Harmonic Functions on Graphs

Taught by Yuval

Mathcamp 2016

1 Motivation

1.1 Random Walks

Let's say you're checked out to Walmart and have no idea how to get back to camp. Because of this, you pick the obvious strategy of walking back randomly: at each intersection, you roll a die and decide which direction to go based on the outcome. Since you want to be sure to get back by the end of sign-in, you want to know how long it'll take, on average, for this random walk to reach the dorms.

Example 1. Suppose Waterville looks like this:

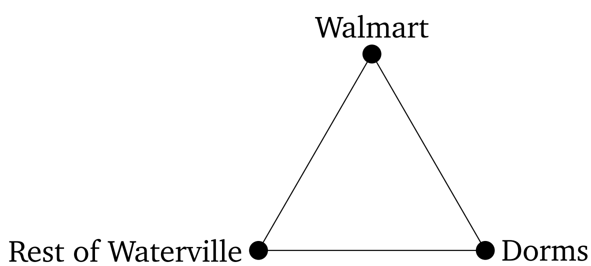

We can try to directly compute how long it'll take a random walk to get from Walmart to the dorms, on average. Before you take any steps, you're at Walmart with probability 1. After taking a single step, you're at the dorms with probability $\frac{1}{2}$ and at the rest of Waterville with probability $\frac{1}{2}$. After taking 2 steps, things are a bit more complicated, but we can figure it out: if you originally got to the dorms after one step, you stopped. If not, then you must have gotten to the rest of Waterville after 1 step, so after 2 steps you are at the dorms with probability $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$ and back at Walmart with probability $\frac{1}{4}$. Continuing this process, we can make the following table (which will look familiar if you were in Zach's colloquium in Week 1):

Probability of being at the given vertex

# of steps	Walmart	Rest of Waterville	Dorms
0	1	0	0
1	0	$\frac{1}{2}$	$\frac{1}{2}$
2	$\frac{1}{4}$	0	$\frac{1}{4}$
3	0	$\frac{1}{8}$	$\frac{1}{8}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
n	0 or 2^{-n}	0 or 2^{-n}	2^{-n}

Using this, we see that

$$\begin{aligned}
(\text{Expected time to get to dorms}) &= \sum_{n=0}^{\infty} n \cdot \text{Prob(you reach dorms at step } n) \\
&= \sum_{n=1}^{\infty} \frac{n}{2^n} \\
&= 2
\end{aligned}$$

So the answer is 2. However, this way of doing things was pretty terrible: it required us to recognize a pattern and sum a somewhat complicated infinite series.

We can do better. Let the *hitting time* $H(a,b)$ denote how long it takes, on average, for a random walk to go from vertex a to vertex b; with this notation, we are trying to calculate the hitting time $H(W,D)$, where W is Walmart and D is the dorms. The first step will take us to vertex D with probability $\frac{1}{2}$ and to vertex R with probability $\frac{1}{2}$, where R denotes the rest of Waterville. Therefore,

$$H(W,D) = 1 + \frac{1}{2}H(D,D) + \frac{1}{2}H(R,D)$$

where the addition of 1 records the fact that we took a single step to get to the position we're in. Note that $H(D,D) = 0$, since a random walk starting at the dorms will get to the dorms in zero steps. Similarly, we can see that

$$H(R,D) = 1 + \frac{1}{2}H(W,D) + \frac{1}{2}H(D,D) = 1 + \frac{1}{2}H(W,D)$$

Therefore, we can conclude that

$$\begin{aligned}
H(W,D) &= 1 + \frac{1}{2}H(D,D) + \frac{1}{2}H(R,D) \\
&= 1 + \frac{1}{2} \cdot 0 + \frac{1}{2}\left(1 + \frac{1}{2}H(W,D)\right) \\
&= \frac{3}{2} + \frac{1}{4}H(W,D)
\end{aligned}$$

We can rewrite this as

$$\frac{3}{4}H(W,D) = \frac{3}{2}$$

which we can solve to give us $H(W,D) = 2$. Note that we got the same answer as before, but that this solution style was much simpler; in particular, we used no infinite series.

Great! We now know how to calculate the hitting time, right? Let's apply our newfound knowledge to a different model of Waterville.

Example 2. Suppose Waterville looks like this:

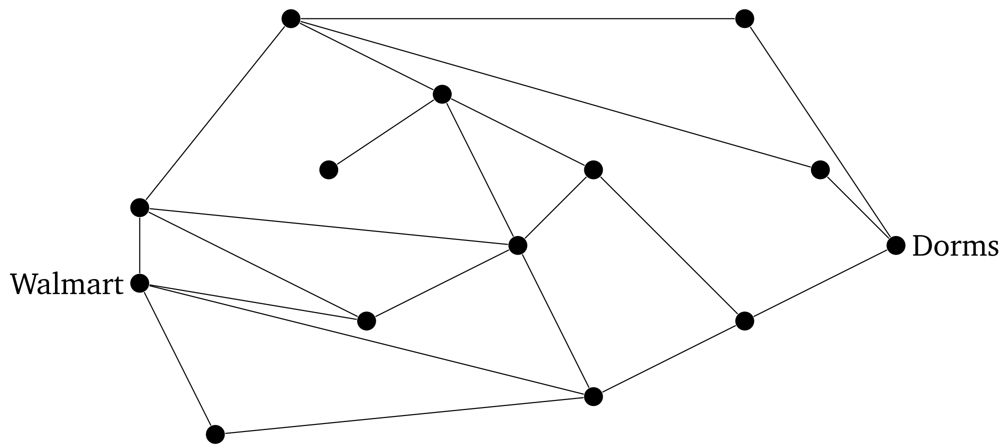

Um... I definitely don't want to sum an infinite series associated to this graph, and I also don't want to go through our second, recursive method of calculating $H(W,D)$. Is there a better way?

