Random Walks and Diffusion

Daniel A. Spielman September 28, 2010

8.1 Overview

We introduce random walks and diffusion on graphs. These are among the most-studied and best-behaved processes on graphs. They will play a crucial role in many of the analyses of graphs that we will perform in the last segment of the course. We will understand many other processes by comparison with these.

8.2 Random Walks

In this lecture, we will consider random walks on undirected graphs. Let's begin with the definitions. Let $G = (V, E)$ be an undirected graph. A random walk on a graph is a process that begins at some vertex, and at each time step moves to another vertex. The vertex the walk moves to is chosen uniformly at random among the neighbors of the present vertex. There are many things we can measure about a random walk, including

1. What is the limiting distribution of the random walk?

2. How long should it take before the walk approaches the limiting distribution?

3. How long should it take before the walk hits every vertex?

3. If we start a walk at s, what is the expected number of steps before it hits t?

Instead of keeping track of where a particular random walk is in a particular simulation, we usually measure the *probability distribution* of the random walk. That is, the probability at a given step that it is at any given vertex. We keep treat probability distributions on the vertices as vectors $\boldsymbol{p} \in \mathbb{R}^n$. I will sometimes write $\boldsymbol{p} \in \mathbb{R}^V$ to emphasize that $\boldsymbol{p}$ is a vector indexed by the vertices of the graph, or I may even write $\boldsymbol{p} : V \to \mathbb{R}$. I will write $\boldsymbol{p}(u)$ to indicate the value of $\boldsymbol{p}$ at a vertex u–that is the probability of being at vertex u in distribution $\boldsymbol{p}$. A probability vector $\boldsymbol{p}$ should satisfy $\boldsymbol{p}(u) \geq 0$, for all $u \in V$, and

$$\sum_u \boldsymbol{p}(u) = 1.$$

We let $\boldsymbol{p}_0$ denote the initial probability distribution. That is, $\boldsymbol{p}_0$ encodes where we start the random walk. If we start the random walk at a vertex v, then

$$\boldsymbol{p}_0(w) = \begin{cases} 1 & \text{if } w = v \\ 0 & \text{otherwise.} \end{cases}$$

But, we could start the walk at a vertex chosen from some distribution.

We will let the vector $\boldsymbol{p}_t \in \mathbb{R}^n$ denote the probability distribution at time t. To derive a $\boldsymbol{p}_{t+1}$ from $\boldsymbol{p}_t$, note that the probability of being at a vertex u at time $t+1$ is the sum over the neighbors v of u of the probability that the walk was at v at time t, times the probability it moved from v to u in time $t+1$. Algebraically, we have

$$\boldsymbol{p}_{t+1}(u) = \sum_{v:(u,v)\in E} \boldsymbol{p}_t(v)/d(v), \tag{8.1}$$

where $d(v)$ is the degree of vertex v.

We will often consider lazy random walks, which are the variant of random walks that stay put with probability $1/2$ at each time step, and walk to a random neighbor the other half of the time. These evolve according to the equation

$$\boldsymbol{p}_{t+1}(u) = (1/2)\boldsymbol{p}_t(u) + (1/2) \sum_{v:(u,v)\in E} \boldsymbol{p}_t(v)/d(v). \tag{8.2}$$

8.3 Diffusion

There are a few types of diffusion that people study in a graph, but the most common is closely related to random walks. In a diffusion process, we imagine that we have some substance that can occupy the vertices, such as a gas or fluid. At each time step, some of the substance diffuses out of each vertex. If we say that half the substance stays at a vertex at each time step, and the other half is distributed equally among its neighboring vertices, then the distribution of the substance will evolve according to equation (8.2). That is, probability mass in the lazy random walk obeys this diffusion equation.

People consider finer time steps in which smaller fractions of the mass leave the vertices. In the limit, this results in continuous random walks. But, that is not a topic for another day.

A variation of diffusion that I enjoy involves placing both matter and anti-matter at vertices. These cancel each other out when they meet. We can then ask question likes "if we put 1 unit of matter at v and 1 unit of anti-matter at u, how much matter will be left after time t?"

8.4 Matrix form

The right way to understand the behavior of random walks is through linear algebra. First, we will re-write equation (8.1) in matrix form. It will help to identify the vertices with the set $\{1, \ldots, n\}$.

We then define the adjacency matrix, $\boldsymbol{A}$, of the graph by setting $\boldsymbol{A}(u,v)$ to 1 if $(u,v) \in E$ and 0 otherwise (note that I write $\boldsymbol{A}(u,v)$ to indicate entry of $\boldsymbol{A}$ in row u and column v). We then define the matrix $\boldsymbol{D}$ to be the diagonal matrix of degrees. That is $\boldsymbol{D}(u,u) = d(u)$, and $\boldsymbol{D}(u,v) = 0$ for $u \neq v$. Note that $\boldsymbol{D}^{-1}$ is also a diagonal matrix, and its uth diagonal entry is $1/d(u)$.

