Graphs and Networks Lecture 12

Harmonic Functions on Graphs, Three Ways

Daniel A. Spielman October 12, 2010

12.1 Overview

We will examine physical metaphors for graphs. We begin with a spring model and then discuss resistor networks.

There is a third interesting approach using random walks. I didn't get to cover it in class, but I include some notes about it at the end of these notes.

12.2 Interpolation on a graph

Say you have a graph $G = (V, E)$, a special set of vertices W, and a function $f : W \rightarrow \mathbb{R}$ whose values you only know at W. You would like to guess the values of f at the vertices not in W. This is a common problem in the field of Machine Learning, and we will see one reasonable solution to it today.

To begin, imagine that every edge is a rubber band (or spring), and that every vertex is a little ring to which all its edge rubber bands are connected. For each vertex $v \in W$, we nail the corresponding ring down onto the real line at $f(v)$. We could let the other rings settle into position, and guess that $f(u)$ is the location of ring u. Hey, why not?

This approach has many reasonable properties, and we will spend this lecture exploring them.

12.3 Hooke's Law

We can use Hooke's law to figure out where all the rings in $V - W$ should wind up. Assume that each rubber band is an idea spring with spring constant 1. Actually, if you have weights on the edges, you could make these the spring constants. If a rubber band connects vertices u and v, then Hooke's law tells us that the force it exerts at node u is in the direction of v and is proportional to the distance between u and v. Let $\boldsymbol{x}(u)$ be the position of each vertex u. Then the force the rubber band between u and v exerts on u is That is,

$$\boldsymbol{x}(v) - \boldsymbol{x}(u).$$

In a stable configuration, all of the vertices that have not been nailed down must experience a zero net force. That is

$$\sum_{v:(u,v)\in E} \boldsymbol{x}(v)\boldsymbol{x}(u) = 0 \implies \sum_{v:(u,v)\in E} \boldsymbol{x}(v) = d(u)\boldsymbol{x}(u). \implies \frac{1}{d(u)} \sum_{v:(u,v)\in E} \boldsymbol{x}(v) = \boldsymbol{x}(u).$$

So, each vertex that is not nailed down is the average of its neighbors.

In the weighted case, we would have for each $u \in V - W$

$$\frac{1}{d(u)} \sum_{v:(u,v)\in E} w_{u,v}\boldsymbol{x}(v) = \boldsymbol{x}(u), \tag{12.1}$$

where here we define

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

to be the weighted degree of node u.

A function $\boldsymbol{x}$ that satisfies these equations for each vertex $u \in V - W$ is said to be *harmonic* on $V - W$.

During this lecture, we will prove that the solution to these equations is unique provided that the graph is connected and that W has at least one vertex. This is probably obvious to those of you who believe in Physics. But I'm just a mathematician who thinks of Physics as a nice story to explain pretty equations. So, I need to prove it.

For example, let me point out that we have not yet even proved that there is a solution to the harmonic equations (12.1).

12.4 Energy

Physics also tells us that the nodes will settle into the position that minimizes the potential energy. The potential energy of an ideal linear spring with constant w when stretched to length l is

$$\frac{1}{2}wl^2.$$

So, the potential energy in a configuration $\boldsymbol{x}$ is given by

$$\mathcal{E}(\boldsymbol{x}) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{(u,v)\in E} w_{u,v}(\boldsymbol{x}(u) - \boldsymbol{x}(v))^2. \tag{12.2}$$

Note that the energy always has to be at least 0. This tells us that at least one minimum must exist. We now use this fact to see that the harmonic equations (12.1) do in fact have a solution.

