

GRAPH THEORY

Lecture notes for 1MA170

Fabian Burghart

Fall 2022


Contents

Lecture 1	Introduction: Bridges of Königsberg	2
Lecture 2	Simple graphs and subgraphs	7
Lecture 3	Trees	13
Lecture 4	Counting spanning trees	18
Lecture 5	Weights and distances	24
Lecture 6	Hamilton cycles	30
Lecture 8	The max-flow-min-cut theorem	35
Lecture 9	Matchings	40
Lecture 10	Connectivity	45
Lecture 11	Planarity	50
Lecture 12	Vertex colourings	55
Lecture 13	More on colourings	60
Lecture 14	Edge-colourings and Ramsey Theory	64
Lecture 16	Szemerédi's regularity lemma	68
Lecture 17	The Rado graph	73
Lecture 18	The Erdős-Rényi random graph	78
Lecture 19	More on random graphs	83

Lecture 1 Introduction: Bridges of Königsberg

In the 18th century, the Prussian city of Königsberg was divided into four different landmasses by the Pregel river: The islands Kneiphof and Lomse and the mainland to either side of the river. Those landmasses were connected by seven bridges:

- Two bridges each going from either part of the mainland to Kneiphof
- One bridge each going from either part of the mainland to Lomse
- and one more bridge connecting Kneiphof and Lomse




(Today, the situation is a bit different: The city, after remaining a German exclave until World War II then fell to Russia and is now known under its Russian name, Kaliningrad. During the war, two of the bridges were destroyed (one between either part of the mainland to Kneiphof). However, a new bridge connecting Oktyabrsky – formerly Lomse – to the southern mainland was constructed).

Given this situation, the mathematician Leonard Euler became invested in determining whether one could give a tour through the city traversing all bridges exactly once. In 1736, he gave a rigorous proof that one cannot, in a paper that is now regarded as the founding moment of graph theory. His argument was based on the following abstractions: Since the exact position of a walker on a landmass is irrelevant, we can represent the landmasses by abstract points (which we call *vertices*) and the bridges by line segments (which we call *edges*) connecting those points. These are the basic ingredients for what we call a *multigraph* in modern language:

Definition 1. A (multi-)graph G is a triple (V, E, ι) consisting of a set V of vertices, a set E of edges, and a map $\iota : E \rightarrow \mathcal{P}_1(V) \cup \mathcal{P}_2(V)$ assigning to each edge e its endpoints¹

In this way, after Euler's abstraction we obtain something like this:

¹Here, we used the following notation: For a set S , the set of all subsets with i elements is denoted by $\mathcal{P}_i(S)$. Thus, an edge can have either one or two distinct endpoints.



In modern language we have $V = \{1, 2, 3, 4\}$, $E = \{a, b, \dots, g\}$ and

$$\begin{aligned} \iota(a) = \iota(b) &= \{1, 3\}, & \iota(c) = \iota(d) &= \{2, 3\} \\ \iota(e) &= \{1, 4\}, & \iota(f) &= \{2, 4\}, & \iota(g) &= \{3, 4\}, \end{aligned}$$

although more often than not specifying a graph by a picture is more convenient (note though that the graph merely captures the combinatorial information of which edges have which endpoints – a drawing necessarily contains a lot more information. More on that next lecture!).

The definition above explicitly allows $\iota(e) = \{v\}$ for some $e \in E, v \in V$ – thus, an edge might connect a vertex to itself. In this case, we say that the edge is a *loop*. Perhaps more relevant to the bridges of Königsberg, we say that two edges $e, e' \in E$ are *parallel* if $\iota(e) = \iota(e')$. The two endpoints of a common edge are called *adjacent* or *neighbours*. We also sometimes say that a vertex v is *incident* to an edge e , by which we mean that v is an endpoint of e . Finally, a graph is *finite* if $|V| + |E| < \infty$.

Having successfully encoded the topography of Königsberg into a mathematical structure, how do we proceed with the notion of a city tour?

Definition 2. Let $G = (V, E, \iota)$ be a graph. A *walk* of length k is a sequence $v_0 e_1 v_1 e_2 v_2 \dots e_k v_k$ where $e_1, \dots, e_k \in E$ and $v_0, v_1, \dots, v_k \in V$, such that $\iota(e_i) = \{v_{i-1}, v_i\}$ for all $i = 1, \dots, k$. A *trail* is a walk that uses no edge twice. A *path* is a walk that uses no vertex twice.

Moreover, a *circuit* is a trail where the first and last vertex coincide, and a *cycle* is a circuit where those vertices are the only one coinciding.

Thus every path is a trail, and every cycle is a circuit, but the converse is not true.

Definition 3. A trail that uses every edge in the graph exactly once is called an *Eulerian trail*. Analogously, a circuit using every edge exactly once is called an *Eulerian circuit* (or Euler tour). If a graph admits an Eulerian circuit, it is simply called *Eulerian*.

Hence, we are specifically interested in whether or not the graph we obtained earlier for Königsberg admits an Eulerian trail (or even an Eulerian circuit).


However, our answer will be in form of a general characterisation of such graphs, for which we need yet another two concepts:

Definition 4. Let $G = (V, E, \iota)$ be a graph. A vertex x is *connected to* a vertex y if there is a walk (equivalently, a trail or a path) starting in x and ending in y . If this is satisfied for all $x, y \in V$, then we say G is connected.

Remark 5. Every vertex is connected to itself via a “lazy” walk using no edges at all. It thus follows that connectedness is an equivalence relation on V , and the equivalence classes are called the (*connected*) *components* of G .

Definition 6. Let $G = (V, E, \iota)$ be a graph and $v \in V$ a vertex. The degree $\deg(v)$ is the number of edges v is incident to, with loops counted twice.

For example, in this graph, $\deg(1) = 3, \deg(2) = 5$, and $\deg(3) = 2$:



Let us now (as Euler did 286 years ago) imagine a connected graph with an Eulerian circuit. As we traverse along that circuit, we have to leave every visited vertex by a different edge than we entered (even if we have loops!). Hence, the degree of every vertex must be even. This is the first implication in the following theorem:

Theorem 7 (Euler, 1736). *A finite connected graph is Eulerian if and only if all its vertex degrees are even.*

Proof. We already showed the easy direction, so assume now that $G = (V, E, \iota)$ is a finite connected graph with only even vertex degrees. To simplify the argument, assume also that the graph contains no loops. Consider a trail $\mathcal{T} = v_0 e_1 v_1 e_2 v_2 \dots e_k v_k$ of maximal length (i.e. such that no edge is used twice, and such that there is no strictly longer trail in G). Since this trail cannot be extended by an additional edge at v_k , all edges incident to v_k must already be in \mathcal{T} . We now show that this implies $v_0 = v_k$.

Assume that there are $2s$ ($s \in \mathbb{N}$) edges incident to v_k . One of those edges is required to be e_k . Where do the other $2s - 1$ go? If v_k occurs as an internal vertex of \mathcal{T} , that is, say $v_j = v_k$ for $1 \leq j < k$, then e_j and e_{j+1} must both be edges incident to v_k . Thus, any internal occurrence of v_k takes up 2 edges (We are really using here that there are no self-loops). This explains where

$2s - 2$ of the edges at v_k go, leaving one edge – and that one has to go out of v_0 to $v_1 \neq v_k$. Hence $v_0 = v_k$, and the trail \mathcal{T} is really a circuit.

Next, we prove that this trail \mathcal{T} is Eulerian: Assume for a contradiction that it wasn't. Then there is an edge e , say with endpoints $\iota(e) = \{x, y\}$, that is not on \mathcal{T} , but one of its endpoints, say $x \in V$, is – this follows from connectedness. Hence, x occurs in some position on \mathcal{T} , say $x = v_i$. Now consider the following trail:

$$\mathcal{T}' := y e x e_{i+1} v_{i+1} \dots e_k v_k e_1 v_1 \dots e_i v_i.$$

This is indeed a trail, since e was not on \mathcal{T} , and it contains $k + 1$ edges, contradicting the assumption that \mathcal{T} is of maximal length. Therefore \mathcal{T} is the desired Eulerian trail.

Finally, what happens if G contains loops? Simply cut off the loops (this doesn't change the connectedness or the parity of the degrees by Definition 6) and consider the Eulerian trail \mathcal{T} obtained from the proof in the loop-less case. Now walk alongside \mathcal{T} . As \mathcal{T} necessarily passes through all vertices in the graph, stop whenever you reach a vertex for the first time, and traverse all loops at that vertex, before continuing with \mathcal{T} . \square

We also have the following corollary:

Corollary 8. *A finite connected graph admits an Eulerian trail if and only if either 0 or 2 of its vertices have odd degree.*

Proof. If all vertices have even degree, there is nothing left to show. Otherwise, introduce a new edge connecting the two vertices with odd degree. In the modified graph, all degrees are even, hence there is an Eulerian circuit. Removing the newly inserted edge from the circuit turns it into an Eulerian trail.

Conversely, if a graph admits an Eulerian trail that is closed, i.e. a circuit, we simply apply the theorem. Otherwise, such an Eulerian trail must have a starting and an ending vertex both having odd degree. \square

In particular, the bridges of Königsberg did not allow for walking through the city along an Eulerian trail, as all four vertices in the corresponding graph have odd degree. (If you're curious, check out Kaliningrad on google maps to see whether it's possible nowadays).

Corollary 8 might raise the question “What happens if a graph has exactly one vertex of odd degree?” – as it turns out, this can't happen:

Lemma 9 (Handshake lemma). *Let $G = (V, E, \iota)$ be a finite graph. Then*

$$2|E| = \sum_{v \in V} \deg(v) \tag{1}$$


In particular, G must contain an even number of vertices of odd degree.

Proof. We use a trick called *double counting*: On the one hand, every edge contributes two half-edges, thus there are $2|E|$ half-edges in total. On the other hand, every vertex contributes $\deg(v)$ half-edges by definition, and thus both sides of equation (1) count the same quantity. \square

Food for thought: If $G = (V, E, \iota)$ is a connected graph with 4 vertices of odd degree, can we always partition the edge set as $E = E_1 \uplus E_2$ such that $(V, E_1, \iota|_{E_1})$ and $(V, E_2, \iota|_{E_2})$ admit Eulerian trails? In other words, is it possible to colour every edge either red or blue such that the graph with red edges and the graph with blue edges each contain an Eulerian trail? What if we have $2k$ vertices of odd degree (for $k \geq 2$)?

Lecture 2 Simple graphs and subgraphs

For the bridges of Königsberg problem last lecture, we considered multigraphs, as the relevant information was *how* the landmasses are connected with each other. However, often in applications it is only relevant *if* vertices are neighbours to each other. To illustrate this, consider the *Towers of Hanoi* (Here, n disks of descending sizes are stacked on one of three pegs. The aim is to move the entire stack to another peg, by only moving one disk at a time. Moreover, smaller disks have to go on top of larger disks). While it is relatively easy to solve the Towers of Hanoi, more interesting questions (such as the fastest solve path) require a better framework for analysis. In this case, we can label the pegs by 1,2,3, and encode any state by a unique string of length n over the digits $\{1,2,3\}$ (where the first digit gives the position of the largest disk, \dots , and the last digit indicates the position of the smallest disk). Moreover, the allowed moves (all of them are reversible) induce a graph structure on the set of those strings by connecting two strings if the corresponding states can be transformed by a single move, giving rise to the Hanoi graph H^n (here shown for $n = 3$):



In other words, we're often in a situation where multiple edges or loops don't matter, so we could ignore them.

