CSC 2421H : Graphs, Matrices, and Optimization	Lecture 5 : 15 Oct 2018
Lazy Random Walks & Expanders	
Lecturer: Sushant Sachdeva	*Scribe: Sara Tang*

1 Finding ν_2 of the Dumbbell Graph

Definition 1.1. *Define a* ***dumbbell graph*** *of size* $2n$ *to be two complete graphs* K_n *connected by a single edge:*

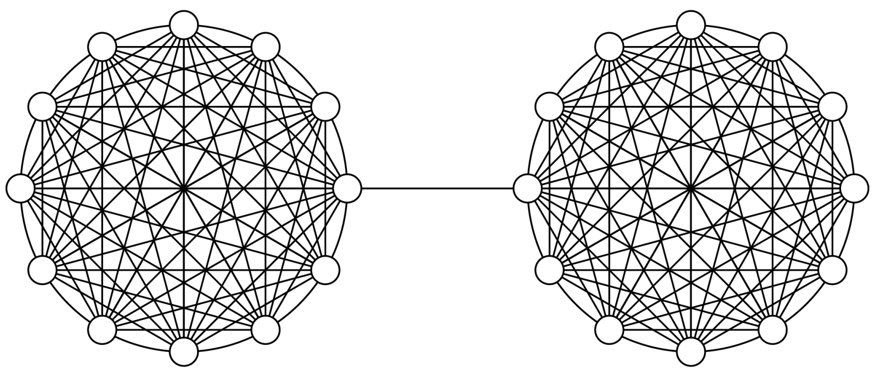

Figure 1: An example of a dumbbell graph.

It is not hard to see that $\Phi = \Theta(1/n^2)$, where $|V| = 2n$: simply cut along the edge connecting the complete graphs. By Cheeger's inequality, it follows that $\nu_2 \leq \Theta(1/n^2)$.

We wish to show that this is a tight bound, ie. $\nu_2 = \Omega(1/n^2)$. Intuitively, this is true, because we know from the main result of last week that the number of lazy random walk steps t it takes to be close to the stationary distribution is directly proportional to $1/\nu_2$.

Indeed, there is probability $1/n$ to get to the spectral vertex (the vertex connecting the two complete graphs), and from the spectral vertex, there is probability $1/n$ chance to cross to the other K_n. But let's do this rigorously.

Lemma 1.2. *Let* $G = (V, E)$*,* $u, v \in V$*. Suppose there exists a path* $P(u, v)$ *from* u *to* v *with length* 3*, as demonstrated in Figure 2. (Note that* (u, v) *need not be an edge in* G*.) Let* $L_{u,v}$ *be the Laplacian of the* (u, v)*-line graph, and* $L_{P(u,v)}$ *be the Laplacian of the* $P(u, v)$*-path graph.*

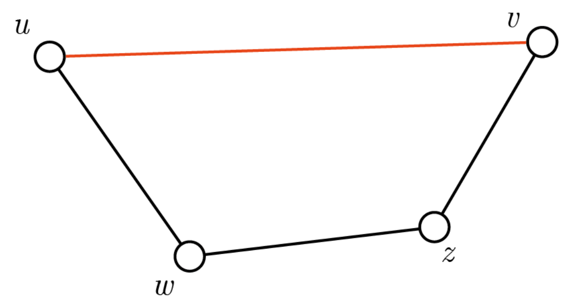

Figure 2: $P(u,v)$ is a path of length 3 via (u,w,z,v).

Then

$$(\mathbb{1}_u - \mathbb{1}_v)(\mathbb{1}_u - \mathbb{1}_v)^\top = L_{u,v} \preceq 3L_{P(u,v)}$$

Proof. By Cauchy-Schwarz,

$$\begin{aligned}
((x_u - x_w) + (x_w - x_z) + (x_z - x_v))^2 &\leq 3((x_u - x_w)^2 + (x_w - x_z)^2 + (x_z - x_v)^2) && \forall x \\
(x_u - x_v)^2 &\leq 3((x_u - x_w)^2 + (x_w - x_z)^2 + (x_z - x_v)^2) && \forall x \\
L_{u,v} &\preceq 3(L_{u,w} + L_{w,z} + L_{z,v}) = 3L_{P(u,v)}
\end{aligned}$$

□

The following is a useful lemma that establishes a lower bound on ν_2 given λ_2 (and the maximum degree of a graph, $d_{\max}$).