For any $\boldsymbol{x}$ that minimizes the energy, the partial derivative of the energy with respect to each variable must be zero. In this case, the variables are $\boldsymbol{x}(u)$ for $u \in V - W$. The partial derivative with respect to $\boldsymbol{x}(u)$ is

$$\frac{1}{2} \sum_{v:(u,v)\in E} w_{u,v}2(\boldsymbol{x}(u) - \boldsymbol{x}(v)) = \sum_{v:(u,v)\in E} w_{u,v}(\boldsymbol{x}(u) - \boldsymbol{x}(v)).$$

Setting this to zero gives (12.1), and another way to see that the solution should be harmonic. Since we know that there is at least one $\boldsymbol{x}$ that minimizes the energy, we now know that there is at least one $\boldsymbol{x}$ that satisfies the harmonic equations. In a few moments, we will see that both solutions are in fact unique, provided that the graph is connected.

Theorem 12.4.1. *If G is a connected graph with all edges weights positive and $|W| \geq 1$, then the minimizer of* (12.2) *is unique.*

Proof. Let's assume by way of contradiction that there are two diiferent minimizers of (12.2), $\boldsymbol{x}$ and $\boldsymbol{y}$. We will prove that for every edge (u, v)

$$\boldsymbol{x}(u) - \boldsymbol{x}(v) = \boldsymbol{y}(u) - \boldsymbol{y}(v). \tag{12.3}$$

Once you know this, it is easy to use the connectivity of the graph to prove that $\boldsymbol{x} = \boldsymbol{y}$.

Consider taking the average of $\boldsymbol{x}$ and $\boldsymbol{y}$:

$$\boldsymbol{z} = \frac{1}{2}(\boldsymbol{x} + \boldsymbol{y}).$$

It is clear that $\boldsymbol{z}$ takes the same values as $\boldsymbol{x}$ and $\boldsymbol{y}$ on W. We now show that if there is some edge (u, v) for which (12.3) fails to hold, then $\mathcal{E}(\boldsymbol{z}) \leq \mathcal{E}(\boldsymbol{x})$, a contradiction.

We will use the following inequality, which I quickly prove.

$$\left(\frac{a+b}{2}\right)^2 \leq \frac{1}{2}a^2 + \frac{1}{2}b^2,$$

with equality if and only if $a = b$. To see this, we subtract the left side from the right side, giving

$$\frac{1}{2}a^2 + \frac{1}{2}b^2 - \left(\frac{a+b}{2}\right)^2 = \frac{1}{4}a^2 + \frac{1}{4}b^2 - \frac{1}{2}ab = \frac{1}{4}(a-b)^2.$$

Using this inequality, we see that

$$\begin{aligned}
\mathcal{E}(\boldsymbol{z}) &= \sum_{(u,v)\in E} w_{u,v}(\boldsymbol{z}(u) - \boldsymbol{z}(v))^2 \\
&= \sum_{(u,v)\in E} w_{u,v}\left(\frac{1}{2}(\boldsymbol{x}(u) - \boldsymbol{x}(v)) + \frac{1}{2}(\boldsymbol{y}(u) - \boldsymbol{y}(v))\right)^2 \\
&\leq \frac{1}{2}\sum_{(u,v)\in E} w_{u,v}\left((\boldsymbol{x}(u) - \boldsymbol{x}(v))\right)^2 + \frac{1}{2}\sum_{(u,v)\in E} w_{u,v}\left((\boldsymbol{y}(u) - \boldsymbol{y}(v))\right)^2 \\
&= \frac{1}{2}\mathcal{E}(\boldsymbol{x}) + \frac{1}{2}\mathcal{E}(\boldsymbol{y}) = \mathcal{E}(\boldsymbol{x}).
\end{aligned}$$

Moreover, the inequality is strict if there is some edge (u, v) for which (12.3) is violated. $\square$

So, we now know that the minimizer of the energy is unique. Does this tell us that the solution to the harmonic equations is also unique?

Not necessarily. The harmonic equations would be satisfied by any local minimizer of the energy. But, we can show that there are no local minimizers other than the global minimum, which implies that the solution to the harmonic equations is unique as well.

