MA431 Spectral Graph Theory: Lecture 5

Ahmad Abdi Neil Olver*

10 The cut and cycle spaces

Let $G = (V, E)$ be a **connected** undirected graph with n vertices and m edges, and let $\vec{G} = (V, \vec{E})$ be any orientation of G. The arbitrariness of the orientation may be a bit disconcerting, but everything we discuss will in fact be insensitive to the choice of orientation (just as was the case when we used it to give one definition of the Laplacian). We will often make use of the natural bijection between $\vec{E}$ and E without comment; so for example, we will generally treat $\mathbb{R}^E$ and $\mathbb{R}^{\vec{E}}$ as the same space. We will also define $\overleftrightarrow{G} = (V, \overleftrightarrow{E})$ to be the *bidirection* of G, where each edge is replaced by both orientations.

Throughout, we will use B to denote the vertex-edge incidence matrix associated with $\vec{G}$. That is, $B \in \mathbb{R}^{V \times \vec{E}} \simeq \mathbb{R}^{V \times E}$, with

$$B_{v,e} = \begin{cases} 1 & \text{if } v \text{ is the head of } e, \\ -1 & \text{if } v \text{ is the tail of } e, \\ 0 & \text{otherwise.} \end{cases}$$

A *flow* on G (or on $\vec{G}$; we won't distinguish) is simply any vector $f \in \mathbb{R}^E$ (or in $\mathbb{R}^{\vec{E}}$—again, we won't distinguish). Note that we do *not* require that $f_e \geq 0$ in a flow. While a positive flow on an edge $(u, v) \in \vec{E}$ should be interpreted as flow from u to v, a negative flow on this edge should be interpreted as flow in the reverse direction, from v to u. The purpose of an orientation is purely to indicate in which direction on an undirected edge a positive flow traverses.

The *net flow* into a node $v \in V$ induced by a flow $f \in \mathbb{R}^E$ is simply the total entering flow less the total leaving flow. We define ∇ to be the operator from $\mathbb{R}^E$ to $\mathbb{R}^V$ which maps a flow f to the vector b, where b_v is the net flow into v for each $v \in V$. That is,

$$(\nabla f)_v = \sum_{e=(u,v) \in \vec{E}} f_e - \sum_{e=(v,w) \in \vec{E}} f_e.$$

In the standard basis, the matrix representing ∇ is simply B. We will often write ∇f_v in place of $(\nabla f)_v$, as there can be no confusion.

A *circulation* is simply a flow f with $\nabla f_v = 0$ for all $v \in V$. More generally, given a *demand vector* $b \in \mathbb{R}^V$ with $\sum_{v \in V} b_v = 0$, we might be interested in flows which correctly match this demand vector, i.e., which satisfy $\nabla f = b$.

We will need some notation. Define, for any edge $e \in \overleftrightarrow{E}$, the vector $\chi^e \in \mathbb{R}^E$ by

$$(\chi^e)_a = \begin{cases} 1 & \text{if } a = e, \\ -1 & \text{if the reverse of } a \text{ is } e, \\ 0 & \text{otherwise.} \end{cases}$$

This is a *signed* characteristic vector of the edge e. For any $F \subseteq \overleftrightarrow{E}$, define $\chi(F) := \sum_{e \in F} \chi^e$. Given a set $S \subseteq V$, with $S \neq \emptyset$ and $S \neq V$, the *cut* associated with S, denoted by $\delta^+(S)$, is the set of arcs $(v, w) \in \overleftrightarrow{E}$ with $v \in S$ and $w \notin S$.

We now define two subspaces of $\mathbb{R}^E$. The subspaces do depend on the choice of orientation, but again, not in any important way.

The *cycle space* $W^\diamond$ is the set of all circulations in G, that is,

$$W^\diamond := \{f \in \mathbb{R}^E : \nabla f = \mathbf{0}\}.$$

We define the *cut space* (sometimes called the *star space*) as simply the orthogonal complement of $W^\diamond$, and denote it by $W^\star$: so $W^\star = (W^\diamond)^\perp$, and $W^\diamond \oplus W^\star = \mathbb{R}^E$. The reason for the names will become clear soon.

