Probabilistic Methods in Graph Theory	Instructor: Paddy
Lecture 3: Infinite Graphs	
Week 1 of 1	*Mathcamp 2010*

Consider the following property:

Definition. Let (‡) denote the following property of graphs: 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.

Is there a graph on infinitely many vertices that satisfies this property?

So: one natural question we could ask about our property (‡), then, is the following: if we take a random graph on $\aleph_0$-many vertices, what is the likelihood that we would get a graph that satisfies (‡)?

To answer this question, let's first define just what we mean by a random graph on $\aleph_0$:

Definition. $G_{\aleph_0,p}$ is the random graph on the vertex set $\mathbb{N}$ formed by doing the following: given a biased coin that comes up heads with probability p and tails with probability $1-p$, flip this coin for every pair of distinct natural numbers $\{x, y\}$. If it comes up heads, add this edge to our graph; else, do not add this edge.

Given this definition, we have the following rather remarkable result:

Theorem 1 *If G is a random graph of the form $G_{\aleph_0,p}$, for $p \neq 0, 1$, then G satisfies (‡) with probability 1.*

Proof. Choose any pair of finite disjoint subsets U, W in $V(G)$. Then, for any vertex $v \in V(G), v \notin U \cup W$, let A_v be the event that v is connected to all of U and none of W. Then, we have that

$$Pr(A_v) = p^{|U|} \cdot (1-p)^{|V|} > 0.$$

Because the probability that A_v doesn't happen plus the probability that A_v **does** happen must sum to 1, we then know that

$$Pr(\text{not } A_v) = 1 - p^{|U|} \cdot (1-p)^{|V|} = \lambda < 1,$$

for some constant $\lambda \in (0, 1)$.

Thus, we know that the probability of k different vertices $v_1, \ldots v_k$ all failing to satisfy A_v is λ^k, which goes to 0 as k increases! So we can specifically bound this probability above by ϵ, for any $\epsilon > 0$, by simply looking at enough vertices.

Now, note that there are only countably many pairs of finite disjoint subsets of $\mathbb{N}$; consequently, we can enumerate all such pairs in a list $\{(U_i, W_i)\}_{i=1}^{\infty}$, and bound the probability of (U_i, W_i) failing to 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 of these events failing is bounded above by the sum

$$\sum_{n=1}^{\infty} \frac{\epsilon}{2^n} = \epsilon;$$

so thie probability of our graph satisfying property (‡) is greater than $1-\epsilon$, for any $\epsilon > 0$; i.e. the probability of our graph satisfying this property is 1! So, almost every random graph satisfies property (‡).

So: the probabilistic method is a fantastically useful way to show the existence of graphs with certain properties! However, it's not so great for actually providing concrete examples of such graphs; typically, an application of probabilistic ideas will only tell you that most graphs have your property, not what one such graph might actually look like.

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:

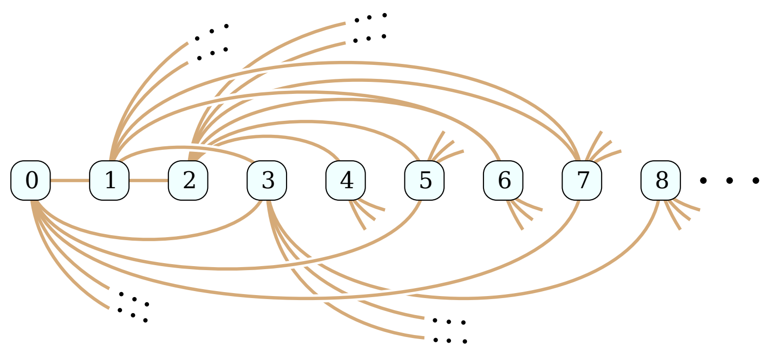

- Start by defining $R_0 = K_1$, the graph with a single vertex.
- If R_k is defined, let R_{k+1} be defined by the following: take R_k, and add a new vertex v_U for every possible subset U of R_k's vertices. Now, add an edge from v_U to every element in U, and to no other vertices in R_k.
- Let $R =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

We claim that this 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 value n such that U, V are both in fact 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 (‡).*

[1]An isomorphism of two graphs $G = (V, E), G' = (V', E')$ is a bijection $\phi : V \to V'$ such that $\{u, v\}$ is an edge in V iff $\{\phi(u), \phi(v)\}$ is an edge in V'.

Proof. To see this, take any two graphs $G = (V, E), G' = (V', E')$ on $\aleph_0$-many vertices that satisfy (‡); we will exhibit an isomorphism ϕ"$V \to V'$ between them.

To do this: fix some ordering $\{v_i\}_{i=1}^{\infty}$ of V's vertices, and do the same for V'. We start with ϕ 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, and
  - 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 $v' \in V'$ such that v' is adjacent to all of the vertices in $\phi(U)$ and no other yet-defined vertices in V' that ϕ hits yet (as both sets are stil finite.)
- At even steps: do the exact same thing as above, except switch V and V'.

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; so it suffices to prove that it preserves edges.

To see why this is true: take any edge $\{u, v\}$ in V, and assume (WLOG) that ϕ was defined on u before it defined on v. Then, when we defined $\phi(u)$, we did it in only one of two ways:

- We defined $\phi(u)$ at an odd stage. In this case, when we defined $\phi(u)$, we defined $\phi(u)$ so that it would be adjacent to all of u's neighbors that we've already defined ϕ on – i.e. v! So we know that $\{\phi(u), \phi(v)\}$ is an edge.
- We defined $\phi(u)$ at an even stage. In this case, we again picked u so that, amongst the already-mapped-to elements of V, it would be adjacent to **only** those elements $w \in V$ so that $\{\phi(u), \phi(v)\}$ are adjacent! So, because $\{u, v\}$ is an edge, so is $\{\phi(u), \phi(v)\}$.

As ϕ is a bijection, the above logic easily goes the other way: so $\{u, v\}$ is an edge in E iff $\{\phi(u), \phi(v)\}$ is an edge in E'. Consequently, 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 are isomorphic.*

(... wait, what?)