Lemma 1.3.

$$\nu_2 \geq \frac{\lambda_2}{d_{\max}}.$$

Proof. Recall by Courant-Fischer,

$$\nu_2 = \min_{\mathbf{x}^\top(\mathbf{D}^{1/2}\mathbb{1})=0} \frac{\mathbf{x}^\top \mathbf{N} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}} = \min_{\substack{T \subset \mathbb{R}^V \\ \dim(T)=2}} \max_{x \in T} \frac{\mathbf{x}^\top \mathbf{N} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}}$$

We can substitute $x = D^{1/2}y$. Since $x \in T$, it follows that $y \in \mathbf{D}^{1/2}T$:

$$= \min_{\substack{T \subset \mathbb{R}^V \\ \dim(T)=2}} \max_{y \in \mathbf{D}^{1/2}T} \frac{\mathbf{y}^\top \mathbf{L} \mathbf{y}}{\mathbf{y}^\top \mathbf{D} \mathbf{y}}$$

Let $\hat{T} = \mathbf{D}^{1/2}T$. Note that we can assume that $\mathbf{D}$ is one-to-one and onto, and hence $\dim(\hat{T}) = 2$. Otherwise, one of the vertices has degree zero, and hence $\lambda_2 = 0$. Thus,

$$
\begin{aligned}
&= \min_{\substack{\hat{T} \subset \mathbb{R}^V \\ \dim(\hat{T})=2}} \max_{y \in \hat{T}} \frac{\mathbf{y}^\top \mathbf{L} \mathbf{y}}{\mathbf{y}^\top \mathbf{D} \mathbf{y}} \\
&\geq \min_{\substack{\hat{T} \subset \mathbb{R}^V \\ \dim(\hat{T})=2}} \max_{y \in \hat{T}} \frac{\mathbf{y}^\top \mathbf{L} \mathbf{y}}{d_{\max} \mathbf{y}^\top \mathbf{y}} \\
&= \frac{1}{d_{\max}} \min_{\substack{\hat{T} \subset \mathbb{R}^V \\ \dim(\hat{T})=2}} \max_{y \in \hat{T}} \frac{\mathbf{y}^\top \mathbf{L} \mathbf{y}}{\mathbf{y}^\top \mathbf{y}} \\
&= \frac{\lambda_2}{d_{\max}}
\end{aligned}
$$

□

Theorem 1.4. *$\nu_2 = \Omega(1/n^2)$ for a dumbbell graph.*

Proof. We begin by finding a bound on $\lambda_2(L_G)$, then apply the second lemma to establish the relationship between λ_2 and ν_2.

Observe that for any $u, v \in V$, there is a path $P(u, v)$ from u to v of length 3. Thus, we can apply our first lemma:

$$L_{u,v} \preceq 3L_{P(u,v)} \preceq 3L_G$$

where the second $\preceq$ follows because we are simply adding more squares to $L_{P(u,v)}$. By summing over all pairs $(u, v) \in V \times V$, $u \neq v$,

$$L_{K_{2n}} \preceq 3\binom{2n}{2} L_G$$

By a corollary of Courant-Fischer, we know that

$$\lambda_2(L_{K_{2n}}) \leq \lambda_2\left(3\binom{2n}{2} L_G\right) = 3\binom{2n}{2}\lambda(L_G)$$

Recall $\lambda_2(L_{K_n}) = n$, so $\lambda_2(L_{K_{2n}}) = 2n$. Furthermore, $3\binom{2n}{2} = \Theta(n^2)$. Thus,

$$\lambda_2(L_G) = \Omega\left(\frac{1}{n}\right)$$

Applying the second lemma, and the fact that $d_{\max} = n$ for the dumbbell graph,

$$\nu_2(L_G) \geq \frac{\lambda_2(L_G)}{n} = \Omega\left(\frac{1}{n^2}\right)$$

□

2 Expanders

To motivate the following section, we notice that it would be nice to have graphs that approximate cliques, in the sense that these graphs quickly approach their stationary distribution but that are also sparser than cliques.

