar. $\Diamond$

Recall from Theorem 21.2 that every edge of G belongs to exactly two peripheral cycles, that G has a unique embedding in the plane, and that the peripheral cycles correspond precisely to the face boundaries of the unique embedding.

Claim 4. Let $e = uv \in E - C$, and let C_1, C_2 be the two peripheral cycles containing e. Then the line through $\phi(u), \phi(v)$ strictly separates $\{\phi(w) : w \in V(C_1) - \{u, v\}\}$ from $\{\phi(w') : w' \in V(C_2) - \{u, v\}\}$.

Proof of Claim. Suppose otherwise. Then there exist $a \in \mathbb{R}^2$ and $\beta \in \mathbb{R}$ such that $a^\top \phi(u) = a^\top \phi(v) = \beta$, and both $V(C_1) - \{u, v\}, V(C_2) - \{u, v\}$ intersect $\{w \in V : a^\top \phi(w) \geq \beta\}$. We shall construct paths P, Q on different sides of the line and are therefore vertex-disjoint, where P connects $V(C_1) - \{u, v\}$ to $V(C_2) - \{u, v\}$, and Q is a uv-path different from $\{e\}$, thereby contradicting the Key Lemma, Lemma 21.4.

To this end, let $X := \{w \in V : a^\top \phi(w) > \beta\}$ and $Y := \{w \in V : a^\top \phi(w) < \beta\}$. Since $e = uv \in E - C$, it follows (from Claim 1) that X, Y are nonempty. By Claim 2, $G[X], G[Y]$ are connected subgraphs. Nondegeneracy, together with local convexity, implies that each of u, v has a neighbour in Y. As a result, since $G[Y]$ is connected, there exists a uv-path Q different from $\{e\}$ whose internal vertices are in Y. The path P is constructed in a similar fashion on the X side. More precisely, let w_1 be a vertex of $V(C_1) - \{u, v\}$ such that $a^\top \phi(w_1) \geq \beta$, and let w_2 be a vertex of $V(C_2) - \{u, v\}$ such that $a^\top \phi(w_2) \geq \beta$. A similar argument as above tells us that there exists an $w_1 w_2$-path P whose internal vertices are in X. Since $\{w_1, w_2\} \cap \{u, v\} = \emptyset$ and $X \cap Y = \emptyset$, it follows that P, Q are the desired vertex-disjoint paths, thereby yielding a contradiction. $\diamondsuit$

Claim 5. Every peripheral cycle is mapped, under ϕ, to a strictly convex polygon.

Proof of Claim. This follows from Claim 4. $\diamondsuit$

Let S be the set of points in $\mathrm{conv}(\phi(V(C)))$ that do not lie on any line segment connecting the ϕ-images of adjacent vertices. For each point $p \in S$, let $n(p)$ be the number of peripheral cycles $C' \neq C$ such that $p \in \mathrm{relint}(\mathrm{conv}(\phi(V(C'))))$.

Claim 6. $n(p) = 1$ for each $p \in S$.

Proof of Claim. Consider the ϕ-image of C, a strictly convex polygon. Pick a non-corner point q on the boundary of this polygon. Observe that $n(p) = 1$ for every point $p \in S$ that is sufficiently close to q: $n(p) \geq 1$ follows from the fact that the edge of G containing q belongs to a peripheral cycle different from C, while $n(p) \leq 1$ follows from Claim 1. To prove the claim for every point in S, we shall leverage Claim 4.

Pick an arbitrary point $p \in S$. Pick a line ℓ through p that does not pass through any $\phi(v), v \in V$, and let q_1, q_2 be the intersection points of ℓ with the ϕ-image of C. Observe that $p \in \ell[q_1, q_2]$. What we just observed tells us $n(p') = 1$ for every $p' \in S \cap \ell[q_1, q_2]$ that is sufficiently close to q_1 or q_2. As we move in $S \cap \ell[q_1, q_2]$ from a neighbourhood of q_1 to a neighbourhood of q_2, the function $n(\cdot)$ does not change unless we "jump" over $(\phi(u), \phi(v))$ for some $uv \in E - C$. However, by Claim 4, $n(\cdot)$ remains the same whenever we jump over any $(\phi(u), \phi(v)), uv \in E - C$. Consequently, $n(\cdot)$ remains the same, 1, as we move in $S \cap \ell[q_1, q_2]$ from a neighbourhood of q_1 to a neighbourhood of q_2, implying in turn that $n(p) = 1$, as required. $\diamondsuit$

Claim 7. The ϕ-images of any two distinct edges of G are internally disjoint.

Proof of Claim. Exercise. $\diamondsuit$

Claim 7 finishes the proof. $\square$

Theorem 21.6. *Let $G = (V, E)$ be a 3-connected planar graph, let $C \subseteq E$ be a peripheral cycle, and assume that $\phi' : V(C) \to \mathbb{R}^2$ maps C to a strictly convex polygon. Then the barycentric extension of ϕ' yields a locally convex embedding of (G, C), and thus a straight-line plane embedding of G.*

Proof. That the barycentric extension of ϕ' yields a locally convex embedding of (G, C) is immediate. Once there, Theorem 21.5 tells us that this embedding must be a plane embedding of G. $\square$

Acknowledgements

Proofs of the preliminaries in §21.2 can be looked up, for example, in [1], Chapter 4 (peripheral cycles are referred to as *non-separating induced cycles*). The drawing method of §21 originally comes from Tutte [3]. Our presentation closely followed Geelen [2].

References

[1] R. Diestel. *Graph Theory, 5th Ed.* Springer-Verlag, Heidelberg, 2016/17.

[2] J. Geelen. On how to draw a graph. May 2012.

[3] W. T. Tutte. How to draw a graph. *Proc. London Math. Soc.*, 3(13):743–768, 1963.