LTCC Course: Graph Theory January-February 2013
§4 Probabilistic Methods and Random Graphs

Peter Allen*

October 21, 2013

1 References

Chapter 11 of Diestel is a good source for most of what we will cover in this course. Rather more material is to be found in Chapter VII of Bollobás.

For those interested in exploring this topic further, there are some excellent textbooks devoted entirely to this subject.

The classic text on random graphs is:

- B. Bollobás, *Random Graphs*, Cambridge Univerity Press, 2nd Edition (2001). See `http://www.cup.cam.ac.uk/us/catalogue/catalogue.asp?isbn=9780521797221`

The first edition, from 1985, is also a very good source, but the 2nd edition is naturally more up-to-date.

Another excellent book with the same title is:

- S. Janson, T. Łuczak and A. Ruciński, *Random Graphs*, Wiley (2000). See `http://eu.wiley.com/WileyCDA/WileyTitle/productCd-0471175412, descCd-authorInfo.html`

Finally, another classic text with a rather broader scope is:

- N. Alon and J. Spencer, *The Probabilistic Method*, Wiley, 3rd Edition (2008). See `http://eu.wiley.com/WileyCDA/WileyTitle/productCd-0470170204.html`

Earlier editions are also excellent.

2 Random Methods

Is there, for all large values of n, a graph on n vertices with no clique of size $5\log n$, and no independent set of size $5\log n$?

Is there a graph with some number n of vertices, no cycles of length less than 100, and no independent sets of size greater than $n/100$? Note that such a graph has chromatic number at least 100.

Is there, for all large values of n, a graph on n vertices with maximum degree at most $n^{2/3}$, and, for every pair (U,V) of sets of vertices with $|U|,|V| \geq n^{1/2}$, there is an edge from U to V?

The answer to each of these questions is yes, but how might one go about proving that?

*These notes are based on Graham Brightwell's notes for LTCC Graph Theory Courses in 2009-12.

The natural first reaction on being confronted by questions like these is to try and *construct* graphs with the required combination of properties. In each case, one of the two required properties demands that the graph have rather few edges, and the other demands rather many. It's hard to strike a balance, and any attempt to base a construction around some nice structure seems doomed to failure[1].

For each of the problems above, by far the best way to solve the problem is not to give an explicit construction at all, but instead to "construct" the graph "at random", and show that, with positive probability, the random graph constructed has the required combination of properties.

As a first, striking, example, let's consider the first question. Here, what we do is to build our random graph in the most naive possible way. We take some (large) set of vertices, say $V = [n] = \{1, \dots, n\}$. Then, for each pair of vertices, we toss a fair coin, and put an edge between the pair if we get a head.

Let's calculate the probability that a particular set C of c vertices forms a clique. This means that each of the $\binom{c}{2}$ coin-tosses corresponding to the pairs in the clique was a head, an event with probability $2^{-\binom{c}{2}}$.

Now, the *expected* number of cliques of size c is given by:

$$\binom{n}{c} 2^{-\binom{c}{2}} \leq \frac{n^c}{c!} 2^{-\binom{c}{2}} \leq \left(\frac{ne}{c2^{(c-1)/2}} \right)^c.$$

If $c \geq 2\log_2 n$, then $n \leq 2^{c/2}$, and so the expected number of c-cliques is at most $(e\sqrt{2}/c)^c$. For $c \geq 5$, this expectation is at most $1/3$. This implies that the probability that the random graph has a clique of size c is at most $1/3$. The same calculation shows that the probability that the graph has an independent set of size c is also at most $1/3$.

So, with probability at least $1/3$, the random graph on $n \geq 5$ vertices has neither a clique nor an independent set of size as large as $2\log_2 n$.

The idea: We can calculate expectation using 'linearity of expectation' and estimate probabilities from expectations.

This delightfully simple argument was first given by Paul Erdős in 1947, and it is one of the first instances of the successful use of the "probabilistic" method in combinatorics.

One can, and indeed Erdős did, recast this entire proof as a "counting argument". Of the $2^{\binom{n}{2}}$ graphs with vertex set $[n]$, the number of them in the class A_C with a clique or independent set on the set C of size c is ..., and then effectively the same calculation shows that there are some graphs in none of the sets A_C.