Theorem 12.4.2. *The only local minimizer of* (12.2) *is the global minimizer.*

Proof. I'll just sketch the proof. First, let $\boldsymbol{z}$ be the global minimizer of (12.2), and let $\boldsymbol{x}$ be any other vector. Consider moving $\boldsymbol{x}$ a little bit in the direction of $\boldsymbol{z}$, say to

$$\boldsymbol{x}_\epsilon \stackrel{\text{def}}{=} (1-\epsilon)\boldsymbol{x} + \epsilon \boldsymbol{z}.$$

The point $\boldsymbol{x}_\epsilon$ still has the correct values on the nodes in W. We can use an argument like the one from the previous proof to show that $\mathcal{E}(\boldsymbol{x}_\epsilon)$ is less than $\mathcal{E}(\boldsymbol{x})$ by an amount that is a linear function of ϵ. This follows from the identity

$$(1-\epsilon)a^2 + \epsilon b^2 - ((1-\epsilon)a + \epsilon b)^2 = \epsilon(1-\epsilon)(a-b)^2.$$

So, $\boldsymbol{x}$ cannot be a local minimizer of the energy if it is different from $\boldsymbol{z}$. Alternatively put, the partial derivative of the energy in the direction $\boldsymbol{z} - \boldsymbol{x}$ is negative. So, there must be some $u \in V - W$ such that the partial derivative of the energy with respect to $\boldsymbol{x}(u)$ is non-zero. □

If there is time, I will give one more concrete way of showing that the solution of the harmonic equations must be unique. We do it by taking the difference of the two harmonic functions. Such a function will be zero at each vertex in W. It then seems obvious that it should be zero everywhere. This is certainly obvious if you accept the physical spring model.

Lemma 12.4.3. *Let $\boldsymbol{x}$ and $\boldsymbol{y}$ be two functions that are both harmonic on $V - W$ and such that for all $w \in W$*

$$\boldsymbol{x}(w) = \boldsymbol{y}(w).$$

Then, for all $u \in V - W$,

$$\boldsymbol{x}(u) = \boldsymbol{y}(u).$$

Proof. Let $\boldsymbol{z} = \boldsymbol{x} - \boldsymbol{y}$. So, $\boldsymbol{z}(w) = 0$ for all $w \in W$, and we need to show that $\boldsymbol{z}(u) = 0$ for all $u \in V$.

Let v be such that $\boldsymbol{z}(u) \leq \boldsymbol{z}(v)$ for all $u \in V$. We have

$$\boldsymbol{z}(v) = \sum_{u:(v,u)\in E} \boldsymbol{z}(u) w_{u,v} / d(v).$$

As $\sum_u w_{u,v}/d(v) = 1$ and $w_{u,v} > 0$, the right-hand side is a weighted average of the values $\boldsymbol{z}(u)$. As it is also equal to the maximum of $\boldsymbol{z}$, all the values $\boldsymbol{z}(u)$ must be equal. So, $\boldsymbol{z}(u) = \boldsymbol{z}(v)$ for every neighbor u of v. By induction on paths in the graph, $\boldsymbol{z}(u) = \boldsymbol{z}(v)$ for every vertex u reachable from v. As this holds for some vertex $w \in W$, we have $\boldsymbol{z}(v) = \boldsymbol{z}(w) = 0$, and $\boldsymbol{z}(u) = 0$ for all vertices u that are reachable from v, which is all of V provided that the graph is connected. □

12.5 Resistor Networks

Given a graph, we can treat each edge as a resistor. If the graph is unweighted, we will assume that the resistor has resistance 1. If an edge e has weight $w(e)$, we will give the corresponding resistor resistance $r(e) = 1/w(e)$. The reason is that when the weight of an edge is very small, the edge is barely there, so it should correspond to very high resistance. Having no edge corresponds to having a resistor of infinite resistance.