Definition 10. A *simple graph* is a multigraph without double edges or loops. Equivalently, it is a pair $G = (V, E)$ where $E \subseteq \mathcal{P}_2(V)$.

Let's discuss this definition: A multigraph $G = (V, E, \iota)$ has no double edges or loops iff the incidence map ι sends no edge to a single-vertex set, and

doesn't send any two edges to the same set of vertices. Hence, a graph is simple iff ι is an injective map $E \rightarrow \mathcal{P}_2(V)$. But in that case, we can directly identify an edge e with its image $\iota(e)$, so we can just as well assume $E \subseteq \mathcal{P}_2(V)$.

Since every simple graph is a multigraph, everything we did in the last lecture also applies to simple graphs, and we will use the same terminology as introduced then.

Most of the remaining lectures will be concerned with simple graphs, but we will occasionally require multigraphs as well. One of the great mathematical benefits of simple graphs is that they are a much more restricted class of graphs, allowing us to make statements such as:

Lemma 11. *Any simple graph on n vertices has at most $\binom{n}{2}$ edges.*

Proof. There are $\binom{n}{2}$ different 2-element subsets of V , hence $|E| \leq \binom{n}{2}$. \square

From this lemma, one might suspect that there exist only a finite number of different simple graphs on n vertices. This is true, but depends on what we mean by "different" – graphs (V, E) and (V', E') might "look the same", but be different according to definition simply because the sets V and V' are not the same. There are two ways to circumvent this - one by fixing the underlying vertex set, leading to labelled graphs, the other by formalising what is meant by "looking the same", leading to the notion of isomorphic graphs.

Definition 12. A *labelled graph* is a graph with a fixed vertex set, commonly (if V is finite) $V = \{1, 2, \dots, n\}$.

Corollary 13. *There are $2^{\binom{n}{2}}$ different labelled graphs on n vertices.*

Proof. Since $V = \{1, \dots, n\}$ is fixed, graphs on this vertex set are the same iff their edge sets $E \subseteq \mathcal{P}_2(V)$ coincide. Conversely, any subset $E \subseteq \mathcal{P}_2(V)$ defines a labelled graph on V . Thus, the labelled graphs on n vertices are in bijection with the power set of $\mathcal{P}_2(V)$, hence there are $2^{\binom{n}{2}}$ of them. \square


Definition 14. Let $G = (V, E)$ and $G' = (V', E')$ be simple graphs. A *morphism* $\varphi : G \rightarrow G'$ of simple graphs is a map $\varphi : V \rightarrow V'$ such that $\{\varphi(v), \varphi(w)\} \in E'$ whenever $\{v, w\} \in E$.

Warning: This notion of graph morphisms is only viable for simple graphs. While it is possible to define graph morphisms for multigraphs as well, the definition is considerably different.

Observe that every simple graph $G = (V, E)$ comes with an *identity morphism* id_G : Indeed, the map $\text{id} : V \rightarrow V$ which sends every vertex to itself satisfies

the condition in Definition 14. Moreover, if G, G', G'' are simple graphs with vertex sets V, V', V'' respectively, and $\varphi : G \rightarrow G'$ and $\varphi' : G' \rightarrow G''$ are morphisms, then there is a morphism $\varphi' \circ \varphi : G \rightarrow G''$ given by the map $\varphi' \circ \varphi : V \rightarrow V''$. (The students knowing what a category is might now verify that simple graphs with morphisms do indeed form a category, all other students don't need to care as it will play no further role in this course).

Example 15. Here is an example of a graph morphism, the dashed arrow indicating which vertices in the left graph are sent to which vertices on the right:




Definition 16. Two simple graphs G, G' as above are *isomorphic* if there is a morphism $\varphi : G \rightarrow G'$ and a morphism $\psi : G' \rightarrow G$ such that $\varphi \circ \psi = \text{id}_{G'}$ and $\psi \circ \varphi = \text{id}_G$. In this case, we call φ an *isomorphism*.

It follows from this definition that a morphism φ is an isomorphism if and only if it is a bijection on vertices, such that additionally $\{\varphi(v), \varphi(w)\} \in e'$ if and only if $\{v, w\} \in E$.


Being isomorphic provides an equivalence relation on the class of all simple graphs, and we will not always distinguish rigorously between an isomorphism class (i.e. an equivalence class under this relation) and a graph. So, a statement like “ G contains a copy of G' ” will mean: “ G contains a graph isomorphic to G' ”.

Observe that graph-theoretic notions are typically invariant under isomorphisms, this includes e.g. vertex degrees, connectedness, Eulerianity, and many more we have yet to encounter.

Example 17. To further illustrate the point of isomorphic graphs, here is the complete list of 11 non-isomorphic graphs on 4 vertices:



As we can see, the number of non-isomorphic graphs on 4 vertices (11) is different from the number of labelled graphs on 4 vertices, which is $2^{\binom{4}{2}} = 64$. For example, the following two graphs are isomorphic, but different as labelled graphs:



There is no exact formula for the number g_n of non-isomorphic graphs on n vertices, however there are asymptotic expansions, like the following, where we write $f(n) \sim g(n)$ to express $\lim_{n \rightarrow \infty} f(n)/g(n) = 1$.


Theorem 18. *The number g_n satisfies*

$$g_n = \frac{2^{\binom{n}{2}}}{n!} \left(1 + \frac{n^2 - n}{2^{n-1}} + O\left(\frac{n^3}{2^{3n/2}}\right) \right).$$

In particular, $g_n \sim 2^{\binom{n}{2}}/n!$.

Let us turn our attention towards special graphs, that is, isomorphism classes of graphs that turn up very often:

- The complete graphs on n vertices, denoted K_n : Those are the graphs containing all $\binom{n}{2}$ potential edges.
- The path graphs of length ℓ , denoted P_ℓ : If we label the vertices by $0, 1, \dots, \ell$ from left to right, our edges are precisely of the form $\{i - 1, i\}$ for all $i = 1, \dots, \ell$.
- The cycle graphs on n vertices, denoted C_n : This is a path of length $n - 1$ with an additional edge joining the first and last vertex.
- The complete bipartite graphs on $a + b$ vertices, denoted $K_{a,b}$: This graph's vertex set is partitioned into two disjoint sets, $V = V_1 \uplus V_2$, where $|V_1| = a$ and $|V_2| = b$. There are no edges between two vertices both belonging to the same set of the partition, but all possible edges between V_1 and V_2 .
- The complete multipartite graphs K_{a_1, \dots, a_r} : Fix $r \in \mathbb{N}$, and for a given number of vertices n choose r positive integers a_1, \dots, a_r such that $a_1 + \dots + a_r = n$. The graph's vertex set is a partition $V = \bigsqcup_{i=1}^r V_i$ with $|V_i| = a_i$, where two vertices are neighbours iff they belong to different sets of this partition. For $r = 2$, we get the complete bipartite graphs, for $r = n$ we get the complete graph on n vertices.



The figure above shows, from left to right, K_6, P_3, C_5 and $K_{2,3}$.

Proposition 19. *The complete multipartite graph K_{a_1, \dots, a_r} on n vertices has $\frac{1}{2}(n^2 - a_1^2 - \dots - a_r^2)$ edges.*

Proof. We use the handshaking lemma (Lemma 9) from Lecture 1. Since a vertex in V_i has one edge to every vertex not in V_i , it has degree $n - a_i$, and there are a_i such vertices. Hence,

$$2|E| = \sum_{i=1}^r a_i(n - a_i) = n \sum_{i=1}^r a_i - \sum_{i=1}^r a_i^2 = n^2 - \sum_{i=1}^r a_i^2,$$

proving the claim. □

In particular, $K_{a,b}$ has $\frac{1}{2}(n^2 - a^2 - b^2) = ab$ edges, and $K_n = K_{1,\dots,1}$ has $\frac{1}{2}(n^2 - n) = \binom{n}{2}$ edges.

Let us now introduce the important concept of subgraphs:

Definition 20. Let $G = (V, E)$ be a simple graph. A simple graph $H = (V', E')$ is a *subgraph* of G if $V' \subseteq V$ and $E' \subseteq E$. An *induced subgraph* is a subgraph $H = (V', E')$ of G such that $E' = \{\{x, y\} \in E : x, y \in V'\}$.

Several remarks are necessary for the proper appreciation of this definition:

Remark 21. The definition requires H to be a graph in its own right. It is thus not good enough to choose any subsets $V' \subseteq V$ and $E' \subseteq E$ to obtain a subgraph; one has to make sure these choices are compatible: including an edge in E' requires that its endpoints are included in V' .

Remark 22. By definition of an induced subgraph H , it is enough to choose V' , since then the edge set of H must contain all edges in G with both their endpoints in V' . For this reason, we will also write $G[V']$ for the subgraph induced by V' .

Remark 23. The notion of (induced) subgraphs extends to multigraphs in intuitively the same meaning, but for a formal definition, one has to make sure that the edges are incident to the same vertices in both graphs. That is, for $H = (V', E', \iota')$ to be a sub-multigraph of $G = (V, E, \iota)$, we have to require that H is a multigraph, that $V' \subseteq V$ and $E' \subseteq E$, but also that $\iota' = \iota|_{E'}$.

What we introduced above as induced subgraph is sometimes called the vertex-induced subgraph. Indeed, there is the dual-notion of the edge-induced subgraph:

Definition 24. Let $G = (V, E)$ be a simple graph, $E' \subseteq E$ a subset. The edge-induced subgraph $G \langle E' \rangle$ is the subgraph (V', E') where $V' := \{v \in V : v \text{ is incident to some } e \in E'\}$.

Finally, the subgraphs that include the entire vertex set of the host graph are often of special interest and therefore deserve a name of their own:

Definition 25. A subgraph $H = (V', E')$ of $G = (V, E)$ is called *spanning subgraph* if $V = V'$.

The same definition applies when G or G and H are multigraphs. If G is a simple graph, then the only subgraph of G that is both spanning and induced is G itself. It follows that a spanning subgraph H of G is uniquely determined by the choice of the set $E' \subseteq E$, hence G has $2^{|E|}$ spanning subgraphs.

Lecture 3 Trees

Definition 26. A *tree* is a graph $T = (V, E)$ that is both connected and contains no cycles.

Observe how in a multigraph, any loops or parallel edges introduce cycles immediately, so restricting the definition to simple graphs makes no difference.

Trees occur everywhere! As an example, consider QuickSort, an algorithm sorting a list of numbers into the correct (say, ascending) order. The algorithm works as follows:

1. Fix an arbitrary pivot element from the list.
2. Compare all non-pivot elements a with the pivot, and if the pivot is larger, move a to the left of the pivot; but if the pivot is smaller, move a to the right of the pivot. (After this step, at least the pivot element will be in its correct position, but the lists to its left and right might not be).
3. Repeat steps 1 and 2 for both sublists to either side of the pivot.

Example 27. Consider the list $(3, 7, 1, 4, 9, 8, 6, 2, 5)$, and let's choose 7 as our pivot element. After rearranging, the list might look like


$$(3, 1, 4, 6, 2, 5, 7, 9, 8)$$

– this depends on how exactly the partitioning step is implemented. Now, the 7 is in correct position, and we are left to sort $(3, 1, 4, 6, 2, 5)$ and $(9, 8)$ independently of each other, and might choose 3 and 8 as pivots for the left and right part, respectively. After rearranging with respect to the pivots, we could then end up with

$$(1, 2, 3, 4, 6, 5, 7, 8, 9)$$

and it only remains to sort the part to the right of 3 and to the left of 7, so $(4, 6, 5)$. Let's pick 5 as the pivot element, and obtain $(4, 5, 6)$ after rearranging, which sorts the list as required.