The result we've proved is usually stated in terms of *Ramsey numbers*, a topic we'll return to next week. It says that, for all $c \in \mathbb{N}$, there is a graph on $n = \frac{c}{e} 2^{(c-1)/2}(1 + o(1))$ vertices containing neither a clique nor an independent set of size c. (Here the $o(1)$ refers to a term that tends to zero as $c \to \infty$.) This is exactly what it means to say that the *Ramsey number* $R(c, c)$ is at least $\frac{c}{e} 2^{(c-1)/2}(1 + o(1))$.

3 Sperner's Theorem

This result isn't about graphs at all. (Apologies.)

[1] In fact there are clever constructions of 'pseudorandom' graphs which solve the last two questions, but these usually use deep theorems from algebraic geometry or number theory in their proofs - and they are still not good enough for the first question!

Consider the Boolean cube Q^n. This is the collection of all 2^n subsets of the set $[n]$, ordered by inclusion. An *antichain* or *Sperner family* in Q^n is a family of subsets of $[n]$, no one of which is contained in another.

One way to construct a large antichain is to take the family of all subsets of $[n]$ of size exactly $\lfloor n/2 \rfloor$. Is this the largest? A famous theorem of Sperner from 1927 says that it is.

A stronger result is the so-called *LYM Inequality*, which says that, if $\mathcal{A}$ is any antichain in Q^n, then

$$\sum_{A \in \mathcal{A}} \binom{n}{|A|}^{-1} \leq 1.$$

This implies Sperner's Theorem, as the largest binomial coefficient is $\binom{n}{\lfloor n/2 \rfloor}$.

Here's a snappy proof of the LYM Inequality.

Take any antichain $\mathcal{A}$ in Q^n. Now take a random ordering $(x_1, x_2, \ldots, x_n)$ of $[n]$, and consider the *chain* $\mathcal{C}$ of sets $C_0, C_1, \ldots, C_n$, where $C_j = \{x_1, \ldots, x_j\}$. Notice that, for each pair of elements of $\mathcal{C}$, one is contained in the other; thus $\mathcal{C}$ intersects $\mathcal{A}$ in at most one element. Hence it is certainly the case that $\mathbb{E}|\mathcal{C} \cap \mathcal{A}| \leq 1$.

Now, for each $A \in \mathcal{A}$, the probability that A is in the random chain $\mathcal{C}$ is exactly $1/\binom{n}{|A|}$, since $\mathcal{C}$ includes exactly one set with $|A|$ elements, and it's equally likely to be any of them.

So we have

$$1 \geq \mathbb{E}|\mathcal{A} \cap \mathcal{C}| = \sum_{A \in \mathcal{A}} \mathbb{P}(A \in \mathcal{C}) = \sum_{A \in \mathcal{A}} \frac{1}{\binom{n}{|A|}},$$

and we are done.

The idea: We find a clever random variable whose expectation is exactly the sum we want to bound, and whose expectation is 'obviously' at most one.

4 Some probabilistic inequalities

So far we only used very simple probabilistic tools: the linearity of expectation together with the fact that if the expectation of a (non-negative) random variable is smaller than one, then it takes the value zero with non-zero probability.

There are three important inequalities we will need for more difficult results. All deal with a non-negative random variable X.

Markov's inequality: For each $a > 0$, we have $\mathbb{P}(X \geq a) \leq \frac{\mathbb{E}X}{a}$.

Chebyshev's inequality: For each $k > 1$, we have

$$\mathbb{P}\Big(|X - \mathbb{E}X|^2 \geq k^2\big(\mathbb{E}X^2 - (\mathbb{E}X)^2\big)\Big) \leq \tfrac{1}{k^2}\,.$$

Chernoff's inequality: Suppose $X_1, \ldots, X_n$ are independent Bernoulli random variables, with $\mathbb{P}(X_i = 1) = p_i$ (we do not need that the p_i are identical: but if they are we say X is a Binomial random variable). Let $\mu = \sum_i p_i$ and $X = \sum_i X_i$, so that by linearity of expectation we have $\mu = \mathbb{E}X$.

Then for each $0 < \delta < 1$, we have

$$\mathbb{P}\big(|X - \mu| \geq \delta\mu\big) < 2e^{-\frac{\delta^2\mu}{3}}\,.$$

All three inequalities are fairly easy to prove. The first two just follow from the definition of expectation, while the last follows from computing the expectation of e^{tX} for a suitable t and applying Markov's inequality.