First, let's see an alternative description of $W^\diamond$, as well as a description of one possible basis. A reminder that given a spanning tree T[1] of G, every edge e not in T has an associated *fundamental cycle*, the cycle consisting of e along with the path in T between the endpoints of e. We will consider this as a directed cycle in $\overleftrightarrow{G}$, oriented so that e is included in the forward direction. There is also the *fundamental cut* associated with any edge $e \in T$: removing e from T splits it into two connected components, partitioning V into $S_e, V \setminus S_e$, where S_e contains the tail of e; the fundamental cut is $\delta^+(S_e)$.

Lemma 10.1. *The following statements about $W^\diamond$ hold.*

1. $W^\diamond = \operatorname{span}(\{\chi(C) : C$ *is a directed cycle in* $\overleftrightarrow{G}\})$.
2. *Given any spanning tree T, and taking C_e to be the fundamental cycle associated with e for each $e \notin T$, $\{\chi(C_e) : e \in E \setminus T\}$ is a basis for $W^\diamond$.*
3. $\dim(W^\diamond) = m - n + 1$.

Proof. For any directed cycle C in $\overleftrightarrow{G}$, $\nabla(\chi(C)) = 0$, and hence $\chi(C) \in W^\diamond$. Since C_e does not contain e' for any $e \neq e' \in E \setminus T$, $\{\chi(C_e) : e \in E \setminus T\}$ is certainly linearly independent. It remains to show that $\dim(W^\diamond) = m - n + 1$; since a spanning tree has $n - 1$ edges, it then follows that $\{\chi(C_e) : e \in E \setminus T\}$ is indeed a basis, and so all parts of the lemma follow.

$W^\diamond$ is the kernel of the ∇ operator, and so it suffices to show that the rank of this operator, or equivalently the rank of B, is $n - 1$. We show this by demonstrating that the kernel of $B^\top$ is 1-dimensional. If $B^\top \alpha = 0$, we must have that $\alpha_u = \alpha_v$ for every $\{u, v\} \in E$, and hence by the connectivity of G, α is a multiple of $\mathbf{1}$. Further, clearly $B^\top \mathbf{1} = 0$. □

[1]We will not distinguish between a spanning tree and its set of edges.

Now we move on to the cut space, which will in fact be the more important subspace for us. Both names (cut space as well star space) should become clear after this lemma.

Lemma 10.2. *The following statements about $W^\star$ hold.*

1. $W^\star = \operatorname{span}(\{\chi(\delta^+(S)) : \emptyset \subsetneq S \subsetneq V\})$.

2. $W^\star = \operatorname{span}(\{\chi(\delta^+(\{r\})) : r \in V\})$.

3. *Given any spanning tree T, and taking $\delta^+(S_e)$ to be the fundamental cut associated with e for each $e \in T$, $\{\chi(\delta^+(S_e)) : e \in T\}$ is a basis for $W^\star$.*

4. $\{\chi(\delta^+(\{r\})) : r \in V \setminus \{t\}\}$ *is a basis, for any choice of $t \in V$.*

5. $\dim(W^\star) = n - 1$.

Proof. Exercise. □

The grad operator (the *gradient*) is a linear map from $\mathbb{R}^V$ to $\mathbb{R}^E$; for any $\pi \in \mathbb{R}^V$, $f = \operatorname{grad} \pi$ is defined by $f_e = \pi_w - \pi_v$ for every $e = (v, w) \in \vec{E}$. The matrix of this linear operator in the standard basis is simply $B^\top$. As such, ∇ and grad are adjoint operators: $\langle \nabla f, \pi \rangle = \langle f, \operatorname{grad} \pi \rangle$.

Lemma 10.3. *The range of* grad *is precisely $W^\star$.*

Proof. $W^\diamond$ is the kernel of the ∇ operator. Thus $W^\star = (W^\diamond)^\perp$ is the range of the adjoint operator grad. □

We have the following relation to the Laplacian of G, which we will denote by L throughout.

Lemma 10.4. *The linear operator which is represented by L in the standard basis is precisely* $\nabla \operatorname{grad}$.

