Lectures 13: Random walks on graphs and electrical networks CSE 525, Fall 2016
Instructor: James R. Lee

1 Random walks

Let $G = (V, E)$ be an undirected graph. The random walk on G is a Markov chain on V that, at each time step, moves to a uniformly random neighbor of the current vertex.

Ffsor $x \in V$, use d_x to denote the degree of vertex x. Then more formally, *random walk on G* is the following process $\{X_t\}$. We start at at some node $X_0 = v_0 \in V$. Then if $X_t = v$, we put $X_{t+1} = w$ with probability $1/d_v$ for every neighbor w of v.

1.1 Hitting times and cover times

One can study many natural properties of the random walk. For two vertices $u, v \in V$, we define the *hitting time H_{uv} from u to v* as the expected number of steps for the random walk to hit v when started at u. Formally, define the random variable $T = \min\{t \geqslant 0 : X_t = v\}$. Then $H_{uv} = \mathbb{E}[T \mid X_0 = u]$.

The *cover time of G starting from u* is the quantity $\mathrm{cov}_u(G)$ which is the expected number of steps needed to visit every vertex of G started at u. Again, we can define this formally: Let $T = \min\{t \geqslant 0 : \{X_0, X_1, \ldots, X_t\} = V\}$. Then $\mathrm{cov}_u(G) = \mathbb{E}[T \mid X_0 = u]$. Finally, we define the *cover time of G* as $\mathrm{cov}(G) = \max_{u \in V} \mathrm{cov}_u(G)$.

1.2 Random walks and electrical networks

It turns out that random walks (on undirected graphs) are very closely related to electrical networks. We recall the basics of such networks now. Again, we let $G = (V, E)$ be a connected, undirected graph which we think of as an electrical circuit with unit resistors on every edge.

If we create a potential difference at two vertices (by, say, connecting the positive and negative terminals of a battery), then we induce an electrical flow in the graph. Between every two nodes u, v there is a *potential* $\phi_{u,v} \in \mathbb{R}$. Electrical networks satisfying the following three laws.

(K1) The flow into every node equals the flow out.

(K2) For sum of the potential differences around any cycle is equal to zero.

(Ohm) The current flowing from u to v on an edge $e = \{u, v\}$ is precisely $\frac{\phi_{u,v}}{r_{uv}}$ where r_{uv} is the resistance of $\{u, v\}$. [In other words, $V = iR$.]

In our setting, all resistances are equal to one, but one can define things more generally. [If we put conductances c_{uv} on the edges $\{u, v\} \in E$, then the corresponding random walk would operate as follows: If $X_t = u$ then $X_{t+1} = v$ with probability $\frac{c_{uv}}{\sum_{v \in V} c_{uv}}$ for every neighbor v of u. In that case, we would have $r_{uv} = 1/c_{uv}$.]

Remark 1.1. In fact, (K2) is related to a somewhat more general fact. The potential differences are given—naturally—by differences in a potential. There exists a map $\varphi : V \to \mathbb{R}$ such that $\phi_{u,v} = \varphi(u) - \varphi(v)$. If G is connected, then the potential φ is uniquely defined up to a tranlation.

To define the potential φ, put $\varphi(v_0) = 0$ for some fixed node v_0. Now for any $v \in V$ and any path $\gamma = \langle v_0, v_1, v_2, \ldots, v_k = v \rangle$ in G, we can define $\varphi(v) = \phi_{v_0,v_1} + \phi_{v_1,v_2} + \cdots + \phi_{v_{k-1},v_k}$. This is well-defined—independent of the choice of path γ—since by (K2), the potential differences around every cycle sum to zero.

Finally, we make an important definition: The *effective resistance* $R_{\text{eff}}(u, v)$ *between two nodes* $u, v \in V$ is defined to be the necessary potential difference created between u and v to induce a current of one unit to flow between them. If we imagine the entire graph G acting as a single "wire" between u and v, then $R_{\text{eff}}(u, v)$ denotes the effective resistance of that single wire (recall Ohm's law). We will now prove the following.

Theorem 1.2. *If $G = (V, E)$ has m edges, then for any two nodes $u, v \in V$, we have*

$$H_{uv} + H_{vu} = 2mR_{\text{eff}}(u, v)\,.$$

In order to prove this, we will setup four electrical networks corresponding to the graph G. We label these networks (A)-(D).

(A) We inject d_x units of flow at every vertex $x \in X$, and extract $\sum_{x \in V} d_x = 2m$ units of flow at vertex v.

(B) We inject d_x units of flow at every vertex $x \in X$, and extract $2m$ units of flow at vertex u.

(C) We inject $2m$ units of flow at vertex u and extract d_x units of flow at every vertex $x \in X$.

(D) We inject $2m$ units of flow at vertex u and extract $2m$ units of flow at vertex v.