We might also choose to represent the algorithm in a more picturesque way:



As you can see, the final result is a labelled tree (with some additional properties, like the choice of a root node, and an orientation between left and right children, but we will choose to ignore this).

Lemma 28. *Every finite tree with at least 2 vertices contains at least two vertices of degree 1. Such vertices are called leaves.*

Proof. Let T be a finite tree on at least 2 vertices. Consider a path $P = xe_1x_1e_2\dots y$ of maximum length in T . This path begins with a vertex x , and ends with a different vertex y . Assume that one of x, y (w.l.o.g. x) has degree at least 2. Thus x has a neighbour w different from x_1 . If w is not contained in P , then the path going from w to x and then via P to y is strictly longer than P , contradicting the maximality of P . Hence, w is in P , but then T contains a cycle: Starting at x , follow along P until w , then go back to x via the edge $\{x, w\}$.

From this contradiction, we conclude $\deg(x) = \deg(y) = 1$. □

Proposition 29. *Any tree on n vertices has $n - 1$ edges.*

Proof. We use induction over the number of vertices. There is exactly one tree (up to isomorphism) consisting of exactly one vertex, \bullet , and it contains no edges. Now, assume that the statement is true for all trees on $n \geq 1$ vertices, and consider an arbitrary tree $T = (V, E)$ on $n + 1$ vertices. By Lemma 28, there is a vertex x in T with $\deg(x) = 1$ (so, it has a unique neighbour which we will call y). Remove x from T by considering the induced subgraph $T' := T[V \setminus \{x\}]$. This removes exactly the vertex x and the edge $\{x, y\}$ from T . Then T' has n vertices and therefore $n - 1$ edges, so T had $n + 1$ vertices and n edges, as required. □

Theorem 30 (Cayley). *There are n^{n-2} labelled trees on n vertices.*

Proof. We will provide a bijection between labelled trees on n vertices and sequences of labels of length $n - 2$, so called Prüfer sequences. This bijection

will take the form of two algorithms, one to convert a tree into a sequence, the other to get the tree back from the sequence. Since there are n^{n-2} Prüfer sequences, this will prove the theorem. To avoid the annoying small cases, assume $n \geq 2$.

Algorithm 1. Let T be a tree on n vertices, with labels $\{1, \dots, n\}$. As long as T has at least 3 vertices, remove the leaf with the smallest label and write down the label of its (unique!) neighbour as the next term in the sequence. Stop when there are only 2 vertices in T left.

Algorithm 2. Let $A = (a_1, \dots, a_{n-2})$ be a sequence of numbers $a_s \in 1, \dots, n$ for $s \in 1, \dots, n-2$. To each $i = 1, \dots, n$, let d_i be one plus the number of times i occurs in the Prüfer sequence A . Now, for s from 1 to $n-2$, find the smallest $j \in \{1, \dots, n\}$ such that $d_j = 1$ (and therefore automatically also $j \neq a_s$), draw an edge between a_s and j , and reduce d_{a_s} and d_j by 1 each. In the end, exactly two vertices $u, v \in \{1, \dots, n\}$ remain with $d_u, d_v > 0$ (in fact, $d_u = d_v = 1$), and we connect u and v with one last edge.


To conclude the proof it is necessary to show that applying algorithm 1 and then algorithm 2 to a labelled tree T returns this labelled tree, and that applying algorithm 2 and then algorithm 1 to a Prüfer sequence will return the sequence. This will be done in the exercises. \square

Example 31. The tree we obtained from the introductory example has Prüfer sequence $(1, 3, 5, 5, 3, 7, 8)$.

Theorem 32 (Characterisation of trees). *The following statements are equivalent:*

- (1) T is a tree.
- (2) For any two vertices x, y of T there exists a unique path connecting x and y .
- (3) T is edge-minimal among connected graphs, i.e. removing any edge from T creates at least two connected components.
- (4) T is edge-maximal among cycle-free graphs, i.e. adding an edge between two non-neighbouring vertices introduces a cycle.

Proof. Let $T = (V, E)$ be a simple graph. We will show the following implications:



(1) \implies (2): Assume T is a tree, $x, y \in V$. Since T is connected, there is at least one path from x to y . If there was more than one path, the union of two of those paths would create a cycle in T . Hence, the path is unique.

(2) \implies (3): Consider any edge $\{x, y\} \in E$. By (2), this edge is the unique path from x to y , hence removing it creates one connected component containing x and one containing y .

(2) \implies (4): Consider two non-neighbouring vertices $x, y \in V$. By (2), there is a unique path in T from x to y . Hence, introducing the new edge $\{x, y\}$ creates a cycle (by concatenating the path from x to y with the new edge).


(3) \implies (1): If T is edge-minimal among connected graphs, then T is in particular connected. Assume now that T were to contain a cycle. Deleting any edge on this cycle would not disconnect T , contradicting edge-minimality. Hence T is a tree.

(4) \implies (1): By assumption (4), T is cycle-free. Assume that T is disconnected, then adding an edge between two different connected components does not introduce a cycle, contradicting edge-maximality. Hence T is a tree. \square

Recall from last lecture that given a graph G , a spanning subgraph of G is a subgraph containing all vertices of G .

Definition 33. Let G be a graph (simple or multi-). A *spanning tree* of G is a spanning subgraph that is also a tree.

Example 34. The following graph contains a spanning tree whose edges are marked in red:



If the host graph G is disconnected, then any spanning subgraph must be disconnected as well, so there can be no spanning trees of G . The following

theorem shows that the obviously necessary condition of connectivity is also sufficient:

Theorem 35. *Let G be a connected (multi-)graph. Then G contains a spanning tree.*

The existence of spanning trees for finite connected G is a straightforward consequence of Theorem 32 – if G is finite, there have to be edge-minimal connected and edge-maximal cycle-free subgraphs! For infinite graphs, the situation is less clear, and in fact Theorem 35 is equivalent to the axiom of choice. We will give a proof relying on Zorn’s lemma (itself equivalent to the axiom of choice).

Theorem 36 (Zorn’s lemma). *Let (A, \leq) be a non-empty partially ordered set. A subset $C \subseteq A$ is a chain if for any two elements $c, c' \in C$, we either have $c \leq c'$ or $c' \leq c$. Assume that for every chain C in A there exists an upper bound $b \in A$ (i.e. $c \leq b$ for all $c \in C$). Then, there exists $m \in A$ with $m \leq a \implies m = a$.*

For our purposes, let A be the set of all cycle-free spanning subgraphs H of $G = (V, E, \iota)$, and for $H, H' \in A$, define $H \leq H'$ if H is a subgraph of H' . This gives (A, \leq) the structure of a partially ordered set. Furthermore, A is non-empty since $H = (V, \emptyset)$ is spanning and cycle-free. Let C be a chain in A , consisting of elements $H_i = (V, E_i)$ for $i \in I$. Define $H_b := (V, \bigcup_{i \in I} E_i)$. We want to show that H_b is an upper bound for C .

By construction, H_b is a spanning subgraph of G . Assume it contains a cycle consisting of edges e_1, \dots, e_r . Then, for every $\ell = 1, \dots, r$ there is an $H_{i(\ell)}$ that contains e_ℓ in C . Since C is a chain, one of the graphs $H_{i(\ell)}$ – say, for $\ell = j$ – contains all of the other graphs for different values of ℓ . Hence, $H_{i(j)}$ must contain the cycle e_1, \dots, e_r . This is a contradiction, since $H_{i(j)} \in C \subseteq A$ must be cycle-free. Hence H_b is cycle-free, and $H_b \in A$. Finally, since $E_i \subseteq \bigcup_{i \in I} E_i$, we have $H_i \leq H_b$ for all $H_i \in C$.

But now all the assumptions of Zorn’s lemma are satisfied! Hence there must be a maximal (with respect to \leq , i.e. edge-maximal) cycle-free spanning subgraph H of G , and by Theorem 32, (4), this H is a spanning tree.

Lecture 4 Counting spanning trees

We begin with a definition:

Definition 37. Let G be a labelled graph. The *complexity* of G is the number of spanning trees of G , denoted $t(G)$.


As a first example, note that a labelled tree on n vertices can equivalently be considered a spanning tree of a labelled K_n , so Cayley's theorem immediately yields $t(K_n) = n^{n-2}$.

In this lecture, we want to obtain one of the classical combinatorial theorems in graph theory, answering the question: Given a labelled, connected, simple graph G , how many spanning trees does G have? To provide an answer to this question, we will develop some of the basic language in *algebraic graph theory*.

Let us begin by introducing two ways of encoding a graph algebraically. For the entire lecture, we will assume that $G = (V, E)$ is a simple graph with a finite labelled vertex set $\{1, \dots, n\}$ and an edge set $\{e_1, \dots, e_m\}$.

Definition 38. The *adjacency matrix* A of a graph G is the $n \times n$ -matrix having entries $A_{ij} = 1$ if i and j are neighbours, and $A_{ij} = 0$ otherwise.

Example 39. For example, the graph



has adjacency matrix

$$\begin{pmatrix} 0 & 1 & 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

Observe that, since “being neighbours” is a symmetric relation on vertices, the matrix A is symmetric. Moreover, no vertex is a neighbour to itself, so all diagonal entries of A vanish. Finally, it follows from the spectral theorem for symmetric matrices that all eigenvalues of A are real.

There is another important way to capture the relevant information about a simple graph in a matrix; however, it turns out to be more useful to *orient* the edges for this: So, for every edge e_1, \dots, e_m , designate one of its endpoint as the positive end, and the other as the negative end. Of course, we will later on have to make sure that whatever we say about a graph does not depend on the chosen orientation!

Definition 40. The *incidence matrix* D of a graph G with a fixed orientation is the $n \times m$ -matrix having entries D_{ij} ($1 \leq i \leq n, 1 \leq j \leq m$), where

$$D_{ij} = \begin{cases} 1 & \text{if } i \text{ is the positive end of } v_j \\ -1 & \text{if } i \text{ is the negative end of } v_j \\ 0 & \text{otherwise} \end{cases}$$

Example 41. The matrix

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & -1 & 1 & 1 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 1 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 & 0 & -1 & -1 & -1 \\ 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

is one possible incidence matrix for the graph G from Example 39. The edges were ordered in the following way:

$$\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 5\}, \{5, 6\}, \{6, 1\}, \{1, 4\}, \{1, 5\}, \{2, 5\}, \{3, 5\}.$$

The following lemma collects several useful properties of the incidence matrix:

Lemma 42. *Let D be the incidence matrix of a finite simple graph G . Then:*

- (i) *The sum of any column of D is zero, hence $\text{rank } D \leq n - 1$.*
- (ii) *If G is connected, then $\text{rank } D = n - 1$.*
- (iii) *If G has c components, then $\text{rank } D = n - c$.*

Proof. The column sum of D is zero since in any column, there is exactly one entry 1 and one entry -1 , representing the positive and negative endpoint of the edge to which the column is associated. Hence, taking the sum of all row-vectors in D is a non-trivial linear combination of 0, hence $\text{rank } D \leq n - 1$. To show part (ii), we show that this is (up to scaling) the only non-trivial linear combination of 0 with row-vectors. So, let r_i for $i = 1, \dots, n$ denote the row-vectors of D , and suppose we have a non-trivial linear combination $\sum_{i=1}^n \alpha_i r_i = 0$. Consider a row k for which $\alpha_k \neq 0$. In this row, there is

a non-zero entry in every column corresponding to an edge incident with the vertex k . For each of these columns, there is exactly one other row (say, r_ℓ) with a non-zero entry in that column, and the two entries will be of opposite signs. But, for our linear combination to yield zero, we thus require $\alpha_\ell = \alpha_k$. Hence we just proved that if $\alpha_k \neq 0$, then $\alpha_\ell = \alpha_k$ for all ℓ which are neighbours to k . But if G is connected, this argument extends to all of G , hence the linear combination is a scalar multiple of $\sum_{i=1}^n r_i = 0$. Thus, $\text{rank } D = n - 1$.