Proof. In the standard basis, $\nabla \operatorname{grad} = BB^\top$, which we already saw was one possible definition of the Laplacian. □

Remark 10.5. The Laplacian is also used to describe a differential operator. Given an appropriately smooth real-valued function $g : \mathbb{R}^k \to \mathbb{R}$, the Laplacian of g is written as $\nabla^2 g$; it is a map from $\mathbb{R}^k$ to $\mathbb{R}$ defined by

$$\nabla^2 g = \sum_{i=1}^{k} \frac{\partial^2 g}{\partial x_i^2}.$$

It can be viewed as first computing the gradient of g (often denoted by ∇g), which is

$$h = \nabla g = \left(\frac{\partial g}{\partial x_1}, \frac{\partial g}{\partial x_2}, \ldots, \frac{\partial g}{\partial x_k} \right),$$

and then taking the divergence of h (often denoted by $\nabla \cdot h$), which is

$$\nabla \cdot h = \sum_{i=1}^{k} \frac{\partial h_i}{\partial x_i}.$$

In our discrete setting, we have exactly the same description of the Laplacian, as the divergence of the gradient[2]. The analogy is particularly clear if you consider the graph G describing a grid in k dimensions.

10.1 Projection onto the cut space

Let $P_\star : \mathbb{R}^E \to \mathbb{R}^E$ denote the linear operator that projects orthogonally onto $W^\star$. We will study this operator in some detail.

Let Π be the $m \times m$ matrix that represents $P_\star$ in the basis $\{\chi^e : e \in \vec{E}\}$. That is, $\Pi_{e,e'}$ is the amount of flow on arc e in the flow $P_\star \chi^{e'}$ obtained from the unit flow across e'. Since $P_\star$ is an orthogonal projection and hence self-adjoint, Π is symmetric. We will mostly not need to refer to a specific basis, but we will use this on occasions (and more later).

So let $f \in \mathbb{R}^E$ be some flow, and let $i = P_\star f$. We first make the observation that $\nabla f = \nabla i$, since $f - i \in W^\diamond$. Moreover, i depends only on ∇f; if $\nabla f = \nabla g$, then $P_\star f = P_\star g$.

Now suppose f satisfies $\nabla f = \mathrm{e}_t - \mathrm{e}_s$, for some $s, t \in V$. (In particular, if $(s,t) \in \vec{E}$, we can take $f = \chi^{(s,t)}$.) So:

- i sends one unit of flow from v to w, and satisfies flow conservation (flow in equals flow out) everywhere else.
- By Lemma 10.3, we have that $i = \operatorname{grad} \pi$ for some $\pi \in \mathbb{R}^V$. In other words, $i_e = \pi_w - \pi_v$ for all $e = (v,w) \in \vec{E}$.

What we have described are nothing more than Kirchhoff's famous laws (as well as Ohm's law) for electrical flow. We think of edges of our graph as being resistors, all with unit resistance. The values π represent *electrical potentials*[3] corresponding to the electrical flow i; flow on an edge is precisely given by the potential difference. In the setup described, we have a unit current from s to t; physically speaking, we can imagine that we have a battery set up that is connected to s and to t. In the scaling we have chosen, we have fixed the current to be unit, with the voltage difference between s and t not determined. That is, $\pi_t - \pi_s$ would be the voltage of the battery. Of course, we could scale differently: if we consider $i' = i/(\pi_t - \pi_s)$, then the potentials associated with i' are $\pi' = \pi/(\pi_t - \pi_s)$, and so $\pi'_t - \pi'_s = 1$. So i' would be the electrical flow associated with a battery of unit voltage.

So $P_\star$ is the operator that takes a flow f, and returns the electrical flow that "corresponds" to it, in the sense of having the same net flow as f everywhere. We can also now read off immediately the well-known fact that electrical flows have minimum *energy*: the energy of a flow is just its squared norm.

Lemma 10.6. *For any flow $f \in \mathbb{R}^E$, $i = P_\star f$ is the unique minimizer of $\|g\|$ amongst all flows g satisfying $\nabla g = \nabla f$ and $g = P_\star f$.*

[2] We cannot get away with overloading ∇ for both the gradient and the divergence in the discrete setting without causing unnecessary confusion, so we use it only for the divergence.

