Chapter 16

Random Walks III

By Sariel Har-Peled, December 30, 2015[1]

> "I gave the girl my protection, offering in my equivocal way to be her father. But I came too late, after she had ceased to believe in fathers. I wanted to do what was right, I wanted to make reparation: I will not deny this decent impulse, however mixed with more questionable motives: there must always be a place for penance and reparation. Nevertheless, I should never have allowed the gates of the town to be opened to people who assert that there are higher considerations that those of decency. They exposed her father to her naked and made him gibber with pain, they hurt her and he could not stop them (on a day I spent occupied with the ledgers in my office). Thereafter she was no longer fully human, sister to all of us. Certain sympathies died, certain movements of the heart became no longer possible to her. I too, if I live longer enough in this cell with its ghost not only of the father and the daughter but of the man who even by lamplight did not remove the black discs from his eyes and the subordinate whose work it was to keep the brazier fed, will be touched with the contagion and turned into a create that believes in nothing."
>
> – J. M. Coetzee, Waiting for the Barbarians.

16.1. Random Walks on Graphs

Let $\mathsf{G} = (\mathsf{V}, \mathsf{E})$ be a connected, non-bipartite, undirected graph, with n vertices. We define the natural Markov chain on G, where the transition probability is

$$\mathsf{P}_{uv} = \begin{cases} \frac{1}{d(u)} & \text{if } uv \in \mathsf{E} \\ 0 & \text{otherwise,} \end{cases}$$

where $d(w)$ is the degree of vertex w. Clearly, the resulting Markov chain M_G is irreducible. Note, that the graph must have an odd cycle, and it has a cycle of length 2. Thus, the gcd of the lengths of its cycles is 1. Namely, M_G is aperiodic. Now, by the Fundamental theorem of Markov chains, M_G has a unique stationary distribution π.

Lemma 16.1.1. *For all $v \in \mathsf{V}$, we have $\pi_v = d(v)/2m$.*

Proof: Since π is stationary, and the definition of P_{uv}, we get

$$\pi_v = [\pi \mathbf{P}]_v = \sum_{uv} \pi_u \mathsf{P}_{uv},$$

and this holds for all v. We only need to verify the claimed solution, since there is a unique stationary distribution. Indeed,

$$\frac{d(v)}{2m} = \pi_v = [\pi \mathbf{P}]_v = \sum_{uv} \frac{d(u)}{2m} \frac{1}{d(u)} = \frac{d(v)}{2m},$$

as claimed. ■

Lemma 16.1.2. *For all $v \in \mathsf{V}$, we have $\mathtt{h}_{vv} = 1/\pi_v = 2m/d(v)$.*

Definition 16.1.3. The *hitting time* $\mathtt{h}_{uv}$ is the expected number of steps in a random walk that starts at u and ends upon first reaching v.

The *commute time* between u and v is denoted by $\mathbf{CT}_{uv} = \mathtt{h}_{uv} + \mathtt{h}_{vu}$.

Let $\mathcal{C}_u(G)$ denote the expected length of a walk that starts at u and ends upon visiting every vertex in G at least once. The *cover time* of G denotes by $\mathcal{C}(G)$ is defined by $\mathcal{C}(G) = \max_u \mathcal{C}_u(G)$.

Example 16.1.4 (Lollipop). Let L_{2n} be the $2n$-vertex *lollipop graph*, this graph consists of a clique on n vertices, and a path on the remaining n vertices. There is a vertex u in the clique which is where the path is attached to it. Let v denote the end of the path, see figure on the right.

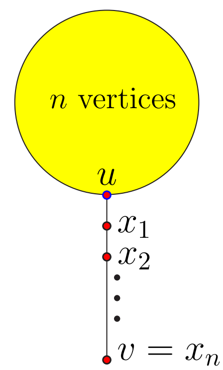

Taking a random walk from u to v requires in expectation $O(n^2)$ steps, as we already saw in class. This ignores the probability of escape – that is, with probability $(n-1)/n$ when at u we enter the clique K_n (instead of the path). As such, it turns out that $\mathtt{h}_{uv} = \Theta(n^3)$, and $\mathtt{h}_{vu} = \Theta(n^2)$. (Thus, hitting times are not symmetric!)