In all cases, the idea is supposed to be that a random variable is 'usually' 'not too far' from its mean. Markov's inequality provides only weak bounds on what 'usually' and 'not too far' should be, but in return we have only to compute expectation. Chebyshev's inequality often tells us much more, but we have to calculate not only the expectation of X but also of X^2, which can be tricky. Finally Chernoff's inequality gives very strong bounds and is easy to use, but only applies in the special case of a sum of independent Bernoulli random variables.

5 Girth and Chromatic Number

This section concerns another classic application of the probabilistic method. It's covered very well in Diestel, for instance, and I refer the reader there for full details.

The *girth* of a graph G is the length of the shortest cycle in G. A graph with high girth, i.e., no short cycles, is "locally" like a tree. Specifically, suppose the girth is at least g, and we take $2k < g$ and let $N_k(v)$ be the number of vertices at distance at most k from vertex v, then all the sets $N_k(v)$, for $v \in V(G)$, induce trees. In particular, all the induced subgraphs on the sets $N_k(v)$ are 2-colourable. Is it possible for the whole graph to have a large chromatic number?

How can we prove that the chromatic number of a graph is large? One way is to show the graph has no large independent sets, as the set of vertices receiving any given colour is an independent set.

We shall prove the following result.

Theorem 5.1. *For each $k \in \mathbb{N}$ and all sufficiently large n, there exists a graph on n vertices with girth greater than k and no independent set of size at least $\lceil n/k \rceil$.*

The proof of this result (also due to Erdős, this time from 1959) is slightly more complicated than the one about Ramsey numbers. For a start, our random process involves the tossing of *biased* coins. Again we fix a vertex set $[n]$: now we put an edge between each pair of vertices with probability p, all choices made independently. Here $p = p(n)$ is a function of n that we can choose to suit our needs. What we have defined here is the standard model $\mathcal{G}(n,p)$ of *random graphs*. We use $G(n,p)$ to denote a random graph chosen according to this method.

We would like to find some p such that (with positive probability, for sufficiently large n) the random graph $G(n,p)$ has neither an independent set of size at least n/k, nor a cycle of length at most k. Unfortunately there is no such p: if $p = 100/n$, it is very likely that $G(n,p)$ contains triangles, and it is very likely that n/k vertices of $G(n,p)$ are isolated (and in particular form an independent set) for some (large) k. If $p < 100/n$, then the situation with isolated vertices only gets worse, and if $p > 100/n$ then there will still probably exist triangles. So the graph we are looking for is not a random graph $G(n,p)$. But the random graph 'almost' works, and we can modify it to make it work.

Proof. Given a natural number k, take $p(n) = n^{-1+\varepsilon}$, where $0 < \varepsilon < 1/k$.

Let $S \subset [n]$ be a fixed set of size $\lceil n/k \rceil$. How close is S to being independent in $G(n,p)$? The number X of edges in S has expectation $p\binom{|S|}{2}$, and X is a random variable to which the Chernoff bound applies. In particular, we have

$$\mathbb{P}\big(|X| < \tfrac{pn^2}{8k^2}\big) < 2e^{-\frac{pn^2}{24k^2}} < e^{-n^{1+\varepsilon/2}} < \tfrac{1}{4}2^{-n}\,.$$

where the last two inequalities hold for all sufficiently large n. Intuitively, this says that the chance that S is even 'near' to being independent is very tiny. Now there are only at most 2^n different choices for $S \subset [n]$. The probability that any one of them contains less than $\frac{pn^2}{8k^2}$ edges in $G(n,p)$ is therefore at most

$$2^n \cdot \tfrac{1}{4} 2^{-n} = \tfrac{1}{4}\,.$$

Finally we conclude that for all sufficiently large n, with probability at least $\frac{3}{4}$, the random graph $G(n,p)$ has the following property. Every set $S \subset [n]$ of size at least $\lceil n/k \rceil$ contains at least $\frac{pn^2}{8k^2} > n$ edges (where the inequality holds for sufficiently large n).