For part (iii), observe that it is possible to re-label the vertex and edge sets in such a way that D is in block-diagonal form, with every block being the incidence matrix of a connected component. The claim then follows from (ii). \square

Lemma 43. *Any square submatrix of an incidence matrix D has determinant 0, 1, or -1 .*

Proof. This is obviously true for square submatrices of size 1×1 . Assume that it holds true for submatrices of size $k \times k$, and consider (if possible within D) a square submatrix M of size $(k + 1) \times (k + 1)$. If every column of M has either two or no non-zero entries, then $\det M = 0$ (in the first case, every column sums to 0, in the second case $M = 0$). Otherwise, there's a column of M having exactly one non-zero entry. Expanding M along this column yields $\det M = \pm \det M'$ where M' is a $k \times k$ -submatrix of D . Hence $\det M \in \{-1, 0, 1\}$. \square

Observe also the following curious fact about D : If $S \subseteq E$ is a set of edges from G , and D_S denotes the matrix consisting of exactly those columns that correspond to edges in S , then D_S is the incidence matrix of the spanning subgraph (V, S) of G . In particular, if S contains exactly $n - 1$ edges, then by part (iii) in the previous lemma D_S has rank $n - 1$ iff (V, S) is a spanning tree of G . The following proposition shows that this is in fact a characterisation of spanning trees:

Proposition 44. *Let $S \subseteq E$ with $|S| = n - 1$. Let M denote any $(n - 1) \times (n - 1)$ -submatrix of the $n \times (n - 1)$ -matrix D_S . Then M is a regular matrix iff (V, S) is a spanning tree of G .*

Proof. We have already observed that D_S has rank $n - 1$ whenever (V, S) is a spanning tree. In that case, removing any row from D_S will create a non-singular square matrix M .

Conversely, if M is non-singular, then D_S contains at least $n - 1$ linearly independent rows and the same number of linearly independent columns. Hence $\text{rank } D_S = n - 1$ and (V, S) must be connected. But a connected graph on n vertices and $n - 1$ edges must be a tree. \square

Lemma 45. *Let G be a finite simple graph. Denote by A its adjacency matrix and by D its incidence matrix (with respect to some fixed orientation). Let Δ be a diagonal $n \times n$ -matrix having diagonal entries $\Delta_{ii} = \deg(i)$ for $1 \leq i \leq n$. Then,*

$$DD^t = \Delta - A$$

and this matrix is called the Laplacian matrix of G . In particular, DD^t is independent from the chosen orientation on G .

Proof. The entry $(DD^t)_{ij}$ is the Euclidean inner product of rows r_i and r_j of D . If $i = j$, this amounts to summing up the squares of the entries in r_i . As observed in the proof of the previous lemma, the non-zero entries in r_i are ± 1 , and there are $\deg(i)$ -many of them. Hence $(DD^t)_{ii} = \deg(i) = (\Delta - A)_{ii}$. For $i \neq j$, the rows r_i and r_j have a non-zero entry in the same column iff this column corresponds to an edge $\{i, j\}$ (and there can be at most one such edge). In that case, the scalar product evaluates $r_i \cdot r_j = (+1)(-1) = -1$, since the entries in this column must be of opposite signs. Hence

$$(DD^t)_{ij} = -1 = (\Delta - A)_{ij}$$

if $\{i, j\}$ is an edge, and

$$(DD^t)_{ij} = 0 = (\Delta - A)_{ij}$$

if $\{i, j\}$ is not an edge. □

Recall that a cofactor M_{ij} of a square-matrix M is the determinant of a submatrix for which one row (row i) and one column (column j) have been removed. The transpose of the matrix containing all those cofactors as its entry is called the adjugate matrix of M , denoted $\text{adj } M$.

Lemma 46. *Let $Q = DD^t = \Delta - A$ be the Laplacian matrix of G . Denote by J the $n \times n$ -matrix all of whose entries are 1. Then $\text{adj } Q$ is a scalar multiple of J .*

Proof. We observe first that $\text{rank } Q = \text{rank } D$. If G is disconnected, then by Lemma 42, $\text{rank } Q < n - 1$, so all cofactors will vanish, and $\text{adj } Q = 0$.

If G is connected, we have $\text{rank } Q = n - 1$. We know from linear algebra that $Q(\text{adj } Q) = (\det Q)I_n$, and the right-hand side is 0 since Q does not have full rank. So, $Q(\text{adj } Q) = 0$, which implies that every column vector of $\text{adj } Q$ is in the kernel of Q . This kernel, however, is one-dimensional by $\text{rank } Q = n - 1$, and spanned by $(1, \dots, 1)^t$: Indeed, taking the inner product between the i -th row of $Q = \Delta - A$ and $(1, \dots, 1)^t$ gives $\deg(i) - \sum_{j \sim i} 1 = 0$, where the sum ranges over all neighbours j of i . It follows that every column of $\text{adj } Q$ is a scalar multiple of $(1, \dots, 1)^t$. As Q is symmetric, $\text{adj } Q$ is as well, and all those scalar multiples have to be the same! □

We are now finally ready to state (and prove) the following theorem:

Theorem 47 (Kirchhoff's matrix-tree-theorem). *Let G be a finite simple graph having Laplacian matrix Q . Then $\text{adj } Q = t(G)J$. In other words, the number of spanning trees in G coincides with any cofactor of Q .*

Equivalently, if $\lambda_1, \dots, \lambda_{n-1}$ are the non-zero eigenvalues of Q , then

$$t(G) = \frac{1}{n} \lambda_1 \cdots \lambda_{n-1}.$$

The prove relies on mostly on the theory developed so far, but also on the following advanced theorem from linear algebra. To state it, assume that A is an $n \times m$ -matrix with $n \leq m$. Then (analogously to incidence matrices) for $S = \{s_1, \dots, s_n\} \subseteq \{1, \dots, m\}$, we write A_S to denote the $n \times n$ -submatrix of A consisting of exactly the columns s_1, \dots, s_n .

Theorem 48 (Cauchy-Binet). *Let A, B be two $n \times m$ -matrices, with $n \leq m$. Then,*

$$\det(AB^t) = \sum_{\substack{S \subseteq \{1, \dots, m\} \\ |S|=n}} \det A_S \det B_S.$$

Proof of the matrix-tree-theorem. We know from Lemma 46 that all cofactors coincide, so it is enough to focus on one particular cofactor. Let D be the incidence matrix of G for a fixed orientation, delete its last row, and denote by \tilde{D} the obtained matrix. Then $\det(\tilde{D}\tilde{D}^t)$ is a cofactor of $Q = DD^t$, and we can express this cofactor with the help of the Cauchy-Binet theorem:

$$\det \tilde{D}\tilde{D}^t = \sum_{\substack{S \subseteq E \\ |S|=n-1}} (\det \tilde{D}_S)^2.$$

The summands on the right-hand side are either 0 or 1, according to Lemma 43, and by Proposition 44 a summand is 1 iff (V, S) is a spanning tree of G . This concludes the proof of the first part. The second part follows from general statements about the characteristic polynomial of a matrix. \square

As an example, let's verify Cayley's formula once again: For $G = K_n$, we have the Laplacian matrix

$$Q = \Delta - A = \begin{pmatrix} n-1 & -1 & \dots & -1 \\ -1 & n-1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & -1 \\ -1 & \dots & -1 & n-1 \end{pmatrix}$$

It can be verified directly that $(1, \dots, 1)^t$ is an eigenvector to the eigenvalue 0, and that each of the $n - 1$ vectors of the form $(0, \dots, 0, 1, -1, 0, \dots, 0)^t$ is an eigenvector to the eigenvalue n . Hence

$$t(K_n) = \frac{1}{n}n^{n-1} = n^{n-2}.$$

Lecture 5 Weights and distances

In the previous two lectures, we proved that in a connected graph G , spanning trees exist, and found a way to count how many there are. The important question left open is: How do we actually find one? As it turns out, we can even answer this question in a more general setting, where we assume that every edge in G comes at a certain cost (the *weight* of an edge), and we want to find a spanning tree with minimal weight. Depending on the application at hand, these weights might be distance, resistance, capacity, For convenience, we will work with simple graphs (but everything is easily generalised to multigraphs).

Definition 49. A *weighted graph* is a finite simple graph $G = (V, E)$ together with a *weight function* $w : E \rightarrow (0, \infty)$. If $H = (V', E')$ is a subgraph of G , then its *weight* is defined as $w(H) := \sum_{e \in E'} w(e)$, the sum of the weights of the edges in H . A *minimum spanning tree* (MST) is a spanning tree T of G such that $w(T)$ is minimal among all spanning trees of G .

For a connected weighted graph G , we know that there exist finitely many spanning trees, and in particular, a minimum spanning tree (albeit not necessarily a unique one).


Lemma 50. *Let G be a connected weighted graph. If w assigns a different weight to every edge, the MST is unique.*

Proof. Assume that w assigns a different weight to every edge, and let $T_1 = (V, E_1), T_2 = (V, E_2)$ be two different minimum spanning trees. In particular, $E_1 \neq E_2$, and the set D , consisting of all the edges of G that are in exactly one of T_1 or T_2 , is nonempty. Pick $e \in D$ with $w(e)$ minimal. W.l.o.g. $e \in E_1$ and $e \notin E_2$. Adding e to T_2 creates a cycle, and somewhere on this cycle there must be an edge e' which is not in E_1 . Hence $e' \in D$. On the one hand, we now have $w(e) < w(e')$ by choice of e . On the other hand, $w(e') \leq w(e)$ as otherwise, T_2 would not be an MST. This gives a contradiction. \square


We will now proceed to present two famous algorithms for finding minimum spanning trees:

Prim's algorithm: Let $G = (V, E)$ be a connected weighted graph. Set $T = (\{v\}, \emptyset)$ for any vertex $v \in V$. As long as T is not a spanning subgraph of G , find an edge e of G between $V(T)$ and $V \setminus V(T)$ of minimal weight. Add this edge together with its endpoint in $V \setminus V(T)$ to T .

Example 51. Prim's algorithm, applied to the graph



takes the following steps, after being initialized in the top middle vertex (some steps have been executed simultaneously):



Theorem 52. Let $G = (V, E)$ be a connected weighted graph with weight function w . Prim's algorithm generates a minimum spanning tree.

Proof. Prim's algorithm generates a connected spanning subgraph T , and it is inductively clear that T contains $n - 1$ edges for $n = |V|$. Hence T is a spanning tree. To show that T is an MST, consider an MST T' . We show that $w(T) \leq w(T')$.


Suppose $T \neq T'$. Consider the earliest edge e that belongs to T but not to T' , according to the order in which Prim's algorithm added edges to T . Partition the vertex set into two parts V_1 and V_2 (and $V_1 \uplus V_2 = V$) such that $T[V_1]$ is the tree that Prim's algorithm had constructed before adding e . In particular, $T[V_1]$ is also a subgraph of T' . Hence there must be an edge f between V_1 and V_2 that occurs in T' , and $f \neq e$. Transform T' by adding the edge e and removing f for it. Since Prim's algorithm chose e over f , we have $w(e) \leq w(f)$, so this transformation did not increase $w(T')$.