Note, that the cover time is not monotone decreasing with the number of edges. Indeed, the path of length n, has cover time $O(n^2)$, but the larger graph L_n has cover time $\Omega(n^3)$.

Example 16.1.5 (More on walking on the Lollipop). To see why $\mathtt{h}_{uv} = \Theta\left(n^3\right)$, number the vertices on the stem $x_1, \ldots, x_n$. Let T_i be the expected time to arrive to the vertex x_i when starting a walk from u. Observe, that surprisingly, $T_1 = \Theta(n^2)$. Indeed, the walk has to visit the vertex u about n times in expectation, till the walk would decide to go to x_1 instead of falling back into the clique. The time between visits to u is in expectation $O(n)$ (assuming the walk is inside the clique).

Now, observe that $T_{2i} = T_i + \Theta(i^2) + \frac{1}{2} T_{2i}$. Indeed, starting with x_i, it takes in expectation $\Theta(i^2)$ steps of the walk to either arrive (with equal probability) at x_{2i} (good), or to get back to u (oopsi). In the later case, the game begins from scratch. As such, we have that

$$T_{2i} = 2T_i + \Theta\left(i^2\right) = 2\left(2T_{i/2} + \Theta\left((i/2)^2\right)\right) + \Theta\left(i^2\right) = \cdots = 2iT_1 + \Theta\left(i^2\right),$$

assuming i is a power of two (why not?). As such, $T_n = nT_1 + \Theta(n^2)$. Since $T_1 = \Theta(n^2)$, we have that $T_n = \Theta(n^3)$.

Definition 16.1.6. A $n \times n$ matrix M is *stochastic* if all its entries are non-negative and for each row i, it holds $\sum_k \mathsf{M}_{ik} = 1$. It is *doubly stochastic* if in addition, for any i, it holds $\sum_k \mathsf{M}_{ki} = 1$.

Lemma 16.1.7. *Let* MC *be a Markov chain, such that transition probability matrix* $\mathbf{P}$ *is* doubly stochastic. *Then, the distribution $u = (1/n, 1/n, \ldots, 1/n)$ is stationary for* MC.

Proof: $[u\mathbf{P}]_i = \sum_{k=1}^{n} \frac{\mathsf{P}_{ki}}{n} = \frac{1}{n}.$ ■

Lemma 16.1.8. *For any edge* $(u \to v) \in \mathsf{E}$*, we have* $\mathtt{h}_{uv} + \mathtt{h}_{vu} \leq 2m$*.*

(Note, that $(u \to v)$ being an edge in the graph is crucial. Indeed, without it a significantly worst case bound holds, see Theorem 16.2.1.)

Proof: Consider a new Markov chain defined by the edges of the graph (where every edge is taken twice as two directed edges), where the current state is the last (directed) edge visited. There are $2m$ edges in the new Markov chain, and the new transition matrix, has $Q_{(u \to v),(v \to w)} = \mathtt{P}_{vw} = \frac{1}{d(v)}$. This matrix is *doubly stochastic*, meaning that not only do the rows sum to one, but the columns sum to one as well. Indeed, for the $(v \to w)$ we have

$$\sum_{x \in V, y \in \Gamma(x)} Q_{(x \to y),(v \to w)} = \sum_{u \in \Gamma(v)} Q_{(u \to v),(v \to w)} = \sum_{u \in \Gamma(v)} \mathtt{P}_{vw} = d(v) \times \frac{1}{d(v)} = 1.$$

Thus, the stationary distribution for this Markov chain is uniform, by Lemma 16.1.7. Namely, the stationary distribution of $e = (u \to v)$ is $\mathtt{h}_{ee} = \pi_e = 1/(2m)$. Thus, the expected time between successive traversals of e is $1/\pi_e = 2m$, by Theorem 16.3.1 (iii).