Now we consider the expectation of $N = N(G)$, the number of cycles in G of length at most k.

$$\mathbb{E}N = \sum_{g=3}^{k} \frac{n(n-1)\cdots(n-g+1)}{2g} p^g \le \frac{1}{2}(k-2)(np)^k.$$

Notice that $(np)^k = n^{\varepsilon k}$, so $\mathbb{E}N = o(n)$. Now, N is a non-negative random variable, and so by Markov's inequality, with $a = n/2$, we find that $\mathbb{P}(N \ge n/2) < 1/2$ for sufficiently large n.

This means that, with probability at least $1/4$, the random graph $G(n,p)$ has the following two properties: $N(G) < n/2$, and every set $S \subset [n]$ of size at least $\lceil n/k \rceil$ contains at least n edges.

This means that there is some graph H with both these properties.

Take such an H. Mark in each cycle of length at most k one edge (these edges need not be distinct) and delete the marked edges. Let the result be H'. Now we have deleted at most $N(H) < n/2$ edges, so any set $S \subset [n]$ of size at least $\lceil n/k \rceil$ still contains at least $n/2$ edges in H'. In other words, H' has no independent sets of size at least $\lceil n/k \rceil$. And it also has no cycles of length at most k, so it is the desired graph. □

The idea: For the p we chose, $G(n,p)$ is 'very far' from having independent sets of size n/k, and we can use Chernoff's inequality which gives us very strong probability bounds to check this. On the other hand, even though $G(n,p)$ contains a few short cycles, it doesn't have many. We can use Markov's inequality to say that with a reasonable probability, $G(n,p)$ has 'not too many more than expectation' short cycles, and we can delete them.

6 Designs

Consider the following problem. Given a complete graph on n vertices, partition the edges into as many edge-disjoint triangles as possible (plus a left-over set).

Is it possible for the leftover set to be empty? A design theorist will recognise this as asking for a *Steiner triple system* on n elements. These are known to exist if and only if n is congruent to 1 or 3 modulo 6 (to see that this is necessary, observe that all vertices have to have even degree, so n must be odd, and the number of edges has to be divisible by 3, which rules out 5 modulo 6).

But what happens if we ask for edge-disjoint copies of K_k instead of K_3? What happens if instead of partitioning the edges of a complete graph K_n we partition the edges of a complete t-uniform hypergraph $K_n^{(t)}$ into copies of $K_k^{(t)}$? These are all still good questions in design theory, but they are very hard. In particular, there is not one single example known of a partition of a 6-uniform hypergraph on n vertices into edge-disjoint copies of $K_7^{(6)}$ for any $n \ge 8$ ($n = 7$ is trivial!). Neither is there a proof that these partitions cannot exist. This problem dates back to Steiner in 1853!

For that matter, what happens if in the original problem n is not congruent to 1 or 3 modulo 6? We cannot hope for a perfect partition, but we can make the leftover set small.

Theorem 6.1 (Rödl, 1985). *For any $1 \le t \le k \in \mathbb{N}$ and $\gamma > 0$, if n is large enough, there is a partition of the edges of K_n^t into edge-disjoint copies of $K_k^{(t)}$ and a leftover set of size at most γn^t.*

We will just consider the case $t = 2, k = 3, n > 3$ here - i.e. the original problem. Note that the $t = 1$ case is trivial, as is the $t = k$ case, as is the $k = n$ case - so this is the smallest non-trivial case, but actually the general case is not much harder. This means we want to find $(1 - 2\gamma)n^2/6$ edge-disjoint triangles in K_n.

At first it is not clear how we can use the probabilistic method here. We have to somehow construct the partition 'randomly', and the obvious way is to select triangles from K_n 'randomly' in some way. Suppose we selected triangles from K_n independently at random with some probability p. Then the expected number of triangles we select is $p\binom{n}{3}$, so we should choose p to be at least $(1 - \gamma)/n$ in order for the expected number of triangles to be at least $(1 - \gamma)n^2/6$. But it is easy to believe that these triangles will not be edge-disjoint: given any one triangle xyz, the expected number of (other) selected triangles containing xy is $p(n - 3)$, so the expected number of selected triangles which share an edge with xyz is $3p(n - 3)$, which is close to 3.