The first equation I recall is

$$V = IR,$$

which says that the potential drop across a resistor is equal to the current flowing over the resistor times the resistance. To apply this in a graph, we will define for each edge (a, b) the current flowing *from a to b* to be $\boldsymbol{i}(a, b)$. As this is a directed quantity, we define

$$\boldsymbol{i}(b, a) = -\boldsymbol{i}(a, b).$$

I now let $\boldsymbol{v} \in \mathbb{R}^V$ be the vector of potentials at vertices. Given these potentials (voltages), we can figure out how much current flows on each edge by the formula:

$$\boldsymbol{i}(a, b) = \frac{1}{r_{a,b}} (\boldsymbol{v}(a) - \boldsymbol{v}(b)) = w_{a,b} (\boldsymbol{v}(a) - \boldsymbol{v}(b)).$$

I would now like to write this equation in matrix form. The one complication is that each edge comes up twice in $\boldsymbol{i}$. So, to treat $\boldsymbol{i}$ as a vector I will have each edge show up exactly once as (a, b) when $a < b$. I now define the *signed edge-vertex adjacency matrix* of the graph $\boldsymbol{U}$ to be the matrix with rows indexed by edges, columns indexed by vertices, such that

$$\boldsymbol{U}((a, b), c) = \begin{cases} 1 & \text{if } a = c \\ -1 & \text{if } b = c \\ 0 & \text{otherwise.} \end{cases}$$

Define $\boldsymbol{W}$ to be the diagonal matrix with rows and columns indexed by edges and the weights of edges on the diagonals. We then have

$$\boldsymbol{i} = \boldsymbol{W} \boldsymbol{U} \boldsymbol{v}.$$

Also recall that resistor networks cannot hold current. So, all the flow entering a vertex a from edges in the graph must exit a to an external source. Let $\boldsymbol{i}_{ext} \in \mathbb{R}^V$ denote the external currents, where $\boldsymbol{i}_{ext}(a)$ is the amount of current entering the graph through node a. We then have

$$\boldsymbol{i}_{ext}(a) = \sum_{b:(a,b) \in E} \boldsymbol{i}(a, b).$$

In matrix form, this becomes

$$\boldsymbol{i}_{ext} = \boldsymbol{U}^T \boldsymbol{i} = \boldsymbol{U}^T \boldsymbol{W} \boldsymbol{U} \boldsymbol{v}. \tag{12.4}$$

The matrix

$$\boldsymbol{L} \stackrel{\text{def}}{=} \boldsymbol{U}^T \boldsymbol{W} \boldsymbol{U}$$

will play a very important role in the study of resistor networks. It is called the Laplacian matrix of the graph.

To better understand the Laplacian matrix, let's compute one of its rows. We will do this using the equations we already have:

$$\begin{aligned}
\boldsymbol{i}_{ext}(a) &= \sum_{b:(a,b)\in E} \boldsymbol{i}(a,b) \\
&= \sum_{b:(a,b)\in E} \frac{1}{r_{ab}}(\boldsymbol{v}(a) - \boldsymbol{v}(b)) \\
&= \sum_{b:(a,b)\in E} w_{ab}(\boldsymbol{v}(a) - \boldsymbol{v}(b)) \\
&= d(a)\boldsymbol{v}(a) - \sum_{b:(a,b)\in E} w_{ab}\boldsymbol{v}(b).
\end{aligned}$$

This gives us the following expression of the entries of $\boldsymbol{L}$:

$$\boldsymbol{L}(a,b) = \begin{cases} d(a) & \text{if } a = b \\ -w_{a,b} & \text{if } (a,b) \in E \\ 0 & \text{otherwise.} \end{cases}$$

In matrix form, we see that $\boldsymbol{L}$ may be expressed as

$$\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{A},$$

where $\boldsymbol{D}$ is the diagonal matrix of weighted degrees and $\boldsymbol{A}$ is the weighted adjacency matrix.