We now define

$$\boldsymbol{W} = \boldsymbol{A}\boldsymbol{D}^{-1}$$

to be the *walk matrix* of the graph. You can compute that

$$\boldsymbol{W}(u,v) = \begin{cases} 1/d(v) & \text{if } (u,v) \in E \\ 0 & \text{otherwise.} \end{cases}$$

We will now see that the update rule (8.1) is equivalent to

$$\boldsymbol{p}_{t+1} = \boldsymbol{W}\boldsymbol{p}_t.$$

This follows from the definition of matrix multiplication, and is obvious once you see it. From the definition of matrix multiplication, we know that

$$\boldsymbol{p}_{t+1}(u) = \sum_v \boldsymbol{W}(u,v)\boldsymbol{p}_t(v)$$

$$= \sum_{v:(u,v)\in E} (1/d(v))\boldsymbol{p}_t(v).$$

You should check that this is identical to (8.1).

This formulation makes it easy to formalize the matter/anti-matter problem I gave a few moments ago. To start with 1 unit of matter at v and 1 unit of anti-matter at u, use the initial vector $\boldsymbol{p}_0$ where

$$\boldsymbol{p}_0(w) = \begin{cases} 1 & \text{if } w = v \\ -1 & \text{if } w = u \\ 0 & \text{otherwise.} \end{cases}$$

After t steps, the distribution of matter and anti-matter is given by $\boldsymbol{W}^t\boldsymbol{p}_0$. The cancellations happen automatically.

The matrix that realizes the lazy random walk is

$$\widehat{\boldsymbol{W}} \stackrel{\text{def}}{=} (1/2)(\boldsymbol{I} + \boldsymbol{W}).$$

8.5 The steady-state distribution

A natural question to ask is whether the distribution of a random walk eventually converges to a steady state. It usually, but not always, does. For example, consider the graph with one edge and two vertices. The random walk will alternate vertices every other step, and never converge to one fixed distribution. A similar thing happens when the graph is bipartite: the walk goes from side

to side every other step. It turns out that bipartite graphs are the only ones on which the random walk does not converge to a steady state. This is one motivation for studying lazy random walks. Lazy random walks always converge to a steady-state distribution. We will prove this next lecture.

For now, we just observe that there is a steady-state distribution and prove that it is unique. The steady-state distribution appears at each vertex with probability proportional to its degree. It is given by the vector $\boldsymbol{\pi}$ which we define by

$$\boldsymbol{\pi}(u) = \frac{d(u)}{\sum_{v \in V} d(v)}.$$

It is easy to verify that $\boldsymbol{\pi}$ is a probability vector, and we will now observe that

$$\boldsymbol{W}\boldsymbol{\pi} = \boldsymbol{\pi}. \tag{8.3}$$

So, if the walk ever reaches the distribution $\boldsymbol{\pi}$, it will remain in that distribution. To see this, it suffices to ignore the denominator and just show that $W\boldsymbol{d} = \boldsymbol{d}$, where $\boldsymbol{d}$ is the vector of degrees. To prove this, note that $D^{-1}\boldsymbol{d} = \boldsymbol{1}$, the all-1's vector, and $\boldsymbol{A}\boldsymbol{1} = \boldsymbol{d}$, as the u-th coordinate of $\boldsymbol{A}\boldsymbol{1}$ is

$$\sum_{v:(u,v)\in E} 1 = d(u).$$

Note that equation (8.3) says that $\boldsymbol{\pi}$ is a right-eigenvector of $\boldsymbol{W}$ of eigenvalue 1. We will exploit this formulation more in the next lecture.

8.6 Uniqueness of the steady-state distribution

I'll now prove to you that if G is connected then the steady-state distribution is unique.

Lemma 8.6.1. *Let G be a connected graph and let $\boldsymbol{p}$ be a probability vector such that $\boldsymbol{W}\boldsymbol{p} = \boldsymbol{p}$. Then $\boldsymbol{p} = \boldsymbol{\pi}$.*

Proof. First, let u be any vertex that maximizes

$$\frac{\boldsymbol{p}(u)}{d(u)}.$$

From the assumption that $\boldsymbol{W}\boldsymbol{p} = \boldsymbol{p}$, we have

$$\boldsymbol{p}(u) = \sum_{v:(u,v)\in E} \frac{\boldsymbol{p}(v)}{d(v)}$$

$$\leq \sum_{v:(u,v)\in E} \frac{\boldsymbol{p}(u)}{d(u)}$$

$$= \boldsymbol{p}(u).$$

Now, the only way we can achieve this equality is if all of the inequalities on the second line are in fact equalities. This means that

$$\frac{\boldsymbol{p}(v)}{d(v)} = \frac{\boldsymbol{p}(u)}{d(u)} \qquad (8.4)$$

for all neighbors v of u.