In order to try to find a better way, let's return to our second method once more. The basic insight we used there was that for any pair of vertices a, b in any graph,

$$H(a,b) = 1 + \frac{1}{\deg(a)} \sum_{c \in N(a)} H(c,b) \tag{1}$$

where $N(a)$ denotes the *neighborhood* of a, i.e. all those vertices adjacent to a, and $\deg(a) = |N(a)|$ is the degree of the vertex a. The reason that this equation holds is precisely the same as the reasoning we used in Example 1. Namely, after taking a single step from a, there is a $1/\deg(a)$ probability that we get to each of the neighbors of a, and from there the expected time to get to b is, by definition, $H(c,b)$. We'll refer to Equation 1 as the *neighbor-averaging property* for random walks: the hitting time from a can be determined from the average of hitting times over all the neighbors of a. This property will turn out to be the key to understanding random walks; are there other natural questions that exhibit a similar neighbor-averaging property?

1.2 Water flow

Suppose we have a network of water pipes modelled by a graph. It is a basic fact[1] from physics that the flow rate of water in a pipe is proportional to the pressure difference between the two

[1] This "fact" is not actually true, but it's a good approximation to true in many simple situations.

endpoints of the pipe. For any vertex v in the graph, let $P(v)$ denote the pressure at that vertex. Then what we have just said is that for any edge $e = uv$,

$$\text{(flow rate along } e\text{)} \propto P(v) - P(u)$$

In addition, note that the total flow into any vertex must equal the total flow out of that vertex, since it's impossible for water to build up at any point in the network. What that means is that if we add up the flow rates over all edges that touch a given vertex, we should get 0. Using our above equation, that tells us that for any vertex v,

$$\sum_{u \in N(v)} (P(v) - P(u)) = 0$$

We can equivalently write this as

$$\begin{aligned} 0 &= \sum_{u \in N(v)} (P(v) - P(u)) \\ &= \sum_{u \in N(v)} P(v) - \sum_{u \in N(v)} P(u) \\ &= |N(v)| P(v) - \sum_{u \in N(v)} P(u) \\ &= \deg(v) P(v) - \sum_{u \in N(v)} P(u) \end{aligned}$$

or equivalently

$$P(v) = \frac{1}{\deg(v)} \sum_{u \in N(v)} P(u)$$

Thus, pressure in a water pipe network also exhibits a neighbor-averaging property. Note that this neighbor-averaging property is slightly different from the one we saw for random walks, because of the $(1+)$ term we saw in equation 1.

1.3 Tiling Rectangles with Squares

A famous question that occupied mathematicians for some time is the question of when we can tile a rectangle by squares, and what sort of restrictions there are on the squares. One such tiling is:

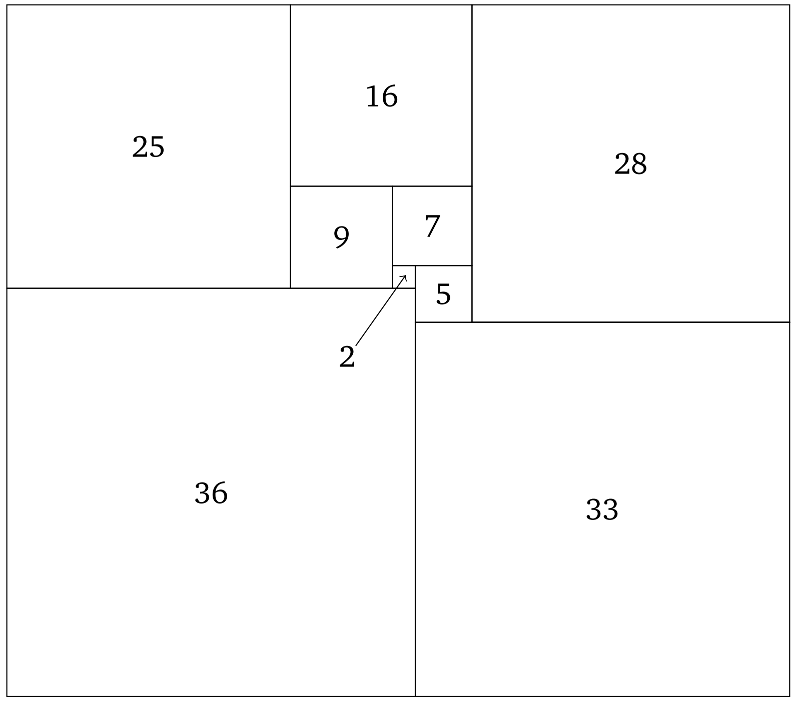

We can associate a graph to this tiling by giving a vertex to each maximal horizontal segment and connecting two such edges if and only if there is a square connecting them:

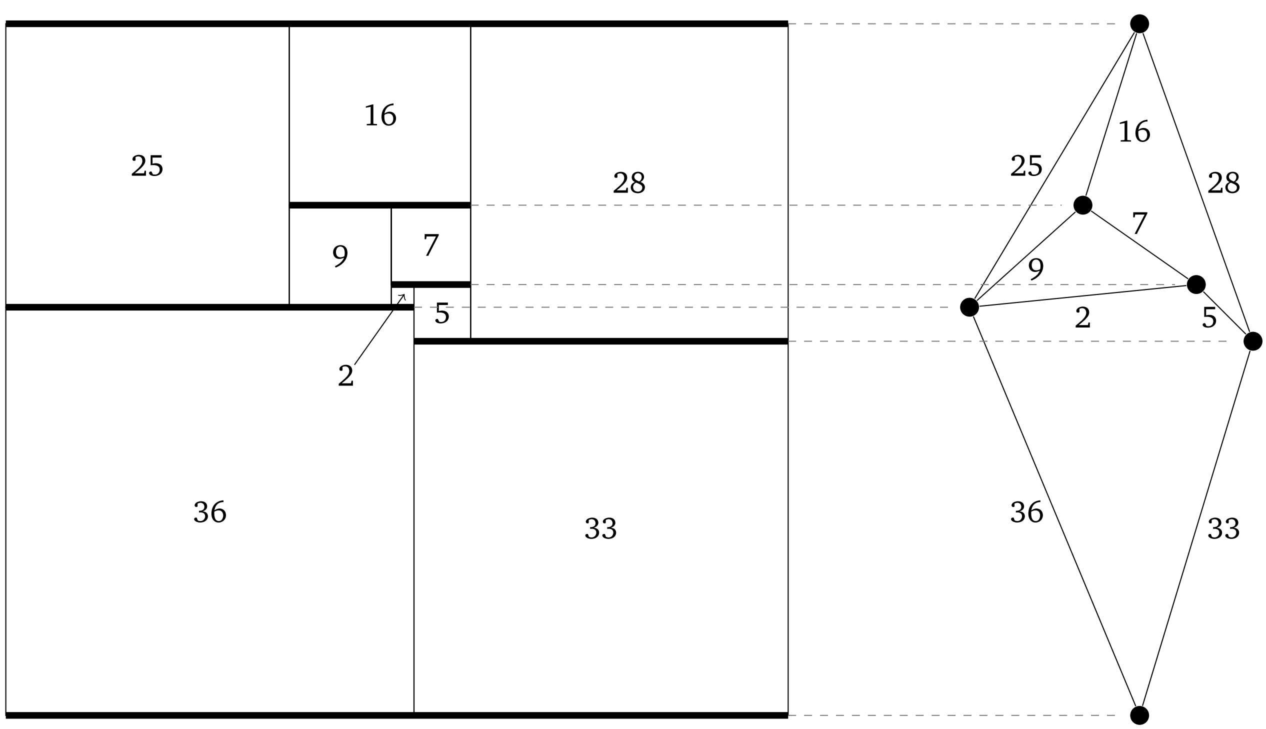

We define a function h on the vertices of this graph by setting $h(v)$ to be the height of vertex v. Then h also has a neighbor-averaging property. The reason is straightforward:

$$\begin{aligned}
\sum_{u\in N(v)} h(u) &= \sum_{\substack{u\in N(v)\\ u \text{ above } v}} h(u) + \sum_{\substack{u\in N(v)\\ u \text{ below } v}} h(u)\\
&= \sum_{\substack{u\in N(v)\\ u \text{ above } v}} (h(v) + \text{side length of square between } u \text{ and } v)\\
&\qquad + \sum_{\substack{u\in N(v)\\ u \text{ below } v}} (h(v) - \text{side length of square between } u \text{ and } v)\\
&= \deg(v)h(v) + (\text{length of segment defining } v) - (\text{length of segment defining } v)\\
&= \deg(v)h(v)
\end{aligned}$$

and dividing by $\deg(v)$ gives us the same neighbor-averaging property as we had for water pipes. Note that we have used the fact that squares are square: specifically, the sum of the side lengths of all the squares touching a horizontal segment from above is precisely the length of that horizontal segment, and similarly for those below.

1.4 Rubber bands

Our final example, and one of the most important ones we will consider, is that of rubber bands. You may have heard (in your physics class, for instance), of Hooke's law, which states that the pulling force an ideal rubber band exerts is proportional to the amount it is stretched. So suppose we replace all of the edges in our graph with ideal rubber bands, nail some of our vertices to the real line, and let the system find its equilibrium position. In equilibrium, the sum of all forces on a point must be zero (for otherwise, it would move); this means that if we let $f(v)$ denote the horizontal position of vertex v, then

$$\sum_{u\in N(v)} (f(u) - f(v)) = 0$$

since the force that the rubber band uv applies on v is proportional to its length, namely $f(u) - f(v)$. We saw this equation before; it is equivalent to

$$f(v) = \frac{1}{\deg(v)} \sum_{u\in N(v)} f(u)$$

And thus, in the equilibrium position, the location of vertices on a rubber band graph also demonstrates a neighbor-averaging property.

2 Harmonic Functions

2.1 Definitions and Disappointments

Let's formalize all of the motivation we did in the previous section.

Definition 1. Let $G = (V, E)$ be a graph. A function $f : V \to \mathbb{R}$ is said to be *harmonic* at $v \in V$ if it has a neighbor-averaging property at v, namely

$$f(v) = \frac{1}{\deg(v)} \sum_{u \in N(v)} f(u)$$

or equivalently

$$\sum_{u \in N(v)} (f(u) - f(v)) = 0$$

If f is harmonic at every $v \in V$, then we say that f is a *harmonic function* on G.

As we saw above, the following functions are harmonic: the pressure function on a network of water pipes, the height function on a square tiling of a rectangle, and the equilibrium position function of a rubber band graph. The hitting time isn't quite a harmonic function, thanks to that pesky (1+) term in Equation 1, but we will soon fix that.