Maybe we can use a 'trick' like the one we used in proving Theorem 5.1, and delete some triangles in order to get a collection of edge-disjoint triangles? The only simple rule which will give us a collection of edge-disjoint triangles is: delete every triangle which shares an edge with another triangle. We can estimate how many selected triangles we have to delete by the number Y of pairs of selected triangles which share an edge. Now we can estimate $\mathbb{E}Y \le p^2n^4$, and by Markov's inequality conclude that (with probability at least $9/10$) we only have to delete $10p^2n^4$ triangles. Unfortunately this is more triangles than we expect to exist if $p \ge (1 - \gamma)/n$!

Via this select-and-delete method, where we select independently with probability p, we can find a collection of at least

$$p\binom{n}{3} - 10p^2n^4$$

edge-disjoint triangles. If we take $p = \varepsilon/n$ for some $\varepsilon > 0$, this is approximately $(\varepsilon - 60\varepsilon^2)\frac{n^2}{6}$ edge-disjoint triangles. But this function is a quadratic in ε whose maximum is (well) below $n^2/6$, so we cannot hope to even get close to the desired number of edge-disjoint triangles. But, if ε is very small then we will delete only a tiny proportion of the triangles we selected.

Rödl's insight is the following. If we perform select-and-delete, with $p = \varepsilon/n$ for some very small $\varepsilon > 0$, then we will get a (small) collection of edge-disjoint triangles T_1, a (tiny) collection of 'wasted' edges contained in the deleted triangles, and a (large) set of 'leftover' edges L_1. But this 'leftover' set will be very well-behaved - it will have the property that 'most' edges are contained in about the same number of triangles in L_1. We can repeat select-and-delete using only the edges of the leftover set to get another collection of edge-disjoint triangles T_2 and another leftover set L_2. By construction $T_1 \cup T_2$ is a collection of edge-disjoint triangles, and we can show that L_2 is still well-behaved. We can keep taking these 'Rödl Nibbles' until we get down to a tiny leftover set and the desired collection of edge-disjoint triangles.

We will only prove the following very simplified version.

Lemma 6.2. *Given $\gamma > 0$ and $\varepsilon > 0$, for all sufficiently small $\delta > 0$, if n is large enough, the following holds. Suppose $D \ge \gamma^2 n/32$. Let G be a graph on n vertices with at least $\gamma n^2/8$ edges in which every edge is in D triangles. Let S be a set of triangles selected independently with probability $p = \varepsilon/D$ from the triangles of G. Let T be the set of triangles in S which share no edge with any other triangle of S. Let G' be obtained from G by deleting any edge contained in a triangle of S. Then with probability at least $1/2$ we have*

- $|S| = (1 \pm \delta)\varepsilon e(G)/3$.
- $|S \setminus T| < 30\varepsilon|S|$.

- *All but at most $\delta e(G)$ edges of G' lie in $(1 \pm \delta)(1-\varepsilon)^2 D$ triangles of G'.*

The 'real' version of this lemma has the weaker condition on G that most of its edges are in about D triangles—so we can apply the 'real' lemma to G':

Lemma 6.3. *Given $\gamma > 0$ and $\varepsilon > 0$, for all sufficiently small $\delta' > 0$ there exists $\delta > 0$ such that for all sufficiently large n the following holds. Suppose $D \geq \gamma^2 n/32$. Let G be a graph on n vertices with at least $\gamma n^2/8$ edges in which all but at most $\delta e(G)$ edges are in $(1 \pm \delta)D$ triangles. Let S be a set of triangles selected independently with probability $p = \varepsilon/D$ from the triangles of G. Let T be the set of triangles in S which share no edge with any other triangle of S. Let G' be obtained from G by deleting any edge contained in a triangle of S. Then with probability at least $1/2$ we have*

- $|S| = (1 \pm \delta')\varepsilon e(G)/3$.
- $|S \setminus T| < 60\varepsilon|S|$.
- *All but at most $\delta' e(G)$ edges of G' lie in $(1 \pm \delta')(1-\varepsilon)^2 D$ triangles of G'.*

This lemma is not conceptually harder to prove. But the weaker condition makes the calculations substantially longer.