Now repeat this procedure until $T = T'$, then $w(T') \geq w(T)$. \square

Remark 53. The runtime of Prim's algorithm depends on the implementation, specifically on how the graph is represented, and how one finds the next edge between $V(T)$ and $V \setminus V(T)$. A good implementation has a runtime of $O(|E| + |V| \log |V|)$.

Kruskal's algorithm: Let $G = (V, E)$ be a connected weighted graph. Let S be a list of edges of G , sorted in order of increasing weight. Let $T = (V, \emptyset)$. While T is disconnected, delete the first entry of S and add this edge to T unless doing so creates a cycle.

Example 54. Kruskal's algorithm, applied to the graph in Example 51, takes the following steps (with edges of the same weight being added simultaneously, if possible):



Theorem 55. Let $G = (V, E)$ be a connected weighted graph with weight function w . Kruskal's algorithm generates a minimum spanning tree.

Proof. Kruskal's algorithm generates a cycle-free spanning subgraph T that is also edge-maximal, and hence a spanning tree. If T is not a minimum spanning tree, let T' denote an MST that has the biggest possible number of edges in common with T . We will construct another MST that has even more edges in common with T to obtain a contradiction.

Let e be the earliest edge (according to Kruskal's algorithm) in T that is not contained in T' . Adding e to T' creates a cycle, and somewhere on this cycle there must be an edge f belonging to T' but not to T . Modifying T' by adding e and removing f yields another spanning tree T'' . Since T' is an MST, $w(T') \leq w(T'')$. However, since Kruskal's algorithm included e but not f , e must have come first in the ordered list S , hence $w(e) \leq w(f)$. Therefore also $w(T'') \leq w(T')$. This is only possible if $w(T') = w(T'')$, so T'' is an MST. However, T'' has more edges in common with T than T' has, which is the desired contradiction. \square

Kruskal's algorithm has a runtime of $O(|E| \log |V|)$.

As mentioned above, we can interpret edge-weights as distances. This gives rise to the following notion of distance in a graph:

Definition 56. Let $G = (V, E)$ be a weighted graph with weight function $w : E \rightarrow (0, \infty)$. For vertices $v, v' \in V$, we define the *graph distance* between

v and v' via

$$d_G(v, v') = \min \left\{ \sum_{e \in E(P)} w(e) : P \text{ is a path from } v \text{ to } v' \right\}$$

where we declare $d_G(v, v') = \infty$ if no such path exists (i.e. if v, v' are from different connected components). This definition extends also to non-weighted graphs by introducing the weight function $w(e) = 1$ for any edge $e \in E$.

As any reasonable notion of distance, graph distance satisfies the axiom of a *metric space*, that is the following three properties:

Lemma 57. *Let $G = (V, E)$ be a connected, weighted graph.*

- (i) *For all $v, v' \in V$, we have $d_G(v, v') \geq 0$. Moreover, $d_G(v, v') = 0$ iff $v = v'$.*
- (ii) *We have $d_G(v, v') = d_G(v', v)$ for all $v, v' \in V$.*
- (iii) *The triangular inequality holds: $d_G(v, v') + d_G(v', v'') \geq d_G(v, v'')$ for all $v, v', v'' \in V$.*

Proof. Statement (i) follows immediately from definition. If P is a path from v to v' , then traversing P in the opposite direction is a path from v' to v , and (ii) follows.

To show the triangular inequality, let P be a path of length $d_G(v, v')$ from v to v' , and let Q be a path of length $d_G(v', v'')$ from v' to v'' . Then by concatenating P and Q (i.e. first traverse all of P , then all of Q) we obtain a walk from v to v'' of length $d_G(v, v') + d_G(v', v'')$. We obtain a path from v to v'' by erasing any cycles this walk may contain, and in doing so, the length can only decrease. Hence there is a path of length at most $d_G(v, v') + d_G(v', v'')$, which shows that $d_G(v, v'') \leq d_G(v, v') + d_G(v', v'')$. \square

Definition 58. The *diameter* of a weighted graph $G = (V, E)$ is the maximum distance between any pair of vertices, i.e.

$$\text{diam}(G) = \max \{ d_G(v, v') : v, v' \in V \}.$$

Again, this definition extends to graphs without weight functions by setting $w(e) = 1$ for any edge $e \in E$.

Given a weighted graph, how can we actually find the distances between vertices? And how can we find the shortest paths, realising those distances? Luckily there's Dijkstra's algorithm. In the form stated here, it gives distances from a selected starting vertex to all other vertices in the graph.

Dijkstra's algorithm: Given a weighted graph $G = (V, E)$, select an initial vertex v_0 , and initialise a distance function $d(v_0, \cdot)$ by

$$d(v_0, v) = \begin{cases} 0 & \text{if } v = v_0 \\ \infty & \text{if } v \neq v_0 \end{cases}$$

Also define a set Q of unvisited vertices, which is initially defined to be $Q = V$. To begin with, consider v_0 to be the currently visited vertex. Now proceed as follows:

1. Remove the currently visited vertex v from the set Q .
2. For all neighbours v' of v in Q , check whether $d(v_0, v) + w(\{v, v'\}) < d(v_0, v')$. If this is the case, then going through v to v' provides a shorter path from v_0 to v' than previously known; hence we update $d(v_0, v')$ to be the smaller value $d(v_0, v) + w(\{v, v'\})$. Otherwise, no change to $d(v_0, \cdot)$ is needed.
3. Declare the new currently visited vertex to be the vertex v in Q with the smallest value $d(v_0, v)$.
4. Repeat steps 1 - 3 until $Q = \emptyset$, then return the distance function $d(v_0, \cdot)$.

Then, the returned distance function coincides with $d_G(v_0, v)$.

Remark 59. Dijkstra's algorithm can easily be modified:

- If only a certain distance $d_G(v, v')$ is needed, run the algorithm with $v_0 = v$, but stop as soon as the currently visited vertex is v' .
- If the actual path realising d_G is of interest, it is possible to modify step 2 as follows: Suppose you have a currently visited vertex v with a neighbour v' in Q , and $d(v_0, v) + w(\{v, v'\}) < d(v_0, v')$. Additionally to updating $d(v_0, v')$, assign to v' the precursor vertex v . Updating this precursor vertex together with $d(v_0, \cdot)$ enables us to backtrack the shortest path after the algorithm has terminated.

Example 60. Here's how Dijkstra's algorithm runs on the graph from Example 51, with vertices being labelled A, B, \dots, F from top left to bottom right:

Step	Current vertex	$d(A, \cdot)$					
		A	B	C	D	E	F
Init	–	0	∞	∞	∞	∞	∞
1	<i>A</i>	0	2	∞	3	∞	∞
2	<i>B</i>	0	2	5	3	6	4
3	<i>D</i>	0	2	5	3	6	4
4	<i>F</i>	0	2	5	3	5	4
5	<i>C</i>	0	2	5	3	5	4
6	<i>E</i>	0	2	5	3	5	4

With the right implementation, Dijkstra's algorithm has a runtime of $O(|E| + |V| \log |V|)$.

Lecture 6 Hamilton cycles

Recall from Definition 2 all the way back in Lecture 1, that a *cycle* in a graph G is a sequence $v_0e_1v_1e_2v_2\dots v_{k-1}e_kv_0$ such that all the vertices v_0, \dots, v_{k-1} and all the edges e_1, \dots, e_k are distinct.

Definition 61. Let $G = (V, E)$ be a finite simple graph. A *Hamilton cycle* in G is a cycle that visits every vertex of G . If G admits a Hamilton cycle, we also simply call G *Hamiltonian*.

Unlike with Eulerian circuits, where we could give a simple characterisation which tells us whether a graph contains one or not, there is no such criterion for hamiltonicity. In fact, the problem of determining whether a graph has a Hamilton cycle is an NP-complete problem.

Deciding hamiltonicity of a graph is a special case of the *travelling salesman problem*, where the task is to find a shortest/minimum weight Hamilton cycle in a weighted graph. This is usually asked on a complete graph to ensure the existence of Hamilton cycles, but this is unnecessary: one can simply draw in all the missing edges and set their weight to infinity. Since finding the shortest Hamilton cycle is (even intuitively) a harder task than determining if there is one, the travelling salesman problem also belongs to the class of NP-complete problems.

However, mathematicians have found sufficient conditions to guarantee the existence of Hamilton cycles, and we will take a look into some of them:

Theorem 62 (Dirac 1952). *Let $G = (V, E)$ be a simple graph on $n \geq 3$ vertices, such that every vertex has degree at least $n/2$. Then G is Hamiltonian.*

Proof. Assume $G = (V, E)$ with $|V| = n \geq 3$ and such that $\min_{v \in V} \deg(v) \geq n/2$. Such G cannot be disconnected, for if it were, vertices in the smallest component would violate the degree condition.


Now consider a path P of maximum length, say $P = v_0e_1v_1e_2v_2\dots v_{k-1}e_kv_k$. This path, being maximal, contains all neighbours of v_0 and all of v_k . However, by the degree condition, there are at least n such neighbours, hence there must be an edge e_i (between v_{i-1} and v_i) such that v_i is a neighbour of v_0 and v_{i-1} is a neighbour of v_k .

From this, construct a cycle C as follows: Starting in v_i , traverse along P all the way to v_k , then take the edge connecting v_k with v_{i-1} . From there, traverse backwards along P all the way to v_0 , and then complete the cycle with the edge from v_0 to v_i .

It remains to show that C is indeed a Hamilton cycle: Assume C is not; then (since G is connected) there exists a vertex v neighboured to a vertex on C (say, to v_j), but not itself on C . Construct a path P' by taking a spanning

tree (i.e. a path) of C having a leaf in v_j , and prolong this path by adding v together with $\{v_j, v\}$ to it. But now P' is a path strictly longer than P , which is a contradiction. Hence C is a Hamilton cycle. \square

Remark 63. The uniform lower bound on the minimal degree $\delta(G) = \min_{v \in V} \deg(v)$ is optimal, i.e. replacing $n/2$ by any smaller integer k allows for counterexamples. Perhaps simplest is the following construction: Set $k := \lfloor (n-1)/2 \rfloor$, and glue together two copies of K_{k-1} by identifying one vertex in one copy with a vertex in the other copy. This graph has minimum degree k , but can not contain a Hamilton cycle (since removing the vertex “in the middle” creates two components – an impossibility for Hamiltonian graphs!).



We immediately get the following generalisation of Dirac’s theorem:

Theorem 64 (Ore, 1960). *Let G be a graph on $n \geq 3$ vertices, such that for any non-adjacent pair of vertices v, w , we have $\deg(v) + \deg(w) \geq n$. Then G is Hamiltonian.*


Proof. The proof works in the same way as the proof for Dirac’s theorem. \square

Even more general is the Bondy-Chvátal theorem, which we will state next (but we require a definition first).

Definition 65. Let $G = (V, E)$ be a finite simple graph. The *closure* $\text{clos}(G)$ of G is defined as the result of the following procedure: For every pair of non-adjacent vertices v, w , draw the edge $\{v, w\}$ if $\deg(v) + \deg(w) \geq n$. Stop when there is no such pair anymore in the graph.

Example 66. If G is a graph on n vertices satisfying the conditions of Ore’s theorem, then $\text{clos}(G) = K_n$.