12.6 Energy dissipation

Recall (from physics) that the energy dissipated in a resistor network with currents $\boldsymbol{i}$ is

$$\mathcal{E}(\boldsymbol{i}) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{(a,b)\in E} \boldsymbol{i}(a,b)^2 r_{a,b} = \frac{1}{2} \sum_{(a,b)\in E} \frac{1}{r_{a,b}}(\boldsymbol{v}(a) - \boldsymbol{v}(b))^2 = \frac{1}{2} \sum_{(a,b)\in E} w_{a,b}(\boldsymbol{v}(a) - \boldsymbol{v}(b))^2.$$

This expression should look familiar. Remember for later that it can never be negative.

Let's see that we can express this in terms of the Laplacian as well. Recall that $\boldsymbol{U}\boldsymbol{v}$ is a vector that gives the potential drop accross every edge. So,

$$(\boldsymbol{U}\boldsymbol{v})^T(\boldsymbol{U}\boldsymbol{v})$$

is the sum of the squares of the potential drops accross all the edges. This is almost the expression that we want. We just need to get in the weights. We do this in one of the following ways:

$$\mathcal{E}(\boldsymbol{i}) = (\boldsymbol{W}^{1/2}\boldsymbol{U}\boldsymbol{v})^T(\boldsymbol{W}^{1/2}\boldsymbol{U}\boldsymbol{v}) = (\boldsymbol{U}\boldsymbol{v})^T\boldsymbol{W}(\boldsymbol{U}\boldsymbol{v}) = \boldsymbol{v}^T\boldsymbol{U}\boldsymbol{W}\boldsymbol{U}\boldsymbol{v} = \boldsymbol{v}^T\boldsymbol{L}\boldsymbol{v}.$$

Let me mention some useful spectral properties of the Laplacian.

Theorem 12.6.1. *The Laplacian matrix of a graph is a positive semi-definite graph. If the graph is connected then the nullspace of its Laplacian is spanned by the constant vector.*

Proof. If λ is an eigenvalue $\boldsymbol{L}$ with eigenvector $\boldsymbol{v}$, the for the current $\boldsymbol{i} = \boldsymbol{W}\boldsymbol{U}\boldsymbol{v}$, we have

$$0 \leq \mathcal{E}(\boldsymbol{i}) = \boldsymbol{v}^T \boldsymbol{L}\boldsymbol{v} = \lambda \boldsymbol{v}^T \boldsymbol{v} = \lambda \|\boldsymbol{v}\|^2 .$$

So, $\lambda \geq 0$. It is clear from all of these definitions that $\boldsymbol{L}\mathbf{1} = \mathbf{0}$. The proof that $\mathbf{1}$ actuall spans the nullspace is completely analogous to the proof that the eigenvalue of a walk matrix has multiplicity 1. □

12.7 Fixing Potentials

We usually don't think of injecting current into a circuit. Rather, we usually attach nodes of a circuit to the terminals of a battery, which induce fixed potentials. Typically, we just attach two vertices to the terminals. Let's call them s and t, and assume that s has been attached to a terminal of potential 1 and that t has been attached to a terminal of potential 0. That is, $\boldsymbol{v}(s) = 1$ and $\boldsymbol{v}(t) = 0$. Then, the potentials all the remaining vertices can be determined by observing that for $x \notin \{s, t\}$, $\boldsymbol{i}_{ext}(x) = 0$ and applying (12.1) to show

$$d(x)\boldsymbol{v}(x) = \sum_{y:(x,y)\in E} \boldsymbol{v}(y)w(x, y). \tag{12.5}$$

Thus, we get another system of linear equations, with boundary conditions fixing the individual variables at s and t. We could also use linear algebra to prove that these equations have a solution. I'll give you a break and use probability instead. We will first show that the equations have a solution, and then prove that it is unique.

12.8 Solving for voltages

We now observe that we can always solve equation (12.4) for $\boldsymbol{v}$, provided that $\mathbf{1}^T \boldsymbol{i}_{ext} = 0$ and the underlying graph is connected. Typically, equations like $\boldsymbol{L}\boldsymbol{v} = \boldsymbol{i}_{ext}$ are solvable, except when the matrix is degenerate. Unfortunately, $\boldsymbol{L}$ does have determinant zero, and so is not necessarily invertible.