Given all of these examples, the following theorem might be a bit surprising:

Theorem 2. *There are no (non-constant) harmonic functions on a (finite connected) graph G.*

Proof. Let $f : V \to \mathbb{R}$ be any non-constant function; we claim that it is not harmonic on all of V. The problem is that the function can't be harmonic at its maximum, since the average of its neighbors is at most the maximum value. However, this is not yet a proof, since it's possible that a maximal vertex has all its neighbors also be maximal. So to make this proof work, we argue as following. Let M be the set of all maximal vertices, i.e.

$$M = \{v \in V : f(v) = \max_{w \in V} f(w)\}$$

Since f is non-constant, $M \neq V$. In addition, since G is a connected graph, there is at least one $v \in M$ that is adjacent to some $v' \notin M$. Then

$$\begin{aligned}
\sum_{u \in N(v)} f(u) &= f(v') + \sum_{u \in N(v) \setminus \{v'\}} f(u) \\
&\leq f(v') + (|N(v)| - 1) \max_{w \in V} f(w) \\
&< |N(v)| \max_{w \in V} f(w) \\
&= \deg(v) f(v)
\end{aligned}$$

where the third line uses the fact that $v' \notin M$, so $f(v') < \max_{w \in V} f(w)$, and the last line uses the fact that $v \in M$, so $f(v) = \max_{w \in V} f(w)$. Dividing by $\deg(v)$ tells us that

$$f(v) > \frac{1}{\deg(v)} \sum_{u \in N(v)} f(u)$$

and thus f is not harmonic at v. $\square$

Note. Observe that the restrictions in the statement of Theorem 2 are crucial. First, any constant function is, in fact, harmonic on all of V. Second, if G is disconnected, then we can make a non-constant harmonic function: make it constant on one connected component and make it some other constant on the other connected component. On the homework, you'll see that the condition on G being finite is also necessary.

Since Theorem 2 guarantees that there are no interesting harmonic functions, it makes sense to give a name to those points on which a function is not harmonic:

Definition 3. For any function $f : V \to \mathbb{R}$, we say that $v \in V$ is a *pole* of f if f is *not* harmonic at v.

Corollary 4. *Every non-constant function f on a finite connected graph has at least two poles.*

Proof. In the proof of Theorem 2, we saw that at least one vertex maximizing f must be a pole of f. Repeating the exact same argument (but reversing all of our inequalities) shows that at least one vertex minimizing f must also be a pole of f. Since f is non-constant, its maximum and minimum must be different. Therefore, these two vertices are distinct, so f has at least two poles. □

Note. Note that I totally lied to you earlier, when I claimed that our motivating examples were harmonic functions. Specifically,

- For the height function of a square tiling, I swept the two poles under the rug: the top and bottom horizontal segments are poles of the function, and the argument I gave for harmonicity totally fails on them.
- In rubber band graphs, remember that we decided to nail some of the vertices to the real line; if we don't do this, all of the rubber bands will shrink to zero length and our entire graph will collapse to a point. However, those vertices nailed to the wall might be poles of the function, since the force-balancing argument does not apply to them.
- In the water pipe model, one of two things could go wrong. If the system is closed, then it will reach equilibrium and there will be no flow, implying that all pressures will be equal and thus our harmonic function will just be constant. If the system is not closed, then there must be water entering and exiting the system at some vertices, and therefore those vertices will be poles of the pressure function.

2.2 Rescuing Harmonic Functions

For a class called "Harmonic Functions on Graphs," Theorem 2 and Corollary 4 are a bit disappointing: if our function is even the tiniest bit interesting (i.e. not constant), then it cannot be harmonic. In other words, there are no interesting harmonic functions. What can we do?

It turns out that we can rescue this, and turn our no-longer-harmonic functions into extremely powerful mathematical tools. The key fact that enables this is the following theorem.

Theorem 5. *Let $G = (V, E)$ be a finite connected graph, and let $B \subseteq V$ be any set of vertices with $|B| \geq 1$. Let $f_0 : B \to \mathbb{R}$ be any function on B. Then there exists a unique function $f : V \to \mathbb{R}$ such that f is harmonic on every $v \in V \setminus B$ and $f(b) = f_0(b)$ for every $b \in B$.*

What this theorem says is that we can pick *any* set of vertices to be our poles and *any* assignment of values to those poles, and we get a *unique* function that extends that assignment and has its poles in our specified set. Such a function f is called the *harmonic extension* of f_0.

Proof of Uniqueness. Unlike many theorems, in this case, proving uniqueness is much easier than proving existence. For suppose we had two functions $f_1, f_2 : V \to \mathbb{R}$ such that f_1, f_2 are both harmonic on $V \setminus B$ and such that $f_1(b) = f_2(b) = f_0(b)$ for all $b \in B$. Then consider the function $g : V \to \mathbb{R}$ defined by $g(v) = f_1(v) - f_2(v)$. Then observe that for any $v \in V \setminus B$,

$$\begin{aligned} g(v) &= f_1(v) - f_2(v) \\ &= \frac{1}{\deg(v)} \sum_{u \in N(v)} f_1(u) - \frac{1}{\deg(v)} \sum_{u \in N(v)} f_2(u) \\ &= \frac{1}{\deg(v)} \sum_{u \in N(v)} (f_1(u) - f_2(u)) \\ &= \frac{1}{\deg(v)} \sum_{u \in N(v)} g(u) \end{aligned}$$