Now we can make two remarks. First, the bounds on $|S|$ and $|S \setminus T|$ yield bounds on $e(G')$: we find that $e(G')$ is very close to $(1-\varepsilon)e(G)$ (the number of edges deleted is at least $3|T|$ and at most $3|S|$). We conclude that after t steps—t nibbles—applying Lemma 6.3 starting from $G_0 = K_n$ we get to a graph G_t in which the number of edges is close to $(1-\varepsilon)^t n^2/2$, and most edges are in about $(1-\varepsilon)^{2t} n$ triangles. If we choose τ such that $\gamma/2 < (1-\varepsilon)^\tau < \gamma$, then we conclude that $e(G_\tau) < \gamma n^2/2$ and that for each $t \leq \tau$, most edges of G_t are in at least $(1-\varepsilon)^{2t} n/2 > \gamma^2 n/8$ triangles. In particular this means the conditions of Lemma 6.3 that $e(G)$ and D should not be too small are satisfied up to τ steps, so we are allowed to make τ nibbles.

Second, at each nibble step we partition the edges of G into three parts: the edges of T (which are edge-disjoint triangles), the edges in S which are not in T (which are 'wasted'), and the edges of G'. Because $|S \setminus T|$ is much smaller than $|T|$, we conclude that the number of 'wasted' edges in each nibble is a tiny fraction of the number of edges covered by T. So after τ nibbles, we have partitioned $E(K_n)$ into a collection of edge-disjoint triangles, a wasted set which is tiny by comparison, and $E(G_\tau)$ which we know has size at most $\gamma n^2/2$. Provided we chose ε sensibly, this means the wasted set and $E(G')$ together account for only at most γn^2 edges, and we conclude that the collection of edge-disjoint triangles we found covers all but γn^2 edges of $E(K_n)$ as desired.

Proof of Lemma 6.2. Given $\gamma > 0$ and $\varepsilon > 0$, we choose any sufficiently small $\delta > 0$.

Let G be an n-vertex graph in which every edge is in D triangles. Let S be obtained by choosing independently triangles from G with probability $p = \varepsilon/D$. Let T be the set of triangles in S which share no edge with any other triangle of S, and let G' be the graph obtained from G by removing all edges in triangles of S.

First we can estimate $|S|$. Observe that G has $De(G)/3$ triangles, and S is obtained by selecting independently from these triangles with probability p. So the expectation of $|S|$ is $pDe(G)/3 = \varepsilon e(G)/3$. By the Chernoff bound, we conclude that

$$\Pr\Big(\big||S| - \varepsilon e(G)/3\big| > \delta\varepsilon e(G)/3\Big) < 2e^{-\delta^2 \varepsilon e(G)/3}$$

which (provided n is large enough) is smaller than 0.1. We conclude that (with probability at least 0.9) the set S has size $|S| = (1 \pm \delta)\varepsilon e(G)/3$ as desired.

Now we can estimate $|S \setminus T|$ by exactly the same method we tried to use above. We estimate how many pairs of triangles we find sharing an edge in S; for each pair we delete two triangles. The number of pairs of triangles sharing an edge in G is just $\binom{D}{2}e(G)$, where we choose the shared edge e then the pair of triangles. The expected number which appear in S is then $p^2\binom{D}{2}e(G) < \varepsilon^2 e(G)/2$. So the expected size of $S \setminus T$ is at most $\varepsilon^2 e(G)$ (since we delete both triangles from each pair). So by Markov's inequality, with probability at least 0.9, the size of $S \setminus T$ is at most $10\varepsilon^2 e(G) < 60\varepsilon|S|$ as desired.

Finally we have to show that most edges of G' are in the 'right' number of triangles of G'. This is the most difficult part of the proof. First, consider some edge uv of G. We condition on the event that uv is also an edge of G': that is, we assume no triangle containing uv is in S. We want to know T_{uv}, the number of triangles containing uv in G'. First we will try to find $\mathbb{E}T_{uv}$: for this we just need to find the probability that a given triangle uvw of G survives to G'. Now there are $D-1$ triangles of G (apart from uvw) using the edge uw, and another (different!) $D-1$ triangles using the edge vw. The triangle uvw survives if and only if none of these triangles are in S. The probability of that occurring is just $(1-p)^{2D-2}$. So we have

$$\mathbb{E}T_{uv} = D(1-p)^{2D-2} \approx (1-\varepsilon)^2 D\,.$$

Now we need to show that T_{uv} takes a value close to its expectation (at least most of the time). It would be very nice if we could say that T_{uv} was a sum of D independent Bernoulli random variables (one for each triangle of G containing uv) and use Chernoff's inequality, but this isn't true: these random variables are not independent, as it can happen that one triangle being selected affects two of the random variables. So we have to use Chebyshev's inequality, and this means we need to estimate $\mathbb{E}T_{uv}^2$.