Observe also how for the following graph, whose closure is K_6 , vertices 3 and 6 can only be connected by an edge after the first few edges were inserted:



Lemma 67. *The closure of a graph does not depend on the order in which the edges are inserted (and is thus well-defined).*

Proof. Let $G = (V, E)$ be a finite simple graph. Assume we construct the closure of G in two different ways, where one way adds the edges e_1, e_2, \dots, e_k to G (in this order), and the other way adds the edges $f_1, f_2, \dots, f_{k'}$ to G – observe that not necessarily $k = k'$.

Then there is a smallest j s.t. $e_j \neq f_j$ (otherwise, one of the sequences contains the entire other sequence, which is impossible). Let G_{j-1} be the graph obtained by adding the edges $e_1 = f_1, \dots, e_{j-1} = f_{j-1}$, and let $e_j = \{v, w\}$. Then, $\deg_{G_{j-1}}(v) + \deg_{G_{j-1}}(w) \geq n$. Since those degrees will not decrease when adding the next edges f_j, \dots , eventually $\{v, w\}$ will be inserted by some f_ℓ . But then we can modify the sequence $f_1, f_2, \dots, f_{k'}$ by inserting the edge f_ℓ instead of f_j , and inserting the edges $f_j, \dots, f_{\ell-1}$ one step later. Now repeat this modification procedure, until the sequences coincide completely. \square

Theorem 68 (Bondy-Chvátal, 1972). *A graph G is Hamiltonian if and only if $\text{clos}(G)$ is.*

Proof. Since G is a spanning subgraph of $\text{clos}(G)$, the latter is Hamiltonian if the former is.

For the other direction, assume we construct $\text{clos}(G)$ from G by adding edges e_1, \dots, e_k , yielding graphs $G_1, \dots, G_{k-1}, G_k = \text{clos}(G)$. If some G_j is Hamiltonian, then all following graphs will be Hamiltonian as well. So assume that for some j , G_{j+1} is Hamiltonian, but G_j is not, and let $e_{j+1} = \{v, w\}$ be the unique edge by which G_j and G_{j+1} differ. In particular, the Hamilton cycle in G_{j+1} must contain $\{v, w\}$. Hence G_j contains a path P from v to w visiting all vertices in G (often called a *Hamilton path*). But since $\{v, w\}$ is the next edge added, we have $\deg_{G_j}(v) + \deg_{G_j}(w) \geq n$, and all neighbours to v and w must lie on P . But now, we can once again use the construction in the proof of Theorem 62, to show that G_j already contained a Hamilton cycle! \square

Another way of generalising the theorems of Dirac and Ore is to replace the uniform condition on the minimum degree, and instead requiring to be above a certain degree sequence:

Definition 69. Let G be a simple graph on n vertices. Then, the *degree sequence* of G is a finite sequence (d_1, \dots, d_n) containing the vertex-degrees of G in non-descending order, i.e. such that $d_1 \leq d_2 \leq \dots \leq d_n$. An arbitrary sequence of integers (a_1, \dots, a_n) is called *Hamiltonian* if all graphs with a degree sequence (d_1, \dots, d_n) satisfying $d_i \geq a_i$ for all $i = 1, \dots, n$ are Hamiltonian.

Observe that isomorphic graphs have the same degree sequence, but having the same degree sequence does not imply being isomorphic.

Theorem 70 (Chvátal, 1972). *Let $n \geq 3$. An integer sequence (a_1, \dots, a_n) with $0 \leq a_1 \leq \dots \leq a_n < n$ is Hamiltonian iff for every $i < n/2$, we have $a_i \leq i \implies a_{n-i} \geq n - i$.*

We turn our attention to two more examples:


The d -dimensional cube. Fix an integer $d \geq 2$. The d -dimensional cube has 2^d vertices, labelled by binary strings of length d – in other words, $V = \{0, 1\}^{\times d}$. Two vertices are neighbours iff the corresponding binary strings differ in exactly one position. It follows immediately from this definition that all vertices have degree d .

Proposition 71. *The d -dimensional cube is Hamiltonian for all $d \geq 2$.*

Proof. We use induction over d . Denote the d -dimensional cube by G_d . For $d = 2$, this is isomorphic to C_4 , which is trivially Hamiltonian, and any Hamilton cycle in G_2 will contain all of the edges, so in particular the edge between $(0, 0)$ and $(1, 0)$. Assume that for some $d \geq 2$, there is a Hamilton cycle in G_d that contains the edge between $(0, 0, \dots, 0)$ and $(1, 0, \dots, 0)$. By deleting this edge from the Hamilton cycle, we get a Hamilton path from $(0, 0, \dots, 0)$ to $(1, 0, \dots, 0)$.

Now consider G_{d+1} . In G_{d+1} , there are is a copy of G_d (namely the subgraph induced by all vertices whose first entry is a 0), and a second disjoint copy of G_d (the subgraph induced by all vertices whose first entry is a 1). Using this, we build a Hamilton cycle including the edge between $(0, 0, \dots, 0)$ and $(1, 0, \dots, 0)$ in G_{d+1} as follows: Start in $(0, 0, 0, \dots, 0)$. Traverse the Hamilton path to $(0, 1, 0, \dots, 0)$ inside the first copy of G_d . Then take the edge to $(1, 1, 0, \dots, 0)$, and traverse the Hamilton path in the second copy of G_d backwards all the way to $(1, 0, 0, \dots, 0)$. Finally, take the edge back to $(0, 0, 0, \dots, 0)$ to close the cycle. \square

Note that this in particular gives an ordering of the first 2^n bit strings such that going from one string to the next requires only one bit-operation. The ordering we constructed in the proof of Proposition 71 is also known as *Gray code*, and has a wide variety of applications.



The Hamiltonian cycle constructed in the proof of Proposition 71 in a 3-cube on the left, and a cycle omitting exactly one vertex in the Petersen graph on the right.

The Petersen graph. The Petersen graph has as vertex set the set of two-element subsets of $\{1, 2, 3, 4, 5\}$, where two subsets are adjacent iff they are disjoint. This gives a graph on 10 vertices, where every vertex has degree 3, and hence there are 15 edges in total.

Proposition 72. *The Petersen graph is not Hamiltonian. However, upon removing any one of its vertices, the remaining graph is, and the Peterson graph admits a Hamilton path.*

Proof. Consider the five edges connecting the inner five vertices with the outer five vertices (in the figure above). Any Hamilton cycle C needs to traverse along either 2 or 4 of those edges. If it contains two of them, they need to connect pairs of adjacent vertices on both the inside and the outside, and this is impossible. So, C contains four of those edges. Assume the edge at $\{3, 5\}$ is not included. This implies that among the 5 outside edges, the two on top and the bottom one are in C , and the other two are not. However, $\{1, 2\}$ needs to lie on C , and this forces all edges in the cycle on vertices $\{1, 2\}, \{4, 5\}, \{1, 3\}, \{2, 5\}, \{3, 4\}$ to be on C , which is absurd.

To see that there exists a cycle through all but one vertex, we first show that the choice of the omitted vertex does not matter. Let $\{i, j\}, \{i', j'\} \subseteq \{1, \dots, 5\}$. Then, there exists a permutation π on $\{1, \dots, 5\}$ mapping i to i' and j to j' . This permutation extends to two-element subsets by mapping $\{k, l\}$ to $\{\pi(k), \pi(l)\}$, and this is in fact a bijection on $\mathcal{P}_2(\{1, \dots, 5\})$ that preserves disjointness. In graph-theoretic language, this shows that for any two vertices $\{i, j\}, \{i', j'\}$ of the Petersen graph, there exists an isomorphism of the Petersen graph with itself that maps $\{i, j\}$ to $\{i', j'\}$ (i.e. the graph is vertex-transitive). Hence the choice of the vertex is irrelevant. Finally, the picture above shows that there is a cycle traversing all vertices but $\{1, 2\}$, and including the edge $\{\{1, 2\}, \{3, 5\}\}$ instead of any of the other two edges incident to $\{3, 5\}$ will result in a Hamilton path. \square

Lecture 8 The max-flow-min-cut theorem

Consider a problem of the following form: We have to transport some material from point A to point B, using e.g. the railway network. We can think of this network as a graph, where e.g. train stations are vertices, and edges are given by the railways in between. However, every railway section only has a certain capacity (limited e.g. by the number of trains, or by their weights, or ...). Given these constraints, how can we transport as much material as possible?

The answer to such a question comes from the study of flows in directed graphs, and will be the topic of this lecture.


Definition 73. A *directed simple graph* is a pair $G = (V, E)$, where $E \subseteq V \times V \setminus \{(v, v) | v \in V\}$. Here, we interpret an edge (v, w) as being an edge from a vertex v to a (different) vertex w . Note that in this way, $(v, w) \neq (w, v)$.

We could also consider directed multigraphs, but for our purposes, directed simple graphs will be completely sufficient.

Definition 74. A *flow network* consists of a directed graph $G = (V, E)$ together with an edge weight function $w : E \rightarrow (0, \infty)$, where two vertices are considered distinguished: a source node $s \in V$ and a sink node $t \in V$.

The function $c : V \times V \rightarrow [0, \infty)$ defined by $c(v, v') = w(v, v')$ if $(v, v') \in E$ and $c(v, v') = 0$ otherwise is called the *capacity function* of the flow network.

For simplification, we shall assume that if (v, v') is an edge in a flow network, then (v', v) is not an edge. This is not a real restriction, however: If both directed edges are present in a graph, we can split the edge (v, v') by introducing a new vertex z in the middle, and set $w(v, z) = w(z, v') := w(v, v')$.



Intuitively, we can think of the flow network being a network of pipes, through which water flows in the direction of the edges from the source to the sink, and where the capacity of an edge tells us how much water per unit of time can pass through this pipe. A *flow* will then be given by the information how much water actually passes through every pipe, and it should satisfy two constraints: First, it should not exceed the capacity; and second, at every vertex v different from source and sink, the total amount of incoming water should equal the total amount of outgoing water. The following definition formalises this:

Definition 75. A flow f on the flow network G with capacity function c is a function $f : V \times V \rightarrow [0, \infty)$ such that the following conditions are satisfied:

- (i) $f(v, v') \leq c(v, v')$ for all $v, v' \in V$ (capacity constraint)
- (ii) For $v \in V \setminus \{s, t\}$, we have $\sum_{x \in V} f(x, v) = \sum_{x \in V} f(v, x)$ (conservation constraint)

The *value* of a flow, denoted by $|f|$, is the total out-flow at the source. In other words, $|f| = \sum_{x \in V} f(s, x) - \sum_{x \in V} f(x, s)$.

It follows from the conservation constraint that $|f|$ is also equal to the total in-flow at the sink, i.e. $|f| = \sum_{x \in V} f(x, t) - \sum_{x \in V} f(t, x)$. Notice that therefore, we can always assume that $|f| \geq 0$ (if it is not, simply swap the roles of s and t).

Imagine now we would cut edges in our network in such a way that we end up with two connected components, one containing the source and the other containing the sink. Any flow in the network must have passed through the removed edges, therefore the sum of the capacities of these edges must be at least as big as the value of the flow.

Definition 76. Let $G = (V, E)$ be a flow network with source s and sink t . An *s-t-cut* is a partition of V into two sets S, T such that $s \in S$ and $t \in T$. The *capacity of the cut* is $c(S, T) = \sum_{(v, v') \in S \times T} c(v, v')$, that is, the sum of the capacities of the cut edges.