We will use the notation $\phi^{(A)}_{x,y}, \phi^{(B)}_{x,y}$, etc. to denote the potential differences in each of these networks.

Lemma 1.3. *For any vertex $u \in V$, we have $H_{uv} = \phi^{(A)}_{u,v}$.*

Proof. Calculate: For $u \neq v$,

$$\begin{aligned} d_u &= \sum_{w \sim u} \phi^{(A)}_{u,w} \\ &= \sum_{w \sim u} \left(\phi^{(A)}_{u,v} - \phi^{(A)}_{w,v}\right) \\ &= d_u \phi^{(A)}_{u,v} - \sum_{w \sim u} \phi^{(A)}_{w,v}\,, \end{aligned}$$

where we have first use (K1), then (K2). Rearranging yields

$$\phi^{(A)}_{u,v} = 1 + \frac{1}{d_u} \sum_{w \sim u} \phi^{(A)}_{w,v}\,.$$

But now the hitting times satisfy the same set of linear equations: For $u \neq v$,

$$H_{uv} = 1 + \frac{1}{d_u} \sum_{w \sim u} H_{wv}\,.$$

We conclude that $H_{uv} = \phi^{(A)}_{u,v}$ as long as this system of linear equations has a unique solution. But consider some other solution H'_{uv} and define $f(u) = H_{uv} - H'_{uv}$. Plugging this into the preceding family of equations yields

$$f(u) = \frac{1}{d_u} \sum_{w \sim u} f(w)\,.$$

Such a map f is called *harmonic,* and it is a well-known fact that every harmonic function f on a finite, connected graph is constant. Since $f(v) = H_{vv} = 0$, this implies that $f \equiv 0$, and hence the family of equations has a unique solution, completing the proof. □

Remark 1.4. To prove that every harmonic function on a finite, connected graph is constant, we can look at the corresponding Laplace operator: $(Lf)(u) = d_u f(u) - \sum_{w \sim u} f(u)$. A function f is harmonic if and only if $Lf = 0$. But we have already seen that, on a connected graph, the Laplacian has rank $n - 1$ and $\ker(L) = \text{span}(1, \dots, 1)$, i.e. the only harmonic functions on our graph are multiples of the constant function.

Define now the *commute time between u and v* as the quantity $C_{uv} = H_{uv} + H_{vu}$. We restate and prove Theorem 1.2.

Theorem 1.5. *In any connected graph with m edges, we have $C_{uv} = 2mR_{\text{eff}}(u, v)$ for every pair of vertices u, v.*

Proof. From Lemma 1.3, we have $H_{uv} = \phi^{(A)}_{u,v}$. By symmetry, we have $H_{vu} = \phi^{(B)}_{v,u}$. Since network C is the reverse of network B, this yields $H_{vu} = \phi^{(C)}_{u,v}$. Finally, since network D is the sum of networks A and C, by linearity we have

$$\phi^{(D)}_{u,v} = \phi^{(C)}_{u,v} + \phi^{(A)}_{u,v} = H_{uv} + H_{vu} = C_{uv} .$$

Finally, note that $R_{\text{eff}}(u, v) = 2m\phi^{(D)}_{u,v}$ by definition, since network D has exactly $2m$ units of current flowing from u to v. This yields the claim of the theorem. □

1.3 Cover times

We can now use Theorem 1.5 to give a universal upper bound on the cover time of any graph.

Theorem 1.6. *For any connected graph $G = (V, E)$, we have* $\text{cov}(G) \leqslant 2|E|(|V| - 1)$.

Proof. Fix a spanning tree T of G. Then we have

$$\text{cov}(G) \leqslant \sum_{\{x,y\} \in E(T)} C_{xy} .$$

The right-hand side can be interpreted as a very particular way of covering the graph G: Start at some node x_0 and "walk" around the edges of the spanning tree in order $x_0, x_1, x_2, \dots, x_{2(n-1)} = x_0$. If we require the walk to first go from x_0 to x_1, then from x_1 to x_2, etc., we get the sum $\sum_{i=0}^{2(n-1)-1} H_{x_i x_{i+1}} = \sum_{\{x,y\} \in E(T)} C_{xy}$. This is one particular way to visit every node of G, so it gives an upper bound on the cover time.

Finally, we note that if $\{x, y\}$ is an edge of the graph, then by Theorem 1.5, we have $C_{xy} = 2|E|R_{\text{eff}}(x, y) \leqslant 2|E|$. Here we use the fact that for every edge $\{x, y\}$ of a graph, the effective resistance is at most the resistance, which is at most one. This completes the proof. □

Remark 1.7. The last stated fact is a special case of the *Rayleigh monotonicity principle.* This states that adding edges to the graph (or, more generally, decreasing the resistance of any edge) cannot increase any effective resistance. In the other direction, removing edges from the graph (or, more generally, increasing the resistance of any edge) cannot decrease any effective resistance. A similar fact is false for hitting times and commute times, as we will see in the next few examples.

