Cheeger constant and conductance

The Cheeger constant of a graph is a measure of its "bottleneckedness". Intuitively, if a graph is well-connected, it should not be possible to cut off a "big" piece by removing a small number of edges. We can interpret "big" either as number of vertices or number of edges. These two views correspond to eigenvalues of the Laplacian and normalized Laplacian respectively.
For a subset $S \subseteq V(G)$, we let the *boundary* be

$$\delta S = \{uv \in E(G) \mid u \in S,\ v \in \overline{S}\}.$$

If we interpret the size of S as number of vertices, we let

$$\theta(S) = \frac{|\delta S|}{|S|}.$$

Then, the Cheeger constant of a graph G is defined as the smallest such cut:

$$\theta_G = \min_{S \subseteq V(G)} \frac{|\delta S|}{\min(|S|, |\overline{S}|)}.$$

If, instead, we are interested in the size of S in terms of edges, we let

$$\phi(S) = \frac{|\delta S|}{\text{vol } S},$$

where

$$\text{vol } S = \sum_{v \in S} d_v.$$

Then, the Cheeger constant, or sometimes called *conductance* in this case, of a graph G is defined as the smallest such cut:

$$\phi_G = \min_{S \subseteq V(G)} \frac{|\delta S|}{\min(\text{vol } S,\ \text{vol } \overline{S})}.$$

Recall that in both the Laplacian and the normalized Laplacian, the smallest eigenvalue is 0, and the multiplicity of 0 is equal to the number of connected components. This implies that the second-smallest eigenvalue is nonzero if and only if the graph is connected. These second-smallest eigenvalues are closely related to the two Cheeger constants. Let $0 = \lambda_0 \leq \lambda_1 \leq \cdots \leq \lambda_{n-1}$ be the eigenvalues of the Laplacian, and $0 = \mu_0 \leq \mu_1 \leq \cdots \leq \mu_{n-1}$ be the eigenvalues of the normalized Laplacian.

Theorem 1. *We have*

$$\frac{\theta_G^2}{2\Delta(G)} \leq \lambda_1 \leq 2\theta_G,$$

and

$$\frac{\phi_G^2}{2} \leq \mu_1 \leq 2\phi_G.$$

In class, we discussed the set-up for the proofs of the lower bounds (on *theta* and *phi*).

Exercise 1. *Prove that $\mu_1 \leq 2\phi_G$.* 3 points

Exercise 2. *Prove that* $\lambda_1 \leq 2\theta_G$. 3 points

Exercise 3. *Find* θ_{Q_n} *and* ϕ_{Q_n}*, where* Q_n *is the* n*-dimensional hypercube.* 1 point

Exercise 4. *Find* θ_{K_n} *and* ϕ_{K_n}*, where* K_n *is the complete graph on* n *vertices.* 2 points

For the upper bound side of the theorem, let's assume that G is a d-regular graph. In this case, the proofs that $\theta(G) \leq \sqrt{2\Delta\lambda_1}$ and $\phi(G) \leq \sqrt{2\mu_1}$ are very similar, because θ and ϕ are equal up to the constant d, and we'll work with $\phi(G)$. We let

$$R(\vec{x}) = \frac{\vec{x}^T \mathcal{L} \vec{x}}{d\vec{x}^T\vec{x}} = \frac{\sum_{uv\in E(G)}(x_u - x_v)^2}{d\sum_{v\in V(G)} x_v^2}.$$

Now, the goal is to show for any $\vec{x} \perp \vec{1}$, there exists an $S \subseteq V(G)$ such that $\phi(S) = O(\sqrt{(}R(\vec{x})))$. We do this with the following two lemmas.

Lemma 2. *For any* $\vec{x} \perp \vec{1}$ *there exist two nonnegative vectors* $\vec{y}$ *and* $\vec{z}$ *on disjoint support such that* $R(\vec{y}), R(\vec{z}) \leq 4R(\vec{x})$.

Exercise 5. *Prove Lemma 2. Let* 2 points

$$y_i = \begin{cases} x_i, & x_i > 0, \\ 0, & \textit{otherwise}, \end{cases} \qquad z_i = \begin{cases} -x_i, & x_i < 0, \\ 0, & \textit{otherwise}. \end{cases}$$

Now, show that

$$\vec{y}^T\mathcal{L}\vec{y}, \vec{z}^T\mathcal{L}\vec{z} \leq \vec{x}^T\mathcal{L}\vec{x}.$$

Then, show that $\vec{y}^T\vec{y} + \vec{z}^T\vec{z} = \vec{x}^T\vec{x}$. *For the proof to work, we need* $\vec{y}^T\vec{y}, \vec{z}^T\vec{z} \geq 1/4$. *If this is not the case for one of them, consider a shift* $\vec{x} + c \times \vec{1}$.

Lemma 3. *For any nonnegative* $\vec{y}$ *there is a set* S *in the support of* $\vec{y}$ *such that* $\phi(S) = O(\sqrt{(}R(\vec{y})))$.

Exercise 6. *Prove Lemma 3. We will use the probabilistic method. You should try this proof in particular if you have some experience with probability. (Otherwise this might be tricky, but I am happy to talk you through it.) Consider the vector* $\vec{y}'$ *given by* $y_i' = y_i^2/\max_j(y_j)$. *Let* t *be a continuous random variable chosen uniformly from the interval* $[0, 1]$. *Let* $S = \{v : y_v' > t\}$. *Show that* 3 points

$$\frac{\mathbb{E}([S, \overline{S}])}{d\mathbb{E}(|S|)} = \frac{\sum_{uv\in E(G)} |x_u - x_v|}{d\sum_{v\in V(G)} x_v}.$$

Show that for nonnegative random variables A *and* B *it holds that* $\mathbb{E}(A)/\mathbb{E}(B) \leq c$ *implies that with nonzero probability* $A/B \leq c$. *Complete the proof from there.*

Finding the second eigenvector $\vec{x}_1$ of L or $\mathcal{L}$ cannot be done in linear time, but there are algorithms that find a vector $\vec{x}_1'$ with $R(\vec{x}_1') = O(R(\vec{x}_1))$ in linear time. From the above results, we see that this can be used to efficiently find small cuts S in networks.