The intuitive argument prior to the definition shows $|f| \leq c(S, T)$ for any flow f and any s-t-cut $V = S \uplus T$. The following lemma shows that equality can only happen in the most extreme case:

Lemma 77. *Let G be a flow network with a flow f and an s-t-cut of V into S, T . Assume $|f| = c(S, T)$. Then $|f|$ is maximal among all flows on G , and $c(S, T)$ is minimal among all s-t-cuts of G .*

Proof. Any flow f' with value $|f'| \geq |f|$ would need to pass through the s-t-cut (S, T) as well, hence $|f'| = |f|$ and f is maximal. Any s-t-cut (S', T') with capacity $c(S', T') \leq c(S, T)$ would need to satisfy $|f| \leq c(S', T')$, hence $c(S', T') = c(S, T)$ and (S, T) is minimal. \square

The following construction is central in the theory of flow networks:

Definition 78. Let $G = (V, E)$ be a flow network with source s and sink t , and let f be a flow on G . The *residual network* G_f is the flow network with


residual capacity c_f on the vertex set V constructed in the following way: For vertices u, v , set

$$c_f(u, v) := \begin{cases} c(u, v) - f(u, v) & \text{if } (u, v) \text{ is an edge in } G \\ f(v, u) & \text{if } (v, u) \text{ is an edge in } G \\ 0 & \text{otherwise} \end{cases}$$

Now, let $E(G_f)$ be the set of pairs (u, v) for which $c_f(u, v) > 0$.

Observe how for this definition to make sense, it is essential to have no parallel edges in opposite direction!

Example 79. Consider a flow network G with the flow f as below, where the edge labels are of the form “flow”/“capacity”. Its residual network G_f is shown on the right, where the dashed arrows are the contributions that come from the flow f , whereas the drawn arrows come from the unused capacity of the edge.



Definition 80. Let G_f be the residual network with respect to a flow f in a flow network G . A v_0 - v_k -path is a sequence $v_0e_1v_1e_2v_2 \dots v_{k-1}e_kv_k$ where $v_0, v_1, \dots, v_{k-1}, v_k$ are vertices, and e_1, \dots, e_k are edges in G_f , such that these edges all point away from v_0 , i.e. $e_i = (v_{i-1}, v_i)$. An s-t-path P is called *augmenting*.

Observe that in the construction of G_f , we only draw the arrows with positive capacity. Hence the residual capacity of any v_0 - v_k -path P , $c_f(P) := \min_{e \in E(P)} c_f(e)$ is positive. In particular, this is true for augmenting paths.

Lemma 81. Let f be a flow in a network $G = (V, E)$ such that G_f admits an augmenting path P . Then the flow f' , defined by


$$f'(u, v) = \begin{cases} f(u, v) + c_f(P) & \text{if } (u, v) \text{ is an edge in } P \\ f(u, v) - c_f(P) & \text{if } (v, u) \text{ is an edge in } P \\ f(u, v) & \text{otherwise} \end{cases}$$

for all edges $(u, v) \in E$, satisfies $|f'| > |f|$.

Sketch of proof. We first check that $0 \leq f'(u, v) \leq c(u, v)$ for all edges $(u, v) \in E$, which involves a case distinction between the first two cases in the definition of f' . If (u, v) is an edge both in P and in E , then $c_f(u, v)$ is the unused capacity of (u, v) , and $c_f(P) \leq c_f(u, v)$. On the other hand, if (v, u) is an edge in P , then $c_f(P) \leq c_f(v, u) = f(u, v)$ by construction of G_f . Moreover, f' satisfy the conservation constraint, since at every interior vertex of P , exactly two edges get changed by $c_f(P)$ (and it is a somewhat tedious task to verify that this is compatible with the direction of the arrows).

Finally, consider the source node s . If f' changes the flow through an outgoing edge in G , then it will increase the flow there by $c_f(P)$, otherwise, it will diminish the flow through an incoming edge G . In both cases, $|f'| = |f| + c_f(P) > |f|$. \square

Example 82. In Example 79, we could have found an augmenting path in G_f by going down from s , then taking the diagonal arrow, and then going down to t . This path would have let to the following new flow f' on G (shown left), and the corresponding residual network $G_{f'}$:



Theorem 83 (Ford-Fulkerson). *Let f be a flow on the network G . Then, t.f.a.e.:*

- (i) f is a maximal flow.
- (ii) G_f contains no augmenting path.
- (iii) There is an s - t -cut (S, T) with $|f| = c(S, T)$.

In particular, the maximal value of a flow equals the minimal capacity of an s - t -cut.

Proof. The implication (iii) \implies (i) is contained in Lemma 77. The implication (i) \implies (ii) is merely the contrapositive formulation of Lemma 81. Hence it only remains to show (ii) \implies (iii).

Assume that G_f does not contain an augmenting path. Define S to be the set of vertices that are either s or can be reached from s by a directed path

in G_f from s . Set $T := V \setminus S$. By assumption $t \in T$, and we have an s-t-cut (S, T) . By construction, every arrow in G from S to T must have its full capacity used by f (otherwise there would still be an arrow in G_f from S to T), and every arrow in G from T to S must have a flow of 0 (otherwise, G_f would contain an arrow in the opposite direction, so from S to T). Therefore, $|f| \geq c(S, T)$, which is only possible if $|f| = c(S, T)$.

The final claim follows via the implication (i) \implies (iii) and Lemma 77. \square

The Ford-Fulkerson theorem makes no statement about the existence of a maximal flow. However, in a finite network $G = (V, E)$ we can regard any flow f as a vector in $\mathbb{R}^{|E|}$, where the e -th entry is from the closed interval $[0, c(e)]$. This implies, together with the conservation constraint, that the set of all flows on G is a compact subset of $\mathbb{R}^{|E|}$, and the map $f \mapsto |f|$ is continuous, hence a maximal flow exists.

Note also that the construction in the proof of the Ford-Fulkerson theorem can be used to obtain a minimum s-t-cut from a given maximum flow f .

Theorem 84 (Integer flow theorem). *If G is a network with an integer capacity function, i.e. $c : V \times V \rightarrow \mathbb{Z}_{\geq 0}$, then there is a maximal integer flow, i.e. a maximal flow such that $f(e) \in \mathbb{Z}$ for every edge e .*

Proof. Lemma 81 also works for integer flows and capacities, then yielding an integral augmenting path. Thus, a non-maximal integer flow can be augmented to another integer flow, and $|f|$ is bounded, hence there must be a maximal integer flow. \square

The theory of Ford-Fulkerson yields the following method of finding maximal flows:

Ford-Fulkerson-“algorithm” Given a network G , start with the 0-flow f . Repeat the following steps:

1. Construct G_f .
2. Find an augmenting path in G_f , then use it to modify f into a flow with larger value, like in Lemma 81. If no augmenting path exists, stop.

This method is guaranteed to yield a maximal flow by Theorem 83, which can, if necessary, be converted into a minimal cut, by performing the construction in the proof of the Ford-Fulkerson theorem. However, the algorithm might not terminate, although it does if the capacities in G are all rational (Why?).

Lecture 9 Matchings

Matchings arise when trying to pair up the vertices in a graph: If it were just for the vertices, this wouldn't pose much of a problem, but we want to do it in such a way that paired vertices are neighbouring. This leads to the following definition:

Definition 85. Let $G = (V, E)$ be a finite simple graph. A *matching* on G is a set $M \subseteq E$, such that no two edges in M have a common endpoint. If $e = \{v, w\}$ is an edge in M , then we also say that the vertices v and w have been matched to each other.

For now, let us confine ourselves with considering *bipartite* graphs, i.e. graphs where we have a partition $V = A \uplus B$ such that there are no edges between two vertices in the same set of the partition. The complete bipartite graphs $K_{a,b}$ are just the edge-maximal bipartite graphs for fixed A, B .

In a bipartite graph on a partitioned vertex set $V = A \uplus B$, we can ask if there is a matching that matches all the vertices in A to some vertex in B . For this to work, any subset $Q \subseteq A$ must have at least $|Q|$ neighbours in B (otherwise, you couldn't even match Q into B !). So, let's denote by $N(Q)$ the set of all vertices that are a neighbour to some $v \in Q$. The next famous theorem states that our obvious necessary condition is sufficient:

Theorem 86 (Hall, "marriage theorem"). *Let $G = (V, E)$ be a finite simple bipartite graph with $V = A \uplus B$. Then G contains a matching of A into B iff $|N(Q)| \geq |Q|$ for all $Q \subseteq A$.*

Proof. It remains to show the sufficiency of the *marriage condition*: Assume that $|N(Q)| \geq |Q|$ for all $Q \subseteq A$. Create a flow network as follows: Introduce a new vertex s which is a neighbour to every vertex in A , and another new vertex t that is a neighbour to every vertex in B . Direct the edges in such a way that everything points away from s and towards t , for edges between A and B orient them so that they go from A to B . Finally, give every edge incident to either s or t capacity 1, and give capacity ∞ to the edges going from A to B (This is for convenience, and can be replaced by a large enough constant depending on G).


For $S' = \{s\}$ and $T' = A \cup B \cup \{t\}$ we have $c(S', T') = |A| < \infty$, hence $c(S, T) < \infty$ for any minimum cut (S, T) . In particular, there cannot be a directed edge (u, v) with $u \in A \cap S$ and $v \in B \cap T$. Hence $N(S \cap A) \subseteq S \cap B$.

With this, let's compute $c(S, T)$:

$$\begin{aligned} c(S, T) &= \sum_{\substack{u \in S \\ v \in T}} c(u, v) = \sum_{v \in T \cap A} c(s, v) + \sum_{u \in S \cap B} c(u, t) \\ &= |T \cap A| + |S \cap B| \geq (|A| - |S \cap A|) + |N(S \cap A)| \\ &\geq |A| - |S \cap A| + |S \cap A| = |A|. \end{aligned}$$

Comparing this with the cut (S', T') shows that a minimum cut has capacity $|A|$, hence there is a maximal integral flow f with value $|A|$. Following along this flow then yields a matching of A into B . \square

To exemplify the construction of the flow network, here is what the network for a $K_{3,2}$ would look like. The unlabelled arrows have infinite capacity.



Corollary 87. *Let $G = (V, E)$ be a finite simple bipartite graph with $V = A \uplus B$. If $|N(Q)| \geq |Q| - d$ for every subset $Q \subseteq A$ and some $d \in \mathbb{N}$, then G contains a matching with $|A| - d$ edges.*

Proof. Add d new vertices to B , each of them neighbouring all vertices in A . This modified graph satisfies the marriage condition and therefore contains a matching of A into the modified B , and at least $|A| - d$ edges of this matching must be in the original graph. Those form the desired matching. \square

Remark 88. The construction of a flow network from a bipartite graph $G = (A \uplus B, E)$ as used in the proof of Hall's theorem gives a bijection between integer flows on the network and matchings on G . Indeed, given an integer flow f from s to t , taking all edges (a, b) with $a \in A$, $b \in B$ and $f(a, b) = 1$ gives a matching on G because the flow through any vertex in G can be at most 1. Moreover, the cardinality of this matching equals the value of the flow, as can be seen by looking at the amount of flow going from A to B .

Hall's theorem answers the question whether we can match every vertex in A to some vertex in B . But if the answer is no, we can still ask for the maximum number of vertices in A that can be matched. König's theorem provides some insight there, but we need another definition first:

Definition 89. Let $G = (V, E)$ be a finite simple graph. A *vertex cover* of G is a subset $S \subseteq V$ such that every edge has an endpoint in S . The *covering number* of G , denoted $\beta(G)$ is the minimum cardinality of any vertex cover of G .

