The Unit Distance Graph and AC	Instructor: Padraic Bartlett
Lecture 1: Infinite Graphs	
Week 5	*Mathcamp 2012*

Consider the following model for creating a "random" graph on n vertices:

- Take n vertices, and label them $\{1, \ldots n\}$.
- For each unordered pair of vertices $\{a, b\}$, flip a coin that comes up heads 1/2 of the time and tails otherwise. If it comes up heads, connect these vertices with an edge; otherwise, do not.

This model for "random" graphs has a number of properties. Amongst other things, we can talk about how "likely" it is that a random graph on n vertices possesses a given property, like "there is a triangle" or "there are no edges in the entire graph."

For example, we can easily describe the likelihood of getting a graph on n vertices with no edges: it's just the probability that every time we flipped a coin in our model, it came up tails. There are as many edges in our graph as there are unordered pairs of vertices in our graph: i.e. $\frac{n(n-1)}{2}$, which you can see by thinking of how many ways you have to choose the first vertex (n), then choosing the second vertex ($n-1$), and then dividing by 2 because we don't care about order. Therefore, the odds of getting such a graph are

$$\left(\frac{1}{2}\right)^{\frac{n(n-1)}{2}},$$

which is vanishingly small for large values of n.

With this as motivation, consider the following property:

Definition. Let (‡) denote the following graph property: we say that a graph G satisfies the property (‡) iff for any pair of finite disjoint subsets $U, W \subset V(G)$, there is some $v \in V(G)$, $v \notin U \cup W$, such that v has an edge to every vertex in U and to no vertices in W.

What kinds of graphs satisfy this property? Well, no finite graph does: simply take $U =$the entire graph and $V = \emptyset$. Then there is no vertex that is not in $U \cup W$, and therefore the above property fails.

But what if we looked at infinite graphs: could we satisfy this property? Relatedly, suppose we studied a "random" graph on $\mathbb{N}$-many vertices: i.e. take $\mathbb{N}$ as your vertex set, for each pair of natural numbers flip a coin, and put an edge between those two elements if and only if your coin comes up heads. How likely is a graph to satisfy this property?

Theorem 1 *If G is a random graph on $\mathbb{N}$ that's generated using the model described above, then G satisfies property (‡) with probability 1 (i.e. the probability that G does **not** satisfy (‡) is 0.)*

Proof. Choose any pair of finite disjoint subsets U, W in $V(G)$. Pick any vertex $v \in V(G), v \notin U \cup W$, and let A_v be the event that v is connected to all of U and none of W. If we let $Pr(A_v)$ denote the probability that A_v occurs, we can easily see that

$$Pr(A_v) = \left(\frac{1}{2}\right)^{|U|} \cdot \left(\frac{1}{2}\right)^{|V|}.$$

The probability that A_v doesn't happen plus the probability that A_v **does** happen must sum to 1 (because we clearly have only two possible outcomes: either A_v does not happen or A_v happens.) Therefore, we then have

$$Pr(\text{not } A_v) = 1 - \left(\frac{1}{2}\right)^{|U|} \cdot \left(\frac{1}{2}\right)^{|V|} < 1.$$

Thus, we know that the probability of k different vertices $v_1, \ldots v_k$ all failing to satisfy A_v is just raising this quantity to the k-th power. Because the quantity above is < 1, taking k-th powers makes this go to zero as k increases! Therefore, for any U, W, we can specifically bound the chances that A_v fails for all of the vertices $v_1, \ldots v_k$ above by above by ϵ, for any $\epsilon > 0$, by simply looking at enough of these vertices $v_1, \ldots v_k$.

Now, notice that there are only countably many pairs of finite disjoint subsets of $\mathbb{N}$. To see why, first notice that for any k, the set of subsets of $\mathbb{N}$ of size k is a countable set: you can prove this using the same method as we did in proof techniques to show that there are countably many subsets of $\mathbb{N}^2$, i.e. by plotting them all as points in $\mathbb{N}^k$, drawing a spiral that starts at the origin and goes through each point, and sending the m-th natural number to the m-th point we hit on our spiral. To extend this to our claim, all we have to do is show that the union of countably many countable sets is countable: to do this, think of each of our countable sets as a copy of $\mathbb{N}$, which we can do because there's a bijection between each countable set and $\mathbb{N}$. Therefore, we can interpret the disjoint union of these countable sets as just $\mathbb{N}^2$, where the first coordinate tells us which countable set we're in and the second coordinate is telling us which element we have in our countable set. $\mathbb{N}^2$ is countable, by the spiral argument we gave above; therefore, the entire collection of these finite subsets is countable! (And therefore pairs of them are also countable, via the same logic.)

Consequently, we can enumerate all such pairs in a list $\{(U_i, W_i)\}_{i=1}^{\infty}$. For each one of these pairs, we proved earlier that we can bound the probability that there is no vertex that hits all of U_i and none of V_i above by any arbitrarily small number that we want. So: pick any $\epsilon > 0$, and bound the probability that (U_i, W_i) does not have a vertex that hits all of U_i and none of W_i by $\epsilon/2^i$, for every i. Then, the probability of any one of these events failing is bounded above by the sum

$$\sum_{n=1}^{\infty} \frac{\epsilon}{2^n} = \epsilon \left(\sum_{n=1}^{\infty} \frac{1}{2^n} \right) = \epsilon.$$

Therefore, the probability that none of these events fails is bounded below by $1 - \epsilon$! If none of these events fail, then our graph satisfies (‡): therefore, we've just shown that almost every random graph satisfies property (‡).

