CCS Discrete II	Professor: Padraic Bartlett

Lecture 9: Infinite Random Graphs

Week 9	*UCSB 2015*

Before we start with the graph theory, we need to take a quick detour into the language of probability: specifically, we need the concept of **independence.**

1 Probability Independent Events

Definition. For any two events A, B that occur with nonzero probability, define $Pr(A$ given $B)$, denoted $Pr(A|B)$, as the likelihood that A happens given that B happens as well. Mathematically, we define this as follows:

$$Pr(A|B) = \frac{Pr(A \cap B)}{Pr(B)}.$$

In other words, we are taking as our probability space all of the events for which B happens, and measuring how many of them also have A happen.

Definition. Take any two events A, B that occur with nonzero probability. We say that A and B are **independent** if knowledge about A is useless in determining knowledge about B. Mathematically, we can express this as follows:

$$Pr(A) = Pr(A|B).$$

Notice that this is equivalent to asking that

$$Pr(A) \cdot Pr(B) = Pr(A \cap B).$$

Definition. Take any n events $A_1, A_2, \ldots A_n$ that each occur with nonzero probability. We say that these n events are are **mutually independent** if knowledge about any of these A_i events is useless in determining knowledge about any other A_j. Mathematically, we can express this as follows: for any $i_1, \ldots i_k$ and $j \neq i_1, \ldots i_k$, we have

$$Pr(A_j) = Pr(A_j | A_{i_1} \cap \ldots \cap A_{i_k}).$$

It is not hard to prove the following results:

Theorem. A collection of n events $A_1, A_2, \ldots A_n$ are mutually independent if and only if for any distinct $i_1, \ldots i_k \subset \{1, \ldots n\}$, we have

$$Pr(A_{i_1} \cap \ldots \cap A_{i_k}) = \prod_{j=1}^{k} A_{i_j}.$$

Theorem. Given any event A in a probability space Ω, let $A^c = \{\omega \in \Omega \mid \omega \notin A\}$ denote the **complement** of A.

A collection of n events $A_1, A_2, \ldots A_n$ are mutually independent if and only if their complements $\{A_1^c, \ldots A_n^c\}$ are mutually independent.

We reserve the proofs of these results for the homework! We close by noting one particularly surprising **non**-result:

Non-theorem. Pairwise independence does not imply independence! In other words, it is possible for a collection of events $A_1, \ldots A_n$ to all be **pairwise independent** (i.e. $Pr(A_i \cap A_j) = Pr(A_i)Pr(A_j)$ for any i, j) but not mutually independent!

Example. There are many, many examples. One of the simplest is the following: consider the probability space generated by rolling two fair six-sided dice, where any pair (i, j) of faces comes up with probability $1/6$.

Consider the following three events:

- A, the event that the first die comes up even.
- B, the event that the second die comes up even.
- C, the event that the sum of the two dice is odd.

Each of these events clearly has probability $1/2$. Moreover, the probability of $A \cap B, A \cap C$ and $B \cap C$ are all clearly $1/4$; in the first case we are asking that both dice come up even, in the second we are asking for (even, odd) and in the third asking for (odd, even), all of which happen $1/4$ of the time. So these events are pairwise independent, as the probability that any two happen is just the products of their individual probabilities.

However, $A \cap B \cap C$ is impossible, as $A \cap B$ holds iff the sum of our two dice is even! So $Pr(A \cap B \cap C) = 0 \neq Pr(A)Pr(B)Pr(C) = 1/8$, and therefore we are not mutually independent.

We use this to study the following property of random graphs:

2 The Rado Property

Definition. Let (‡) denote the **Rado** 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.

Notice that no **finite** graph has this property, because we could just make U equal to all of G, and we would be unable to find any $v \notin U$. But what if we considered an **infinite** random graph?

Definition. Take $\mathbb{N}$ vertices $\{v_i\}_{i=1}^{\infty}$. For each pair of vertices v_i, v_j, flip a coin that comes up heads half of the time and tails the other half of the time, and connect v_i to v_j with an edge if and only if our coin came up heads.

We call the result of this process a **random graph on $\mathbb{N}$ vertices with respect to the uniform distribution,** and denote this as $G_{\mathbb{N},1/2}$ for shorthand.

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

Theorem 1. *If G is a random graph of the form $G_{\mathbb{N},p}$, 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

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

Notice that for any collection $\{v_1, \ldots v_k\}$ of distinct vertices, all of these events $A_{v_1}, \ldots A_{v_n}$ are independent: this is because the coinflips used to determine any A_{v_i} do not matter for any other A_{v_j}, and therefore that $Pr(A_{v_i}) = Pr(A_{v_i} | A_{v_{j_1}} \cap \ldots \cap A_{v_{j_l}})$ for any $i \neq j_1, \ldots j_l$.

Consequently, by our theorem earlier, the complements of these events $A^c_{v_1}, \ldots A^c_{v_k}$ are **also** independent!

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 for any v,

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

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

Thus, by independence, we know that the probability of k different vertices $v_1, \ldots v_k$ all with A_{v_i} failing to hold 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 $\mathbb{N}$ (prove this if you don't see why!) 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:

Definition. The **Rado** graph is the following graph:

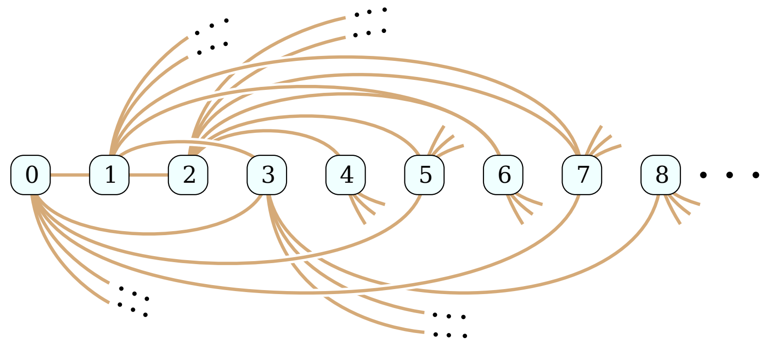

(Shamelessly stolen from Wikipedia.)

- 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 $\mathbb{N}$-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.

We close our study of (‡) with the following proposition:

Proposition 2. *Any two graphs that satisfy* (‡) *are isomorphic*[1].

Proof. To see this, take any two graphs $G = (V, E), G' = (V', E')$ on $\mathbb{N}$-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.)

[1]We say that two graphs G_1, G_2 are **isomorphic** if there is a bijection $\varphi : V(G_1) \to V(G_2)$, such that $\{x, y\}$ is an edge in G_1 if and only if $\{\varphi(x), \varphi(y)\}$ is an edge in G_2. In other words, there is a way to "relabel" the vertices of G_1 so that it looks like G_2: i.e. G_1 and G_2 are the "same," up to isomorphism.

- 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 ϕ that's defined on all of V, V'; we claim that such a map is an isomorphism.

ϕ is 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!

□

Combining our results gives us the following rather surprising result:

Corollary 3. *With probability 1, any two random graphs are isomorphic.*

(... wait, what?)