Thus, g is also harmonic on $V \setminus B$. By the proof of Corollary 4, we know that the maximum and the minimum of g cannot be achieved on the vertices where it is harmonic, so the maximum and minimum of g must both be achieved somewhere in B. However, we also know that for any $b \in B$,

$$g(b) = f_1(b) - f_2(b) = f_0(b) - f_0(b) = 0$$

Therefore, the maximum and minimum of g must both be 0, so g must be the constant zero function. That means that for any $v \in V$,

$$0 = g(v) = f_1(v) - f_2(v)$$

which implies that $f_1 = f_2$. So we do indeed have a unique harmonic extension. □

Proof of existence. First of all, note that the case $|B| = 1$ is already done; namely, a function $f_0 : B \to \mathbb{R}$ is simply a single real number $r \in \mathbb{R}$, and we can extend f_0 to the constant function f whose value is r everywhere. So from now on, we will assume that $|B| \geq 2$.

It turns out that all of those motivating examples we did are really useful for proving existence of harmonic extensions. Because of that, and because this theorem is so important, here are four distinct proofs of existence.

Rubber bands: As before, turn every edge of G into a rubber band. Nail every vertex in B to the wall, with the horizontal position of $b \in B$ given by $f_0(b)$. Then, let the system find its equilibrium position.[2] For $v \in V$, define

$$f(v) = x\text{-position of } v \text{ at equilibrium}$$

Then f certainly extends f_0, since the position of $b \in B$ is precisely where we nailed it down, namely $f_0(b)$. In addition, f is harmonic on $V \setminus B$, by the argument we made previously: for every unnailed vertex, the forces acting on it are balanced, which implies that

$$\sum_{u \in N(v)} (f(u) - f(v)) = 0$$

for any unnailed vertex v. This, as we saw, is equivalent to f being harmonic at v.

Pipes: This one is very similar to the rubber bands one. Make a network of pipes out of your graph, and insist that the pressure at every $b \in B$ be fixed at $f_0(b)$. We can make sure that this pressure stays constant by dynamically inserting or removing water from that node to keep the pressure fixed. Then, again by the same argument as before, we see that the pressure at equilibrium flow will be a harmonic function for all $v \in V \setminus B$.

Random walks: This proof is a bit more intricate than the two above, but it's also significantly less sketchy; in particular, it doesn't rely on any physical intuition or "facts" from physics. Define a function $f : V \to \mathbb{R}$ as follows: for $v \in V$,

$$f(v) = \sum_{b \in B} f_0(b) \operatorname{Prob}(b \text{ is the first vertex in } B \text{ that a random walk from } v \text{ reaches})$$

Then note that f does indeed extend f_0, since b is the first vertex in B that a random walk from b will reach. To prove that f is harmonic, let's first simplify notation: let

$$P(b, v) = \operatorname{Prob}(b \text{ is the first vertex in } B \text{ that a random walk from } v \text{ reaches})$$

so that

$$f(v) = \sum_{b \in B} f_0(b) P(b, v)$$

[2] You can use your physical intuition to convince yourself that this equilibrium exists. For a formal proof, you can consider the energy function and use a multivariable calculus argument to see that it has a unique minimum value. Finally, you can show that at the minimum energy position, all forces are balanced.

Then for $v \in V \setminus S$

$$\begin{aligned}
f(v) &= \sum_{b\in B} f_0(b)P(b,v) \\
&= \sum_{b\in B} f_0(b) \sum_{u\in N(v)} \text{Prob}(u \text{ is the first step of the walk}) \cdot P(b,u) \\
&= \sum_{b\in B} \sum_{u\in N(v)} \frac{1}{\deg(v)} f_0(b)P(b,u) \\
&= \frac{1}{\deg(v)} \sum_{u\in N(v)} \sum_{b\in B} f_0(b)P(b,u) \\
&= \frac{1}{\deg(v)} \sum_{u\in N(v)} f(u)
\end{aligned}$$

which is precisely the harmonicity property. Thus, f is indeed a harmonic extension of f_0.

Linear Algebra: This paragraph is not actually a proof, just a statement that such a proof does exist. One can prove the existence of harmonic extensions using some general techniques from linear algebra and spectral graph theory, but the proof is complicated and unenlightening. In my opinion, the above proofs are the “right” way to think about harmonic extensions: they tell you where the harmonic extension is actually coming from, rather than a general abstract proof that it should exist.

□

3 Using Harmonic Functions

It turns out that the uniqueness of harmonic extensions will be very helpful for us soon. In particular, recall that all of our motivating examples give us harmonic functions, so if harmonic extensions are unique, we must get the *same* harmonic function in all cases. Using this, we can get the following result, which is just one of many like it.

Theorem 6. *Suppose we have a square tiling of a rectangle R, and let $G = (V, E)$ be the associated horizontal-segment graph. Let s denote the bottom edge of the rectangle and t the top edge. Then the ratio between the width and the height of R equals the force required to hold s and t one unit distance apart when all the edges of G are rubber bands.*

Proof. Let $f_0(s) = 0, f_0(t) = 1$, and f the unique harmonic extension of f_0 to V. Then consider

$$\sum_{u\in N(s)} f(u)$$

On the one hand, since the height of square tilings is harmonic, this is the sum over all bottom squares of their heights, which is the same as the sum of their widths, which is just the width of R, given that the height of R is 1. On the other hand, suppose we hold s at position 0 and t at position 1. Then since rubber band placement is harmonic, the force applied on s is also

$$\sum_{u \in N(s)} f(u)$$

So these two are equal. □

This is actually a special case of a much more general theorem, which we won't prove:

Theorem 7. *Let $G = (V, E)$ be a graph, and let $s, t \in V$. Define the following quantities:*

- *F_{st} is the force required to hold s and t distance 1 apart when all edges are rubber bands.*
- *W_{st} is the width of a rectangle when square-tiled according to G with s as the bottom edge and t as the top edge, and height 1.*
- *C_{st} is the commute time between s and t, which is the expected amount of time it'd take a random walk to start at s, reach t, and then return to s.*
- *R_{st} is the effective resistance between s and t when all edges are 1 Ohm resistors (ignore this if you haven't seen resistor diagrams previously).*
- *G' is the graph gotten by identifying vertex s with vertex t, and $T(G)$ denotes the number of spanning trees in G.*

Then

$$F_{st} = W_{st} = \frac{1}{R_{st}} = \frac{2|E|}{C_{st}} = \frac{T(G)}{T(G')}$$

Idea of Proof. For every one of these quantities, it is possible to associate a function to G that is 0 at s, 1 at t, and harmonic everywhere else. Since harmonic extensions are unique, all of these functions must be equal. From that, we can deduce all of these equalities. □

In a very similar vein, we can relate hitting times—our original motivating question—to rubber bands, albeit not exactly through harmonic functions.

Theorem 8. *Pick a vertex b of a graph $G = (V, E)$ and nail it to the wall. To every other vertex x, attach a weight equal to $\deg(x)$. Finally, replace all edges with rubber bands. Let the system find its equilibrium position. Then*

$$H(a, b) = \textit{height different between } a \textit{ and } b$$

Proof. As we saw at the beginning, we have the following recurrence for $H(a, b)$:

$$\begin{aligned} H(a,b) &= 1 + \frac{1}{\deg(a)} \sum_{v \in N(a)} H(v,b) \\ &= \frac{1}{\deg(a)} \left(\deg(a) + \sum_{v \in N(a)} H(v,b) \right) \end{aligned}$$

We can rewrite this as

$$-\deg(a)+\sum_{v\in N(a)}(H(a,b)-H(v,b))=0$$

Now, suppose we we define a function h on the vertices by $h(b)=0$, and $h(x)$ is the equilibrium height of x when all edges are rubber bands and each vertex x has a weight of $\deg(x)$ attached to it. At equilibrium, the sum of the forces on each vertex (apart from b) will be 0. The forces on a given vertex a are a force of $\deg(a)$ pointing down (from the weight attached to a), and a force of $h(a)-h(v)$ for each neighbor v of a (from the rubber band). The fact that these forces are balanced means that

$$-\deg(a)+\sum_{v\in N(a)}(h(v)-h(a))=0$$

Since this is (up to a sign) the same equation as we got above for $H(a,b)$, we see that $h(a)=-H(a,b)$. Since b was nailed at 0, this means that the distance between a and b in this equilibrium position is exactly $H(a,b)$, as desired. $\square$

Note. Strictly speaking, this theorem does not use the uniqueness of harmonic extensions, since the hitting time is not actually harmonic. However, in its use of a neighbor-averaging property and in its close connection to rubber bands (which, we will see, are in some sense the prototypical example of a harmonic function), it is certainly very similar to other harmonic function arguments.

Observe the following fact about Theorem 5: we can write down the following sets of equations

$$\begin{aligned} f(v)&=f_0(v) && \forall v\in B\\ f(v)&=\frac{1}{\deg(v)}\sum_{u\in N(v)}f(u) && \forall v\in V\setminus B\end{aligned}$$

If we treat this as a set of linear equations with variables $f(v)$ for all $v\in V$ (and treat everything else, including the values of f_0, as constants), then we get $|V|$ linear equations and $|V|$ variables. As we know from basic (linear) algebra, if this system of equations has a solution, then we can find it very efficiently using row-reduction.[3] What this means is that in addition to being powerful theoretical tools, harmonic extensions are actually useful in practice: we can find their values very quickly. Note that this argument only works because we already know that this system has a solution.

Another important observation from this system of equations is that all of the coefficients of the equations are rational numbers, assuming that f_0 takes rational values. This implies that

Theorem 9. *If $f_0: B\to\mathbb{Q}$ is a rational-valued function, then its unique harmonic extension f is also rational-valued.*

[3] In algorithmic language, we can find a solution in time $O(|V|^3)$.

Proof. Row-reducing a system of equations is done only by adding, subtracting, multiplying, and dividing by coefficients, so if all the coefficients start out rational, they will never become irrational. □

This simple observation turns out to be enormously powerful. For instance, it implies that the hitting time between two vertices on any graph will always be rational, by Theorem 8. This is pretty surprising when you think back to our original method for calculating hitting time, which involved summing some complicated infinite series. It's totally not obvious that all of these series will always have a rational number as their sum, but Theorem 9 guarantees it. Another important consequence of Theorem 9 is the following famous theorem:

Theorem 10 (Dehn)**.** *A rectangle R can be tiled by squares if and only if the ratio of the side lengths of R is rational.*

Proof. By rescaling, we may assume that R is a $1 \times x$ rectangle. For one direction, suppose that $x = a/b \in \mathbb{Q}$. Then we may divide the rectangle into a grid of squares whose side lengths are all $1/b$, and we are done.