[3] Actually, the negations of electrical potentials—the usual convention is that electrical current flows "downhill" with respect to the potential. Nevertheless we'll just call them potentials.

Proof. Let $i = P_\star f$. For any g with $\nabla g = \nabla f$, $g - i \in W^\diamond$, and so $\langle i, g - i\rangle = 0$. We deduce that $\|g\|^2 = \|i\|^2 + \|g - i\|^2$, and the claim is immediate. □

Definition 10.7. The *effective resistance* of an edge $e = (s,t) \in \vec{E}$ is the potential difference $\pi_t - \pi_s$ of a potential π corresponding to the electrical flow $i = P_\star \chi^e$.

The name is motivated by imagining replacing the graph with a single resistor between the endpoints of e, and setting the resistance of this resistor equal to the effective resistance. Then the potential difference corresponding to a unit current through this resistor is the same as in the original circuit. There are a number of equivalent formulations of the effective resistance (the final one requires a small adjustments in the weighted case that we will discuss later).

- It is the inverse of the total amount of current that flows between the endpoints of e if a unit voltage battery is applied across e (just consider the rescaling $i/(\pi_t - \pi_s)$.
- It is the energy of i. To see this, observe that

$$\langle i, \chi^e\rangle = \langle P_\star i, \chi^e\rangle = \langle i, P_\star \chi^e\rangle = \|i\|^2.$$

- It is the current i_e (since $i_e = \pi_t - \pi_s$), and hence also the entry $\Pi_{e,e}$ of the matrix associated with $P_\star$.

We can also define the effective resistance between any two nodes $u, v \in V$, even if they are not connected by an edge. Take i to be the electrical flow with $\nabla i = \mathrm{e}_v - \mathrm{e}_u$, and let π be associated potentials; the effective resistance between u and v is then simply $\pi_v - \pi_u$.

The effective resistance turns out to be a very important quantity. Edges with large effective resistance can be considered "more important" than ones with small effective resistance: for example, if an edge is a cut in the graph (removing the edge disconnects it), it's effective resistance will be 1, which is as large as possible (why?). We will see some applications later.

It will be very useful to have a concrete description of $P_\star$. How can we actually *compute* $i = P_\star f$? We proceed as follows. First, compute $b = \nabla f$; as already noted, this is the only information we need about f. Let π be a potential associated with the current i we are trying to compute; it suffices to find π, since then $i = \operatorname{grad} \pi$ follows immediately. So given b, we would like to compute π. To see how to do this, let's consider the much easier question of going from π to b: we have that $b = \nabla i = \nabla \operatorname{grad} \pi = L\pi$, by Lemma 10.4.

So to get π from b, we simply invert: $\pi = L^+ b$. (Notice, though, that any π that solves $L\pi = b$ would also work; this would differ from $L^+ b$ by a multiple of the all-ones vector, and shifting all potentials by a constant makes no difference.) Putting this all together,

$$P_\star f = \operatorname{grad} \pi = \operatorname{grad} L^+ b = \operatorname{grad} L^+ \nabla f.$$

In other words, we can write the operator $P_\star$ as $\operatorname{grad} L^+ \nabla$. Or written in terms of the standard basis, we have $\Pi = B^\top L^+ B$. (For maximal possible confusion, this can be further expanded as $\Pi = B^\top (BB^\top)^+ B \ldots$)

Note that this also gives us another way to write the effective resistance of an edge $e = (s,t)$. It is

$$\Pi_{e,e} = (B^\top L^+ B)_{e,e} = (\mathrm{e}_t - \mathrm{e}_s)^\top L^+ (\mathrm{e}_t - \mathrm{e}_s). \tag{1}$$

It is also worth interpreting the equation $L\pi = b$ in the case $b = \mathrm{e}_t - \mathrm{e}_s$. This tells us in particular that π is *harmonic* on $V \setminus \{s,t\}$: for any $v \neq s,t$, we have that

$$\pi_v = \frac{1}{\deg(v)} \sum_{w:\{v,w\}\in E} \pi_w.$$

Remark 10.8. There is a very well-developed theory about solving linear Laplacian systems. That is, given a matrix L that is the Laplacian of some graph, and some vector b, solve the system $Lx = b$. There are algorithms that solve (with some small error that can be specified) these systems in "near-linear time", meaning time $O(m \operatorname{polylog} m)$, which is basically the best you could hope for; $O(m)$ time is needed just to read the graph.