We can use this argument to show that this ratio is fixed for every vertex in the graph. To do that formally, let S be the set of vertices v for which (8.4) holds. Assume by way of contradiction that there is some vertex w not in S. As the graph is connected, it contains a path between u and w. Thus, there must be some edge (y, z) for which $y \in S$ and $z \notin S$. To obtain a contradiction, repeat the argument from the beginning of the proof with y in the role of u. It tells us that

$$\frac{\boldsymbol{p}(y)}{d(y)} = \frac{\boldsymbol{p}(z)}{d(z)},$$

which contradicts our assumption that $y \in S$ and $z \notin S$. □

8.7 Types of Convergence

In the next lecture, we will analyze how quickly the distribution of a random walk converges to the steady state and we will examine the obstacles to fast convergence. But first, we need to figure out what it means to converge.

The simplest measures of convergence reduce this to one number measuring how far $\boldsymbol{p}_t$ is from $\boldsymbol{\pi}$. One can measure this as a norm of the vector $\boldsymbol{p}_t - \boldsymbol{\pi}$. Different norms are useful in different applications. The standard is the Euclidean norm, also called the 2-norm:

$$\|\boldsymbol{x}\|_2 = \sqrt{\sum_u \boldsymbol{x}(u)^2}.$$

This one is so standard that the subscripted 2 is often dropped.

Another standard norm is the 1-norm:

$$\|\boldsymbol{x}\|_1 = \sum_u |\boldsymbol{x}(u)|.$$

This gives a distance called the *total variation* distance between $\boldsymbol{p}_t$ and $\boldsymbol{\pi}$.

The third common norm is called the max-norm or the infinity-norm, and it measures the maximum difference between $\boldsymbol{p}_t$ and $\boldsymbol{\pi}$:

$$\|\boldsymbol{x}\|_\infty = \max_u |\boldsymbol{x}(u)|.$$

8.8 A more holistic approach

I often find it useful to examine a plot of the whole distribution $\boldsymbol{p}_t$. To set this up, I would first like to view the random walk as living on sockets rather than on vertices. Recall that I define a

socket to be the endpoint of an edge where it attaches to a vertex. For example, the edge (u, v) has two sockets: one at u and one at v. If the graph has m edges then it has $2m$ sockets.

When the random walk is at u, it next moves to a random neighbor of u. We can view this as first moving to a random socket attached to u. If it follows this socket to node v, we can then immediately pick the socket attached to v that it will traverse next. Recall that the stead-state distribution $\boldsymbol{\pi}$ visits each vertex with probability proportional to its degree. If we transfer this distribution to the sockets, then we see that it visits each socket with *equal* probability.

Assign a number between 1 and $2m$ to each socket. I will denote a probability distribution on sockets by a vector $\boldsymbol{q} \in \mathbb{R}^{2m}$. Given a probability distribution on the vertices $\boldsymbol{p}$, we can derive the corresponding distribution on sockets by setting $\boldsymbol{q}(s)$ to be $\boldsymbol{p}(u)/d(u)$ when s is a socket attached to vertex u.

I will take this vector $\boldsymbol{q}$, sort it from largest to smallest, and measure its cumulative sum. To be concrete, I will do with a toy vector $\boldsymbol{q}$.

```
>> q = [.1, .2, .3, .4];
>> qs = sort(q,'descend')

qs =

    0.4000    0.3000    0.2000

>> plot([0:4],[0,cumsum(qs)])
>> hold on
>> plot([0:4],[0,cumsum(qs)],'o'
```

Let me state this formally. For each integer k between 0 and $2m$ I define $C_{\boldsymbol{q}}$ to be the sum of the largest k entries in the vector $\boldsymbol{q}$. I then plotted k against $C_{\boldsymbol{q}}(k)$, drawing straight lines between the integral points. When $\boldsymbol{q}$ is the uniform distribution, which is equal to $1/2m$ in each entry, $C_{\boldsymbol{q}}(k) = k/2m$. So, this plot is a straight line.

For every other probability distribution $\boldsymbol{q}$ on sockets, the plot of $C_{\boldsymbol{q}}$ is concave and lies above the straight line from $(0,0)$ to $(2m, 1)$. One can show that $C_{\boldsymbol{q}_{t+1}}$ lies beneath $C_{\boldsymbol{p}_t}$ for every t. If I get to this material in today's lecture, then we will use this approach next lecture to show that $C_{\boldsymbol{q}_t}$ converges to the straight line.