For the converse, suppose we have some tiling of R by squares. Let G be the associated horizontal-segment graph, and let h be the height function on G, which we know is harmonic everywhere but the top and bottom of the rectangle. Moreover, at the two poles, the values of h are 0 and 1, which are rational. So by Theorem 9, we know that h takes on only rational values. Applying this to the bottom-most squares immediately tells us that they all have rational side length. However, this implies that the total length of the bottom segment, which is the sum of all the bottom-most side lengths, must be rational as well. But that length is precisely x, so $x \in \mathbb{Q}$, as desired. □

4 Straight-Line Embeddings

Using harmonic functions, we can prove the following famous theorem:

Theorem 11. *Let G be a planar graph. Then G can be drawn in the plane such that all the edges are straight lines and no pair of edges intersect.*

We will prove this in several steps.

Definition 12. A graph G is called 3-connected if, for any pair of vertices, deleting them does not disconnect the graph. Equivalently, for any subset $S \subseteq V$, there are at least three edges out of S.

Proposition 13. *Let G be a planar graph. We can add edges to G to get a new graph H such that H is still planar, and is also* 3*-connected.*

Proof. We add edges by triangulating each face of the graph until we can't anymore. On your homework, you'll show that this is indeed forms a 3-connected graph. □

Because of this fact, it suffices to prove that any 3-connected planar graph can be embedded in the plane with straight-line edges. For if we add edges to G, then straight-line embed it, then delete all the edges we added, we will clearly end up with a straight-line embedding. We will explicitly construct such an embedding:

Definition 14. The *Tutte embedding* of a 3-connected planar graph G is defined as follows. Pick some face F of G, and nail it to the plane as some convex polygon. Replace all edges with rubber bands and let the system find its equilibrium position. Then the location of each vertex in this equilibrium state is its Tutte embedding position. For a vertex v, we denote this position by $T_v \in \mathbb{R}^2$. Then, we embed an edge uv as the straight line segment between T_u and T_v.

Our goal is now to prove that this Tutte embedding has no pair of edges intersect.

Lemma 15. *In the Tutte embedding, both the x-coordinate and the y-coordinate of T_v form harmonic functions on $V \setminus F$. In other words, for any non-nailed vertex, its x-position and y-position are the average of those of its neighbors.*

Proof. This works by the same force-balancing argument as we originally used on rubber bands. Separate the force on any vertex into its x-component and its y-component. Both of these are proportional to the x-separation and the y-separation, respectively, of v from its neighbors, so we get harmonicity. □

All the difficulty in proving that the Tutte embedding has no intersecting edges lies in the following lemma:

Lemma 16. *Let ℓ be any line intersecting the Tutte embedding. Let U denote the set of vertices that end up on one given side of ℓ. Then the vertices in U induce a connected subgraph of G. In other words, ℓ cuts G into two connected subgraphs.*

Proof. Begin with the nodes of F, i.e. those vertices that we nailed down. Since they were nailed as the vertices of a convex polygon, all the ones in U form a path, and thus are connected. Call this path P. We will show that every vertex in U can be connected to one of the vertices in P, which means that U induces a connected subgraph.

By slightly perturbing the line ℓ (so that the set U is unchanged), we may assume that ℓ is not parallel to any edge. By rotating the entire picture, we may also assume that ℓ is a horizontal line and U is the half above ℓ. Fix some $v \in U$, and consider T_v. We know that the y-position of T_v is the average of the y-positions of its neighbors. We split into two cases:

Case 1: Suppose there exists some neighbor $u \in N(v)$ such that $T_u \neq T_v$; in other words, not the entire neighborhood of v was mapped to the same point (this is what will happen in general). Then, in particular, there is some neighbor v_1 such that $y(T_{v_1}) > y(T_v)$ (by the averaging property and the fact that no edge is parallel to ℓ). Iterate this argument at v_1 to get one of its neighbors v_2 with $y(T_{v_2}) > y(T_{v_1})$. Repeat this to get a path of vertices, each connected to the next via an edge, whose y-coordinates increase. Since G is finite, eventually this process must stop, but that can only happen when we reach a nailed vertex (where the averaging property is false). Then this nailed vertex must be in P, so we have connected v to P, as desired.

Case 2: In this case, every neighbor u of v has the property $T_u = T_v$. Then let H denote the connected subgraph of G consisting of all vertices located at T_v. Then either H contains a nailed vertex, in which case its connectivity implise that v is connected to P. If not, then it is connected to some vertex not in H, and we can repeat the above argument again, since one of these neighbors must be strictly higher than T_v.

□

We will be repeatedly using Lemma 16 to prove that the Tutte embedding has no intersecting edges. Before doing so, we need to deal with a few possible degenerate cases.

Definition 17. Call a vertex v *degenerate* if there is some line ℓ such that $T_v \in \ell$ and $T_u \in \ell$ for all $u \in N(v)$.

Lemma 18. *There are no degenerate vertices in the Tutte embedding.*

Proof. Suppose for contradiction that there were, and fix some $v \in V$ and some line ℓ so that v and all its neighbors are placed on ℓ. Let H denote the connected subgraph of G consisting of all vertices on ℓ whose neighbors are also all on ℓ; so H contains v, and it might contain some of its neighbors, and might contain some of its neighbors' neighbors, and so on.

By 3-connectivity, H has at least 3 edges connecting it to the rest of G. Let N denote the set of neighbors of H outside H itself. Then the above just says that $|N| > 3$. Now, let U_1 and U_2 be the two sets of vertices lying on either side of ℓ; by Lemma 16, we know that both U_1 and U_2 induce connected subgraphs. We claim that for every $a \in N$, a is adjacent to a vertex in U_1 and a vertex in U_2.