Examples.

1. **The path.** Consider first G to be the path on vertices $\{0, 1, \ldots, n\}$. Then $H_{0n} + H_{n0} = C_{0n} = 2nR_{\text{eff}}(0, n) = 2n^2$. Since $H_{0n} = H_{n0}$ by symmetry, we conclude that $H_{0n} = n^2$. Note that Theorem 1.6 implies that $\text{cov}(G) \leqslant 2n^2$, and clearly $\text{cov}(G) \geqslant H_{0n} = n^2$, so the upper bound is off by at most a factor of 2.

2. **The lollipop.** Consider next the "lollipop graph" which is a path of length $n/2$ from u to v with an $n/2$ clique attached to v. We have $H_{uv} + H_{vu} = C_{uv} = \Theta(n^2)R_{\text{eff}}(u, v) = \Theta(n^3)$. On the other hand, we have already seen that $H_{uv} = \Theta(n^2)$. We conclude that $H_{vu} = \Theta(n^3)$, hence $\text{cov}(G) \geqslant \Omega(n^3)$. Again, the bound of Theorem 1.6 is $\text{cov}(G) \leqslant O(n^3)$, so it's tight up to a constant factor here as well.

3. **The complete graph.** Finally, consider the complete graph G on n nodes. In this case, Theorem 1.6 gives $\text{cov}(G) \leqslant O(n^3)$ which is way off from the actual value $\text{cov}(G) = \Theta(n \log n)$ (since this is just the coupon collector problem is flimsy disguise).

1.4 Matthews' bound

The last example shows that sometimes Theorem 1.6 doesn't give such a great upper bound. Fortunately, a relatively simple bound gets us within an $O(\log n)$ factor of the cover time.

Theorem 1.8. *If $G = (V, E)$ is a connected graph and $R = \max_{x,y \in V} R_{\text{eff}}(x, y)$ is the maximum effective resistance in G, then*

$$|E|R \leqslant \text{cov}(G) \leqslant O(\log n)|E|R\,.$$

Proof. One direction is easy:

$$\text{cov}(G) \geqslant \max_{u,v} H_{uv} \geqslant \frac{1}{2} \max_{u,v} C_{uv} = \frac{1}{2} 2|E| \max_{u,v} R_{\text{eff}}(u, v) = |E|R\,.$$

For the other direction, we will examine a random walk of length $2c|E|R \log n$ divided into $\log n$ epochs of length $2c|E|R$. Note that for any vertex v and any epoch i, we have

$$\mathbb{P}[v \text{ unvisited in epoch } i] \leqslant \frac{1}{c}\,.$$

This is because no matter what vertex is the first of epoch i, we know that the hitting time to v is at most $\max_u H_{uv} \leqslant \max_u C_{uv} \leqslant 2|E|R$. Now Markov's inequality tells us that the probability it takes more than $2c|E|R$ steps to hit v is at most $1/c$.

Therefore the probability we don't visit v in any epoch is at most $c^{-\log n} = n^{-\log c}$, and by a union bound, the probability that there is some vertex left unvisited after all the epochs is at most $n^{1-\log c}$.

We conclude that

$$\text{cov}(G) \leqslant 2c|E|R \log n + n^{1-\log c} 2n^3\,,$$

where we have used the trivial bound on the cover time from Theorem 1.6. Choosing c to be a large enough constant makes the second term negligible, yielding

$$\text{cov}(G) \leqslant O(|E|R \log n)\,,$$

as desired. □

One can make some improvements to this "soft" proof, yielding the following stronger bounds.

Theorem 1.9. *For any connected graph $G = (V, E)$, the following holds. Let t_{hit} denote the maximum hitting time in G. Then*

$$\text{cov}(G) \leqslant t_{\text{hit}} \left(1 + \frac{1}{2} + \cdots + \frac{1}{n}\right).$$

Moreover, if we define for any subset $A \subseteq V$, the quantity $t^A_{\text{min}} = \min_{u,v \in A, u \neq v} H_{uv}$, then

$$\text{cov}(G) \geqslant \max_{A \subseteq V} t^A_{\text{min}} \left(1 + \frac{1}{2} + \cdots + \frac{1}{|A| - 1}\right). \tag{1.1}$$

For the proofs, consult Chapter 11 of the Levin-Peres-Wilmer book http://pages.uoregon.edu/dlevin/MARKOV/markovmixing.pdf.

Kahn, Kim, Lovász, and Vu showed that the best lower bound in (1.1) is within an $O(\log \log n)^2$ factor of $\text{cov}(G)$, improving over the $O(\log n)$-approximation in Theorem 1.8. In a paper with Jian Ding and Yuval Peres, we showed that one can keep going and compute an $O(1)$ approximation using a multi-scale generalization of the bound (1.1) based on Talagrand's majorizing measures theory.