In the light of the above comments, it's interesting to note that we can actually construct a graph that satisfies (‡)! In fact, consider the following construction:

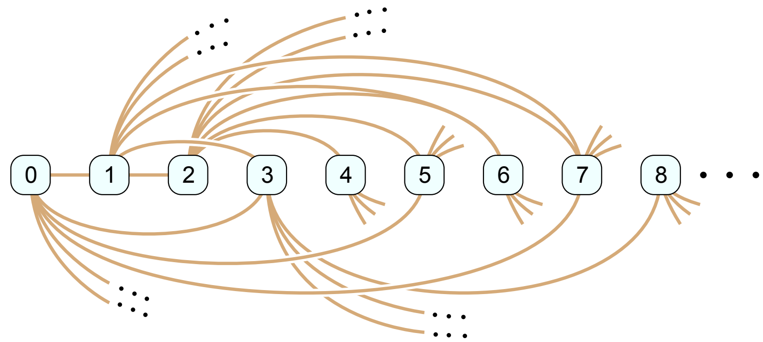

- Start by defining $R_0 = K_1$, the graph with a single vertex.
- If R_k is defined, let R_{k+1} be defined by the following: take R_k, and add a new vertex v_U for every possible subset U of R_k's vertices. Add an edge from v_U to every element in U, and to no other vertices in R_k.
- Let $R = \cup_{k=1}^{\infty} R_k$.

We claim that R is a graph on $\aleph_0$-many vertices that satisfies property (‡). To see why: pick any two finite disjoint subsets U, V of $V(R)$. Because each vertex of R lives in some R_k, we know that there is some finite value n such that U, V are both subsets of R_n, as there are only finitely many elements in $U \cup V$. Then, by construction, we know that there is a vertex v_U in R_{n+1} with an edge to every vertex in U and to none in V.

This graph is known as the Rado graph, and it has the following remarkable property:

Proposition 2 *The Rado graph is the only graph on $\aleph_0$-many vertices, up to isomorphism*[1]*, that satisfies* (‡).

Proof. To see this, take any two graphs $G = (V, E_G), H = (W, E_H)$ on $\aleph_0$-many vertices that satisfy (‡). We will create an isomorphism ϕ" $V \to W$ between these two graphs..

To do this: fix some ordering $\{v_i\}_{i=1}^{\infty}$ of V's vertices. Similarly, order W's vertices as $\{w_i\}_{i=1}^{\infty}$. We start with our isomorphism $\phi : V \to W$ undefined for any values of V, and construct ϕ via the following back-and-forth process:

- At odd steps:
  - Let v be the first vertex under V's ordering that we haven't defined ϕ on.
  - Let U be the collection of all of v's neighbors in V that we currently **have** defined ϕ on.
  - By (‡), we know that there is a $w \in W$ such that w is adjacent to all of the vertices in $\phi(U)$, and is also not adjacent to any other vertices that we have mapped to with ϕ. (We can apply (‡) because both of these sets are finite.)
  - Set $\phi(v) = w$.

[1]An isomorphism of two graphs $G = (V, E_G), H = (W, E_H)$ is a bijection $\phi : V \to W$ such that $\{u, v\}$ is an edge in E_G if and only if $\{\phi(u), \phi(v)\}$ is an edge in E_H.

- At even steps: do the exact same thing as above, except switch V and W! I.e.
  - Let w be the first vertex under W's ordering that we haven't yet mapped to with ϕ.
  - Let U be the collection of all of w's neighbors in W that we currently **have** mapped to with ϕ.
  - By (‡), we know that there is a $v \in V$ such that v is adjacent to the set $\phi^{-1}(U)$ made of the vertices in V that map to U, and v is not adjacent to any other vertices that we have defined ϕ on. (Again, we can apply (‡) because both of these sets are finite.)
  - Set $\phi(v) = w$.

So, in other words, we're starting with ϕ totally undefined; at our first step, we're then just taking ϕ and saying that it maps $v_1 \in V$ to some element in V'. Then, at our second step, we're taking the smallest element in V' that's not $\phi(v_1)$, and mapping it to some element w that either does or does not share an edge with v, depending on whether $\phi(w)$ and $\phi(v)$ share an edge.

By repeating this process, we eventually get a map that's defined on all of V, V'; we claim that such a map is an isomorphism. It's clearly a bijection, as it hits every vertex exactly once by definition. Therefore, it suffices to prove that it preserves edges: i.e. that $\{u, v\}$ is an edge in E_G if and only if $\{\phi(u), \phi(v)\}$ is an edge in E_H.

To see why this is true, take any pair of vertices u, v in V. Assume without any loss of generality that ϕ was defined on u before it defined on v (one of them has to be defined first, so it might as well be u.) Then, when we defined $\phi(v)$, there were only two ways we went about doing it:

- We defined $\phi(v)$ at an odd stage. In this case, when we defined $\phi(v)$, we defined $\phi(v)$ so that it would only be adjacent to the image under ϕ of all of v's neighbors that have already been defined! In particular, this means that we defined $\phi(v)$ to be adjacent to $\phi(u)$ if and only if $\{u, v\}$ was an edge in E_G.

- We defined $\phi(v)$ at an even stage. In this case, we picked v so that it would only be adjacent to every element in

  $$\phi^{-1}(\text{elements currently mapped to by } \phi \text{ that are neighbors of } \phi(v)).$$

  But this means that v is adjacent to $\phi^{-1}(\phi(u)) = u$ if and only if $\{\phi(u), \phi(v)\}$ is an edge in W! So, because $\{u, v\}$ is an edge, so is $\{\phi(u), \phi(v)\}$.

Therefore, we have that ϕ is an isomorphism.

Finally, combining our results gives us the following rather surprising result:

Corollary 3 *With probability 1, any two random graphs on $\mathbb{N}$ are isomorphic, and furthermore isomorphic to the Rado graph. In other words, up to labeling, any random graph on $\mathbb{N}$ is the Rado graph.*

(... wait, what?)