By the definition of H and N, we know that each $a \in N$ has $T_a \in \ell$. On the other hand, since $a \notin H$, a must have a neighbor not on ℓ. Call this neighbor b, and without loss of generality, $b \in U_1$. If a is not a nailed-down vertex, then a's position is the average of its neighbors, so it must also have at least one neighbor in U_2. If a is nailed down, then the two adjacent vertices nailed down will be on opposite sides of ℓ.

Now, recall that H is a connected subgraph by definition, and U_1, U_2 induce connected subgraphs by Lemma 16. So we can contract all of the edges in these subgraphs to get three super-vertices, each of which is connected to every vertex in N. This means that we have found a $K_{3,3}$ embedded in G. Since $K_{3,3}$ is not planar, G cannot be planar either, a contradiction. □

Lemma 19. *Let ab be an edge of G that is not an edge of the face F. Let F_1, F_2 denote the two faces of G that contain ab, and let ℓ denote the line defined by ab. Then all the vertices of F_1 lie on one side of ℓ, and all the vertices of F_2 lie on the other side of ℓ (apart from a and b themselves).*

Proof. Again, by rotating, we may assume that ℓ is a horizontal line. For contradiction, suppose that there was some vertex $c \in F_1$ that lies on the same side of ℓ as some vertex $d \in F_2$; without loss of generality, assume these are both on the upper side of ℓ, or else at least one is on ℓ itself. In this latter case, Lemma 18 guarantees that the vertex on ℓ is not degenerate, so it has at least one neighbor on the upper side of ℓ. By Lemma 16, we know that the upper side induces a connected subgraph, so we regardless get a path connecting c to d that lies entirely on the upper half (either directly or to the neighbor of the vertex that lies on ℓ); call this path P. In addition, since the lower half of ℓ also induces a connected subgraph, we know that there is a path connecting a and b that uses only the lower half; call this path P'. Since one lies entirely on the top half and the other entirely on the bottom half, they must be disjoint.

Now go back to the original planar drawing of G, which we know exists. It must look something like this:

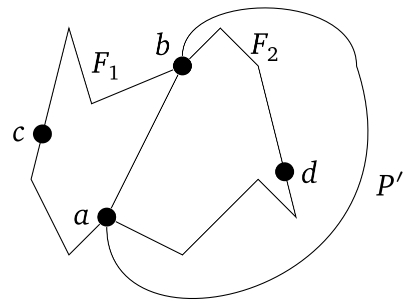

The reason is that we know that F_1, F_2 are faces, so the path P' connecting a and b cannot go through their interiors. But note that in this picture, c and d are separated by $P' \cup \{ab\}$, so the supposedly disjoint path connecting them, P, cannot exist.

This is the contradiction we were looking for, so F_1 and F_2 must indeed be on opposite sides of ℓ in the Tutte embedding. □

Lemma 20. *In the Tutte embedding, the interiors of the faces are disjoint.*

Proof. Pick some point $z \in \mathbb{R}^2$; we want to prove that z is in the interior of at most one face. Since we care about interiors of faces, we may assume that z is not a vertex or an edge in the embedding. Draw a line ℓ going through z that does not go through any T_v for any vertex v. For some faraway points on ℓ, they are not in the interior of any face. As we move towards the embedded points and enter the polygon F given by the nailed vertices, the number of faces whose interior we're in increases to 1. As long as we stay within a face, this number stays at 1, so the only thing that could go wrong is when ℓ crosses an edge. However, by Lemma 19, every time we cross an edge, we exit one face and enter a new face. So along ℓ, the number of faces whose interior we're in stays constant at 1. We eventually reach z, so z is in the interior of exactly one face. □

This is the final piece, and we can now prove Tutte's Theorem.

Theorem 21 (Tutte)**.** *The Tutte embedding is a straight-line embedding with no edges intersecting.*

Proof. Suppose we had two edges intersecting at some point z. Consider the four faces on either sides of these two edges. Some point z' very near z must lie in the interiors of two such faces; by Lemma 20, this can only happen if those two faces are actually the same face. But in that case, we see that our two edges were actually two adjacent edges on the same face, and their intersection must be at a vertex.

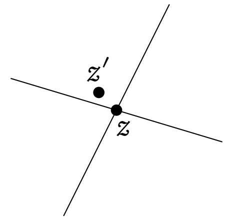

□

Looking back at the proof, we did various complicated steps involving planarity and 3-connectivity; however, the key result, that we used again and again, is Lemma 16: that a line intersecting the Tutte embedding divides the graph into two connected components. This result is the one that fundamentally depended on the averaging property of rubber bands, or equivalently on the harmonicity of the Tutte embedding. That property is what makes the Tutte embedding special, and what makes this proof work.

As a final note, observe that the Tutte embedding can be computed efficiently, since the x and y coordinate of T_v are simply harmonic extensions of the coordinates of the nailed vertices. As we saw, calculating harmonic extensions can be done very efficiently (since it involves solving a system of linear equations), and this implies that Tutte's theorem gives us an efficient algorithm for finding a straight-line embedding, in addition to proving that one exists.

5 Where to go from here

If you want to learn more about any of the topics in this course, a good place to check is László Lovász's book *Geometric Representations of Graphs*, available for free on his website at `http://www.cs.elte.hu/~lovasz/geomrep.pdf`; the book also has an extensive bibliography, if you want to delve even deeper. The paper where most of the square-tiling theorems and examples came from is "The Dissection of Rectangles into Squares," by Brooks, Smith, Stone, and Tutte.