The above shows that solving such a system fast is equivalent to implementing $P_\star$ fast—and this is indeed crucially exploited in these fast algorithms. We won't develop the theory in this course; more information can be found in a monograph by Vishnoi [2].

The following is called the *Rayleigh monotonicity principle*. It says that adding edges to the graph can only decrease the effective resistance.

Theorem 10.9. *Let $G' = (V, E')$ with $E' \supseteq E$. Then for any two nodes $u, v \in V$, the effective resistance between u and v in G' is smaller (or equal to) its value in G.*

Proof. Exercise. □

11 Kirchhoff's effective resistance theorem

So we now know how to compute the effective resistance of an edge, by (1); we simply need to solve a linear system. But let's do something that's usually not very useful, and obtain a "formula" for the effective resistance via Cramer's rule.

Assume vertices are labelled so that the edge of interest is $e = (n, n-1)$; we wish to find some π for which $L\pi = \mathrm{e}_{n-1} - \mathrm{e}_n$, and then the desired effective resistance is $\pi_{n-1} - \pi_n$. Since $\mathbf{1}$ is in the kernel of L, we can restrict our attention to solutions with $\pi_n = 0$ (note that this means we are not choosing $\pi = L^+(\mathrm{e}_{n-1} - \mathrm{e}_n)$; our choice will differ by a shift). We can also drop the last row of the system, since it is a linear combination of the other constraints. Thus, we can solve the system $(L[n])y = \mathrm{e}_{n-1}$; a solution y represents the restriction of π to $\mathbb{R}^{n-1}$, and y_{n-1} is the effective resistance.

We now apply Cramer's rule to the now invertible system. Let L' be obtained from $L[n]$ by replacing the $(n-1)$'th column by e_{n-1}. Then

$$y_{n-1} = \frac{\det(L')}{\det(L[n])} = \frac{\det(L[n][n-1])}{\det(L[n])}.$$

The last equality comes from expanding the determinant of L' on the final column.

But now we can apply the matrix tree theorem. Recall that $\det(L[n])$ is the number of spanning trees of G; we also saw in the proof of the matrix tree theorem that $L[n][n-1]$ is a $(n-2) \times (n-2)$ principal submatrix of the Laplacian of G/e, and hence that $\det(L[n][n-1])$ is the number of spanning trees of G/e. Thus, the effective resistance of e is simply the fraction of spanning trees that use edge e.

Let $\mathcal{T}$ denote the set of all spanning trees of G. A *uniformly random spanning tree* is simply a spanning tree chosen uniformly at random from $\mathcal{T}$. Properties of uniformly random spanning trees and related objects are objects of intense interest to probabilists; we'll only scratch the surface. We have shown

Theorem 11.1. *If T is a uniformly random spanning tree of G, then $\Pr(e \in T) = \Pi_{e,e}$ for any edge $e \in E$.*

There is something a bit opaque about this proof. Both spanning trees and electrical flows are somewhat "combinatorial" objects, satisfying combinatorial or linear constraints. The above proof shows the correspondence via a detour through the algebraic world of determinants, which is difficult to interpret. In the next lecture, we will give a second, much more combinatorial proof of Kirchhoff's effective resistance theorem.

Acknowledgements

The book by Lyons and Peres [1] is an excellent source for a lot of this, from the more probabilistic viewpoint.

References

[1] R. Lyons and Y. Peres. *Probability on Trees and Networks*, volume 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016. Available at `http://rdlyons.pages.iu.edu/`.

[2] N. Vishnoi. $Lx = b$. In *Foundations and Trends in Theoretical Computer Science*, volume 8, pages 1–141. now, 2013.