But, we know a lot about $\boldsymbol{L}$: its nullspace is spanned by $\mathbf{1}$. Equation (**??**) tells us that $\boldsymbol{i}_{ext}$ is orthogonal to the nullspace of $\boldsymbol{L}$, and so, because $\boldsymbol{L}$ is symmetric, $\boldsymbol{i}_{ext}$ is in the range of $\boldsymbol{L}$. So, we can solve (12.4) by just inverting $\boldsymbol{L}$ on its range–which is the space orthogonal to its nullspace. This will provide us with a solution $\boldsymbol{v}$ that is also orthogonal to the nullspace.

Here's a more explicit way of constructing the solution. Let $0 = \gamma_1 < \gamma_2 \leq \cdots \leq \gamma_n$ be the eigenvalues of $\boldsymbol{L}$ and let $\mathbf{1} = \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n$ be a corresponding orthonormal basis of eigenvectors. We have

$$\boldsymbol{i}_{ext} = \sum_{i=1}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right).$$

As

$$\left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right) = 0,$$

we can simplify this to

$$\boldsymbol{i}_{ext} = \sum_{i=2}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right).$$

So, we can just set

$$\boldsymbol{v} = \sum_{i=2}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right) / \gamma_i.$$

This will be a solution to (**??**) because

$$\begin{aligned}
\boldsymbol{L}\boldsymbol{v} &= \boldsymbol{L} \sum_{i=2}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right) / \gamma_i \\
&= \sum_{i=2}^{n} \gamma_i \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right) / \gamma_i \\
&= \sum_{i=2}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right) \\
&= \boldsymbol{i}_{ext}.
\end{aligned}$$

The bottom line is that the conventional wisdom that "one can only solve a system of linear equations if the matrix is invertible" is just too pessimistic. The more optimistic approach that we have taken here is to solve the system by multiplying by the pseudo-inverse instead of the inverse. Since you may find it useful later in life, let me say that the pseudo-inverse of a symmetric matrix is just the inverse on the range of the matrix. For a matrix with eigenvalues $\gamma_1, \ldots, \gamma_n$ and corresponding eigenvectors $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n$, it is given by

$$\boldsymbol{L}^+ \stackrel{\text{def}}{=} \sum_{i:\gamma_i \neq 0} \gamma_i \boldsymbol{u}_i \boldsymbol{u}_i^T.$$

12.9 Random Walks

Consider a random walk that starts at a vertex x, and stops when it finally reaches s or t. Let $F(x)$ be the probability that it hits s before t. It should be clear that F is well-defined, and that by definition $F(s) = 1$ and $F(t) = 0$.

Let's see what happens in between. For a vertex x that is neither s nor t, the probability that it steps to a node y is given by

$$w(x, y)/d(x).$$

We have

$$\begin{aligned} F(x) &= \sum_{y:(x,y)\in E} \Pr\left[\text{the walk goes to } y \text{ after } x\text{, and eventually stops at } s\right] \\ &= \sum_{y:(x,y)\in E} \Pr\left[\text{the walk goes to } y \text{ after } x\right] F(y) \\ &= \sum_{y:(x,y)\in E} F(y)w(x,y)/d(x). \end{aligned}$$

Multiplying through by $d(x)$, we get

$$d(x)F(x) = \sum_{y:(x,y)\in E} F(y)w(x,y).$$

That is, F is a solution to equation (12.5).

Let me point out that one can generalize this technique to find a solution when any number of vertices are fixed to any potentials. Let $W \subset V$ be the set of vertices whose potential we will fix. Consider fixing the potential of any element of W to 1 and the rest to 0. We can then enforce (12.5) at the rest of the vertices by considering the random walk that stops when it reaches any node of W, and considering the probability that it hits s first. Once you know how to build this function, you can build any function on W by taking a linear combination of such functions.