Claim 6.4.

$$\mathbb{E}T_{uv}^2 \le \big(\mathbb{E}T_{uv}\big)^2 + \delta^5 D^2/2\,.$$

We will prove this later. Assuming the claim, we can complete the proof. By Chebyshev's inequality, we have

$$\mathbb{P}\Big(\big|T_{uv} - \mathbb{E}T_{uv}\big|^2 \ge 2\delta^{-1}\big(\mathbb{E}T_{uv}^2 - (\mathbb{E}T_{uv})^2\big)\Big) \le \delta/16\,.$$

Substituting $\mathbb{E}T_{uv}^2$, we get that the probability of $T_{uv} \neq \mathbb{E}T_{uv} \pm \delta^2 D$ is at most $\delta/4$. Since $\delta < 1/4$, and since $D < 2\mathbb{E}T_{uv}$, we conclude that the probability of $T_{uv} \neq (1 \pm \delta/2)\mathbb{E}T_{uv}$ is at most $\delta/4$. Finally, since $\mathbb{E}T_{uv}$ is approximately $(1-\varepsilon)^2 D$, we conclude that the probability of $T_{uv} \neq (1\pm\delta)(1-\varepsilon^2)D$ is at most $\delta/4$.

Now let Y be the random variable counting edges uv of G' for which $T_{uv} \neq (1\pm\delta)(1-\varepsilon^2)D$. The expectation of Y is at most $\delta e(G)/4$, so by Markov's inequality with probability at least $3/4$ we have $Y \le \delta e(G)$. In other words, with probability at least $3/4$, all but at most $\delta e(G)$ edges of G' lie in $(1\pm\delta)(1-\varepsilon)^2 D$ triangles of G'.

Finally the probability that any one of the three claimed items fails is at most $0.1+0.1+0.25 < 0.5$, so with probability at least 0.5 all three items hold as desired. It remains to check the Claim.

Proof of Claim 6.4. Given two events A and B, we define the *covariance* of A and B to be $\mathrm{Cov}(A,B) = \mathbb{E}AB - \mathbb{E}A\mathbb{E}B$. Now if A and B are independent, their covariance is zero. This will help us evaluate $\mathbb{E}T_{uv}^2$.

For each w such that uvw is a triangle of G, let I_w be the random variable which takes value 1 if uvw is a triangle of G', and 0 otherwise. We have the following equality.

$$\mathbb{E}T_{uv}^2 - (\mathbb{E}T_{uv})^2 = \mathbb{E}\Big(\sum_w I_w\Big)^2 - \Big(\mathbb{E}\sum_w I_w\Big)^2 = \mathbb{E}\sum_w I_w^2 + 2\mathbb{E}\sum_{w<w'} I_w I_{w'} - \sum_w (\mathbb{E}I_w)^2 - 2\sum_{w<w'}\mathbb{E}I_w\mathbb{E}I_{w'}\,.$$

and this is bounded above by

$$\mathbb{E}I_w^2 + 2\sum_{w<w'} \mathrm{Cov}(I_w, I_{w'}) = \mathbb{E}T_{uv} + 2\sum_{w<w'} \mathrm{Cov}(I_w, I_{w'})\,, \tag{1}$$

where we left out the third sum and rearranged. Our aim is to show that this last quantity is bounded above by $\delta^5 D^2/2$. Now consider some w and w' which make triangles of G with uv. The events I_w and $I_{w'}$ are independent unless there is a triangle of G which uses both w, w' and one of u and v. In other words, $\mathrm{Cov}(I_w, I_{w'}) = 0$ unless either uww' or vww' (or both) are triangles of G. If one of the two (but not both) is a triangle of G, then

$$\mathrm{Cov}(I_w, I_{w'}) = (1-\varepsilon/D)^{4D-5} - (1-\varepsilon/D)^{4D-4}$$

and if both are triangles, we get

$$\mathrm{Cov}(I_w, I_{w'}) = (1-\varepsilon/D)^{4D-6} - (1-\varepsilon/D)^{4D-4}\,.$$