Definition 2.1. *A d-regular, unweighted graph G is said to be an ε-**expander** if*

$$-\varepsilon d \le \lambda_i(A) \le \varepsilon d \qquad \forall i < n.$$

Equivalently, because $L = dI - A$ and $N = \frac{1}{d}L$,

$$\begin{aligned} (1-\varepsilon)d \le \lambda_j(L) &\le (1+\varepsilon)d && \forall j \ne 1 \\ 1-\varepsilon \le \lambda_j(N) &\le 1+\varepsilon && \forall j \ne 1 \\ (1-\varepsilon)L_{K_n} \preceq \frac{n}{d}L_G &\preceq (1+\varepsilon)L_{K_n} \end{aligned}$$

The last equivalence is not trivial, so we will demonstrate it below:

Lemma 2.2. *G is a d-regular ε-expander iff*

$$(1-\varepsilon)L_{K_n} \preceq \frac{n}{d}L_G \preceq (1+\varepsilon)L_{K_n}$$

Proof. This is equivalent to showing

$$(1-\varepsilon)\mathbf{x}^\top \mathbf{L}_{K_n}\mathbf{x} \le \frac{n}{d}\mathbf{x}^\top \mathbf{L}_G \mathbf{x} \le (1+\varepsilon)\mathbf{x}^\top \mathbf{L}_{K_n}\mathbf{x} \qquad \forall \mathbf{x}$$

Any $\mathbf{x}$ can be expressed as the linear combination of the eigenvectors of $\mathbf{L}_{K_n}$. Thus, $\mathbf{x} = c\mathbb{1} + \mathbf{y}$, where $\mathbf{y}^\top \mathbb{1} = 0$ and $\mathbf{y}$ has eigenvalue n. Restrict the inequality we are trying to prove to the space orthogonal to span$\{\mathbb{1}\}$:

$$\begin{aligned} (1-\varepsilon)n\mathbf{y}^\top \mathbf{y} \le \frac{n}{d}\mathbf{y}^\top \mathbf{L_G}\mathbf{y} &\le (1+\varepsilon)n\mathbf{y}^\top \mathbf{L_{K_n}}\mathbf{y} && \forall y \\ (1-\varepsilon)d\mathbf{y}^\top \mathbf{y} \le \mathbf{y}^\top \mathbf{L_G}\mathbf{y} &\le (1+\varepsilon)d\mathbf{y}^\top \mathbf{L_{K_n}}\mathbf{y} && \forall y \\ (1-\varepsilon)d \le \frac{\mathbf{y}^\top \mathbf{L_G}\mathbf{y}}{\mathbf{y}^\top \mathbf{y}} &\le (1+\varepsilon)d && \forall y \end{aligned}$$

which is true by our second definition of the d-regular ε-expander. □

In particular d-regular ε-expanders are nice because it would otherwise be difficult for such graphs to arise randomly:

Let $\mathcal{G}(n,p)$ be a random graph model with n vertices such that each edge (u,v) occurs independently with probability p.

Then $\mathbb{E}[\text{no. of edges}] = p\binom{n}{2}$ and $\mathbb{E}[\deg v] = (n-1)p$. In other words, if we want d_v to be constant (to mimic a d-regular graph), we would only need $p \sim \frac{1}{n}$. However, for the graph to also be connected we would have to take $p \sim \frac{\log n}{n}$. Therefore, a connected random graph will most likely be more dense than its d-regular counterpart.

Theorem 2.3. *$\forall \varepsilon > 0$. $\exists d(\varepsilon)$ such that there is a "family" of d-regular ε-expanders, ie. an increasingly-sized collection of ε-expanders.*

Corollary 2.4. *Let G be a d-regular ε-expander. Then $\nu_2(G) \geq 1 - \varepsilon = \Omega(1)$ implies we only need $\Theta(\log n)$ steps to get close to stationary, and we only need $\log d$ bits of randomness to choose the next vertex.*

Compare this to K_n, which also takes $\Theta(\log n)$ steps to get close to stationary, but would need $\log n$ bits of randomness to choose the next vertex.