Consider $\mathtt{h}_{uv} + \mathtt{h}_{vu}$ and interpret this as the time to go from u to v and then return to u. Conditioned on the event that the initial entry into u was via the $(v \to u)$, we conclude that the expected time to go from there to v and then finally use $(v \to u)$ is $2m$. The memorylessness property of a Markov chains now allows us to remove the conditioning: since how we arrived to u is not relevant. Thus, the expected time to travel from u to v and back is at most $2m$. ■

16.2. Electrical networks and random walks

A ***resistive electrical network*** is an undirected graph; each edge has ***branch resistance*** associated with it. The electrical flow is determined by two laws: *Kirchhoff's law* (preservation of flow - all the flow coming into a node, leaves it) and *Ohm's law* (the voltage across a resistor equals the product of the resistance times the current through it). Explicitly, Ohm's law states

$$\text{voltage} = \text{resistance} * \text{current}.$$

The *effective resistance* between nodes u and v is the voltage difference between u and v when one ampere is injected into u and removed from v (or injected into v and removed from u). The effective resistance is always bounded by the branch resistance, but it can be much lower.

Given an undirected graph G, let $\mathcal{N}(G)$ be the electrical network defined over G, associating one ohm resistance on the edges of $\mathcal{N}(G)$.

You might now see the connection between a random walk on a graph and electrical network. Intuitively (used in the most unscientific way possible), the electricity, is made out of electrons each one of them is doing a random walk on the electric network. The resistance of an edge, corresponds to the probability of taking the edge. The higher the resistance, the lower the probability that we will travel on this edge. Thus, if the effective resistance $\mathbf{R}_{uv}$ between u and v is low, then there is a good probability that travel from u to v in a random walk, and $\mathtt{h}_{uv}$ would be small.

Theorem 16.2.1. *For any two vertices u and v in G, the commute time $\mathbf{CT}_{uv} = 2m\mathbf{R}_{uv}$, where $\mathbf{R}_{uv}$ is the effective resistance between u and v.*

Proof: Let ϕ_{uv} denote the voltage at u in $\mathcal{N}(G)$ with respected to v, where $d(x)$ amperes of current are injected into each node $x \in \mathsf{V}$, and $2m$ amperes are removed from v. We claim that

$$\mathtt{h}_{uv} = \phi_{uv}.$$

Note, that the voltage on an edge xy is $\phi_{xy} = \phi_{xv} - \phi_{yv}$. Thus, using Kirchhoff's Law and Ohm's Law, we obtain that

$$x \in V \setminus \{v\} \qquad d(x) = \sum_{w \in \Gamma(x)} \text{current}(xw) = \sum_{w \in \Gamma(x)} \frac{\phi_{xw}}{\text{resistance}(xw)} = \sum_{w \in \Gamma(x)} (\phi_{xv} - \phi_{wv}), \tag{16.1}$$

since the resistance of every edge is 1 ohm. (We also have the "trivial" equality that $\phi_{vv} = 0$.) Furthermore, we have only n variables in this system; that is, for every $x \in \mathsf{V}$, we have the variable ϕ_{xv}.

Now, for the random walk interpretation – by the definition of expectation, we have

$$x \in V \setminus \{v\} \qquad \mathtt{h}_{xv} = \frac{1}{d(x)} \sum_{w \in \Gamma(x)} (1 + \mathtt{h}_{wv}) \quad \iff \quad d(x)\,\mathtt{h}_{xv} = \sum_{w \in \Gamma(x)} 1 + \sum_{w \in \Gamma(x)} \mathtt{h}_{wv}$$
$$\iff \quad \sum_{w \in \Gamma(x)} 1 = d(x)\,\mathtt{h}_{xv} - \sum_{w \in \Gamma(x)} \mathtt{h}_{wv} = \sum_{w \in \Gamma(x)} (\mathtt{h}_{xv} - \mathtt{h}_{wv}).$$

Since $d(x) = \sum_{w \in \Gamma(x)} 1$, this is equivalent to

$$x \in V \setminus \{v\} \qquad d(x) = \sum_{w \in \Gamma(x)} (\mathtt{h}_{xv} - \mathtt{h}_{wv}). \tag{16.2}$$

Again, we also have the trivial equality $\mathtt{h}_{vv} = 0$.[2] Note, that this system also has n equalities and n variables.