In either case, we get $\mathrm{Cov}(I_w, I_{w'}) \leq 2\varepsilon/D$. Now for each given w, there are at most $2D$ vertices w' such that one of uww' or vww' is a triangle of G. So (1) can be bounded above by

$$\mathbb{E}T_{uv} + \sum_{w} 4D(2\varepsilon/D) = \mathbb{E}T_{uv} + 8D < \delta^5 D^2/2\,,$$

where the final inequality is true, since $D > \gamma^2 n/16$, for sufficiently large n. □

That proves the Claim, and hence the proof of Lemma 6.2 is complete. □

7 Exercises

There are plenty more exercises in the textbooks!

1. For $k \in \mathbb{N}$, a graph $G = (V, E)$ has *Property* S_k if, for every pair (A, B) of disjoint k-element subsets of V, there is a vertex x of the graph that is adjacent to every vertex of A and no vertex of B.
(a) Find a graph with property S_1.
(b) Show that, for each $k \in \mathbb{N}$, there is a graph with property S_k.

2. A k-uniform *hypergraph* is a pair $H = (V, E)$, where V is a set of *vertices*, and E is a family of k-element subsets of V. (So a 2-uniform hypergraph is just a graph.) A hypergraph $H = (V, E)$ has *Property B* if V can be partitioned into two subsets V_1 and V_2 in such a way that no edge is entirely contained within one of the two sets.
(a) Show that, if $H = (V, E)$ is a k-uniform hypergraph with $|E| < 2^{k-1}$, then H has property B.
(b) Show that, if $H = (V, E)$ is a k-uniform hypergraph such that each edge in E intersects at most d others, and $e(d + 1) \leq 2^{k-1}$, then H has property B.

3. (a) Let $p = n^{-t}$, for $0 < t < 1$, and let k be a fixed natural number. Write down an expression for the expected number of k-cliques in $G(n, p)$. Hence show that, if $t > 2/(k - 1)$, the probability that $G(n, p)$ contains a k-clique tends to zero as $n \to \infty$.

It is also true that, if $t < 2/(k-1)$, then the probability that $G(n,p)$ contains a k-clique tends to one as $n \to \infty$: to prove this, one needs to work with the variance of the number of k-cliques.

(b) Let H denote the graph on five vertices a, b, c, d, e with seven edges: a, b, c, d form a clique, and de is also an edge. For $p = n^{-7/10}$, find the expected number of copies of H in $G(n, p)$. What is

$$\lim_{n\to\infty} \mathbb{P}(G(n,p) \text{ contains a copy of } H)?$$

(c) There is a parameter $b(H)$ of graphs such that, if $p = n^{-t}$ and $t > b(H)$, then the probability that $G(n, p)$ contains a copy of H as a subgraph tends to zero, while if $p = n^{-t}$ and $t < b(H)$, then this probability tends to 1. Based on the calculations in this question, what do you think this parameter $b(H)$ might be?

4. Set $p = n^{-2/5}$, and consider a random graph $G = G(n, p)$.
(a) Show that the degree of any fixed vertex v has a Binomial distribution, and find an upper bound on the probability that this degree is greater than or equal to $n^{2/3}$.
(b) Show that the probability that the maximum degree of G is at most $n^{2/3}$ is at least $2/3$.
(c) Show that, with probability at least $2/3$, for every pair (U, V) of subsets of $V(G)$, with $|U|, |V| \geq n^{1/2}$, there is an edge from U to V.
(d) What can you deduce from (b) and (c)?

5. (a) Try to prove Lemma 6.3. You should find that the only difficult part is to prove that most edges are in the 'right' number of triangles. You will not be able to prove that T_{uv} behaves nicely for every edge $uv \in G$: you will need to assume both that uv happened to lie in about the 'right' number of triangles in G, and that 'most' of those triangles share edges with about the 'right' number of triangles in G. If you make this assumption, you should be able to modify the argument given to show that T_{uv} is likely to be about the 'right' size. Then you will need to show that there cannot be too many edges of G which don't satisfy the assumption.
(b) Try to prove the special case of Theorem 6.1 from Lemma 6.3. The difficulty here is to find out how to set constants in order to make the argument work.
(c) Try to prove Theorem 6.1—or try to understand the argument given in e.g. Alon and Spencer!