Eq. (16.1) and Eq. (16.2) show two systems of linear equalities. Furthermore, if we identify $\mathtt{h}_{uv}$ with ϕ_{xv} then they are exactly the same system of equalities. Furthermore, since Eq. (16.1) represents a physical system, we know that it has a unique solution. This implies that $\phi_{xv} = \mathtt{h}_{xv}$, for all $x \in \mathsf{V}$.

Imagine the network where u is injected with $2m$ amperes, and for all nodes w remove $d(w)$ units from w. In this new network, $\mathtt{h}_{vu} = -\phi'_{vu} = \phi'_{uv}$. Now, since flows behaves linearly, we can superimpose them (i.e., add them up). We have that in this new network $2m$ unites are being injected at u, and $2m$ units are being extracted at v, all other nodes the charge cancel itself out. The voltage difference between u and v in the new network is $\widehat{\phi} = \phi_{uv} + \phi'_{uv} = \mathtt{h}_{uv} + \mathtt{h}_{vu} = \mathbf{CT}_{uv}$. Now, in the new network there are $2m$ amperes going from u to v, and by Ohm's law, we have

$$\widehat{\phi} = \text{voltage} = \text{resistance} * \text{current} = 2m\mathbf{R}_{uv},$$

as claimed. ∎

Example 16.2.2. Recall the lollipop L_n from Exercise 16.1.4. Let u be the connecting vertex between the clique and the stem (i.e., the path). We inject $d(x)$ units of flow for each vertex x of L_n, and collect $2m$ units at u. Next, let $u = x_0, x_1, \ldots, x_n = v$ be the vertices of the stem. Clearly, there are $2(n - i) - 1$ units of electricity flowing on the edge $(x_{i+1} \to x_i)$. Thus, the voltage on this edge is $2(n - i)$, by Ohm's law (every edge has resistance one). The effective resistance from v to u is as such $\Theta(n^2)$, which implies that $\mathtt{h}_{vu} = \Theta(n^2)$.

Similarly, it is easy to show $\mathtt{h}_{uv} = \Theta(n^3)$.

A similar analysis works for the random walk on the integer line in the range 1 to n.

Lemma 16.2.3. *For any n vertex connected graph G, and for all $u, v \in V(G)$, we have $\mathbf{CT}_{uv} < n^3$.*

Proof: The effective resistance between any two nodes in the network is bounded by the length of the shortest path between the two nodes, which is at most $n - 1$. As such, plugging this into Theorem 16.2.1, yields the bound, since $m < n^2$. ∎

[2] In previous lectures, we interpreted $\mathtt{h}_{vv}$ as the expected length of a walk starting at v and coming back to v.

16.3. Tools from previous lecture

Theorem 16.3.1 (Fundamental theorem of Markov chains). *Any irreducible, finite, and aperiodic Markov chain has the following properties.*

(i) *All states are ergodic.*

(ii) *There is a unique stationary distribution* π *such that, for* $1 \le i \le n$*, we have* $\pi_i > 0$.

(iii) *For* $1 \le i \le n$*, we have* $\mathtt{f}_{ii} = 1$ *and* $\mathtt{h}_{ii} = 1/\pi_i$.

(iv) *Let* $N(i,t)$ *be the number of times the Markov chain visits state* i *in* t *steps. Then*

$$\lim_{t\to\infty} \frac{N(i,t)}{t} = \pi_i.$$

Namely, independent of the starting distribution, the process converges to the stationary distribution.

16.4. Bibliographical Notes

A nice survey of the material covered here, is available online at http://arxiv.org/abs/math.PR/0001057 [DS00].

Bibliography

[DS00] P. G. Doyle and J. L. Snell. Random walks and electric networks. *ArXiv Mathematics e-prints*, 2000.