As we saw in Remark 88, constructing the flow network to a bipartite graph relates matchings and flows. A dual relationship exists between vertex covers and s-t-cuts with finite capacity: If (S, T) is a finite capacity cut on the network constructed from $G = (A \uplus B, E)$, then $C := (A \cap T) \cup (B \cap S)$ is a vertex cover of G with $|C| = c(S, T)$ (indeed, any edge (a, b) with $a \in A, b \in B$ that is not covered by C would extend to a path from s to t in the flow network, which is impossible since (S, T) was a cut). Similarly, any vertex cover C of G gives rise to an s-t-cut in the network: Set $S = \{s\} \cup (A \setminus C) \cup (B \cap C)$ and $T = \{t\} \cup (A \cap C) \cup (B \setminus C)$. Again, $c(S, T) = |C|$. Therefore, there is also a bijection between finite capacity cuts (S, T) and vertex covers C on G , with $c(S, T) = |C|$.

Theorem 90 (König's theorem). *Let G be a finite simple bipartite graph. Then the maximum cardinality of a matching on G equals the minimum cardinality of a vertex cover of G .*

Proof. We use the afore-mentioned bijections between matchings and flows, and between cuts and vertex covers. Let M be a matching in G with maximal cardinality $|M|$. Construct the flow network G' to the bipartite graph as in the proof of Hall's theorem. Then, by Remark 88, a maximal flow f on G' has value $|f| = |M|$. By the max-flow-min-cut theorem (Theorem 83), a minimum cut (S, T) on G' must be of capacity $c(S, T) = |M|$. Finally, the bijection constructed after Definition 89 yields $\beta(G) = c(S, T) = |M|$, as required. \square

Observe that even for general graphs, one inequality is true: Every edge in a matching has to be covered by a vertex cover, and no vertex can cover more than one edge of the matching at a time. Hence, for a maximum matching M , the inequality $|M| \leq \beta(G)$ holds for an arbitrary graph G .

Let us now shift our attention towards matchings in arbitrary graphs. In particular, let us focus on the question of when we can match up all vertices in a graph:

Definition 91. A *perfect matching* or *1-factor* M on a finite simple graph $G = (V, E)$ is a matching on G such that every vertex $v \in V$ is an endpoint of an edge $e \in M$.

The name 1-factor comes from the following definition:

Definition 92. Let $k \in \mathbb{N}$. A k -regular graph is a graph where every vertex has degree k . A k -factor of a finite simple graph G is a k -regular spanning subgraph of G .

For example, the 1-regular graphs are precisely those where every connected component is a copy of K_2 and the 2-regular graphs are precisely those where every connected component is a cycle. Moreover, every hamiltonian graph admits a 2-factor (the Hamilton cycle being a 2-regular spanning subgraph).

Clearly, for a finite simple graph $G = (V, E)$ to have a perfect matching, $|V|$ must be even. However, more can be said: We shall write $G - S := G[V \setminus S]$ for the graph obtained after removing a vertex subset $S \subseteq V$. If G admits a perfect matching, then $G - S$ cannot contain more than $|S|$ components on an odd number of vertices (since those components must have had an edge from the matching connecting them to a vertex in S). For brevity, we call those components *odd* and denote the number of odd components in a graph G by $o(G)$.

Theorem 93 (Tutte's 1-factor theorem). *A finite simple graph $G = (V, E)$ admits a perfect matching iff $o(G - S) \leq |S|$ for any $S \subseteq V$.*

Proof. We only have the necessity of the condition $o(G - S) \leq |S|$ left to show. So, let $G = (V, E)$ be a graph satisfying this condition but without perfect matching. Adding an edge to such a graph will not affect the condition: If the new edge connects two even or one even and one odd component in $G - S$, then the $o(G - S)$ doesn't change, and if it connects two odd components, then $o(G - S)$ decreases. Hence, we can assume that G is an edge-maximal graph satisfying $o(G - S) \leq |S|$ for all $S \subseteq V$, but without containing a perfect matching.

Denote $n = |V|$. Observe how considering $S = \emptyset$ implies $o(G) \leq 0$, hence n must be even. Now, choose $S = \{v \in V : \deg(v) = n - 1\}$ (i.e. the set of vertices being a neighbour to all other vertices), and consider $G - S$. We distinguish 2 cases:

1) All components of $G - S$ are complete. Then all of the even components contain a perfect matching, and in each of the odd components, all but one vertex can be matched up. By choice of S , this remaining vertex must have been a neighbour to all vertices in S . Moreover, since there are at most $|S|$ odd components in $G - S$, we can match all of these remaining vertices with vertices from S . This might leave some (an even number, since n is even!) vertices in S unmatched, but by construction, $G[S]$ is complete, hence they can be matched to each other. In this way, we constructed a perfect matching on G .

2) Let now K be a component of $G - S$ that is not complete. In particular, there exist vertices $x, a, b \in V(K)$ such that $\{x, a\}$ and $\{a, b\}$ is an edge, but

$\{x, b\}$ is not. Moreover, $a \notin S$, so there is a vertex $c \in V(G)$ that is not a neighbour to a . By assumption, G was edge-maximal without a perfect matching. Hence, there exist perfect matchings

$$\begin{aligned} M_1 & \text{ for } (V, E \cup \{\{x, b\}\}) \\ M_2 & \text{ for } (V, E \cup \{\{a, c\}\}). \end{aligned}$$

Surely, $\{x, b\} \in M_1$ and $\{a, c\} \in M_2$.

In G , consider a maximal path P starting at c with an edge from M_1 , and then alternating between edges in M_2 and edges in M_1 . This is always possible: Since M_2 matches c with a , but M_1 doesn't have this edge available, M_2 has to match c to some other vertex $d \neq a$. Then M_1 cannot match d to c , as c is already matched, so M_1 matches d to a new vertex $e \neq c$, and so on. This procedure continues until we reach x, a , or b .

Let v denote the last vertex on P . If the last edge on P was from M_1 , then $v = a$ – otherwise there is another vertex M_2 -matched to v that is not on P yet. Then, define C to be the cycle in $(V, E \cup \{\{a, c\}\})$ consisting of P and the edge $\{a, c\}$. Similarly, if the last edge in P is from M_2 , then $v \in \{x, b\}$ (since P is a path in G , the edge $\{x, b\} \in M_1$ is not available to extend P further!). In this case, define C to be the cycle in $(V, E \cup \{\{a, c\}\})$ consisting of P and the edges $\{v, a\}, \{a, c\}$. In both cases, C is a cycle of even length, where every second edge belongs to M_2 . Now modify M_2 by replacing its edges on C by the edges not in C – this will again be every second edge on C , but it will avoid the edge $\{a, c\}$. Together with the edges from M_2 that were not on C , we thus obtain a perfect matching in G using only edges from E , a contradiction. \square

Lecture 10 Connectivity

We have seen what it means for a graph to be connected already as early as lecture 1. The notion of connectivity greatly generalises this, and leads to statements about structural properties of graphs.

Definition 94. Let $k \in \mathbb{Z}_{\geq 0}$. A finite simple graph $G = (V, E)$ is called k -connected if $|V| > k$ and $G - X$ is connected for all sets $X \subseteq V$ with $|X| < k$. The largest k for which G is k -connected is the *connectivity* $\kappa(G)$ of G .

With this definition, every non-empty graph is 0-connected, and every connected graph on at least 2 vertices is 1-connected. Also observe that $\kappa(K_n) = n - 1$ (as no vertex subset can disconnect a complete graph upon its removal), and that $\kappa(G) = 0$ means that either $G = \bullet$ or that G is disconnected.

Definition 95. Let $G = (V, E)$ be a finite simple graph, and let $v, w \in V$ and $A, B \subseteq V$. A set $X \subseteq V$ *separates* v from w if $v, w \notin X$ and if every path from v to w contains at least one vertex of X . We denote the minimum size of a set separating v from w by $\kappa(v, w)$. A set $X \subseteq V$ *separates* A from B if every path from A to B contains at least one vertex of X . Two paths from v to w are called *(vertex)-independent* if they do not share a common vertex besides the two endpoints.

Careful: The definition above makes a subtle but important difference between separating vertices v, w (in which case the separating set X may *not* contain either of the two vertices), and separating sets of vertices A, B (in which case X may intersect A and B). In fact, to separate A from B , one is required to have $A \cap B \subseteq X$, otherwise the lazy path at any vertex in $A \cap B$ would prevent separation.

Proposition 96. For a finite simple graph $G = (V, E)$ that is not a complete graph, we have $\kappa(G) = \min_{x,y} \kappa(x, y)$, where the minimum is taken over all non-adjacent pairs of vertices x, y .


Proof. Since G is not complete, $\kappa(G)$ equals the minimum size of a set $X \subseteq V$ whose removal disconnects the graph. In particular, $G - X$ contains at least two components, so we can choose x_0, y_0 as vertices from different components. Hence x_0, y_0 can be separated by X , and $\min_{x,y} \kappa(x, y) \leq |X| = \kappa(G)$.

Conversely, let x_0, y_0 be two non-adjacent vertices such that $\kappa(x_0, y_0)$ attains its minimum, then there exists a separating set X for x_0 and y_0 with $|X| = \kappa(x_0, y_0) = \min_{x,y} \kappa(x, y)$. In particular, a set of size $|X|$ suffices to disconnect G , hence $\kappa(G) \leq \min_{x,y} \kappa(x, y)$. \square

Recall that we originally defined a graph to be connected if for any two vertices v, w there is a path from v to w . Is there a similar characterisation for k -connectedness? It turns out, the answer is yes:

Theorem 97 (Menger). *Let $G = (V, E)$ be a finite simple graph, and let $v, w \in V$ be non-adjacent. Then the minimum size of a set that separates v from w equals the maximum number of independent paths from v to w .*

Proof. Construct a flow network from G as follows: Replace every edge $\{x, y\} \in E(G)$ by two directed edges (x, y) and (y, x) with infinite capacity each. Then split every vertex $x \in V(G) \setminus \{v, w\}$ into two vertices x_0, x_1 such that every incoming edge to x is incoming to x_0 and every outgoing edge from x is outgoing from x_1 . Also add an edge (x_0, x_1) of capacity 1. Finally, declare v to be the source and w to be the sink.



Analogously to the proofs of Hall's and König's theorem, integer flows through the network correspond to unions of disjoint paths from v to w in the graph (where the value of the flow coincides with the number of disjoint paths). Finite capacity v - w -cuts in the network need to cut through edges of capacity 1 that came from splitting vertices, hence finite capacity cuts correspond to sets separating v from w (where the capacity of the cut coincides with the size of the separating set). Hence, the statement follows from the max-flow-min-cut theorem, Theorem 83. \square

Corollary 98 (Menger, global version). *A finite simple graph $G = (V, E)$ is k -connected if and only if it contains k independent paths between any two vertices.*

Proof. If G is k -connected, then $\kappa(G) \geq k$, hence $\kappa(x, y) \geq k$ for all non-adjacent pairs of vertices by Proposition 96. By Menger's theorem, those vertices are therefore connected by at least k independent paths, and we need to show that there are also enough paths when x and y are neighbours. Assume that there are at most independent $k - 1$ paths from x to y . After removing the edge $\{x, y\}$ from G , resulting in a graph G' , there are at most $k - 2$ paths left. Hence, according to Menger, we can separate x from y by a set X of size at most $k - 2$ in G' . At the same time, G has more than k vertices, so there is a vertex v not in X , and different from x and y . Then v is separated by X from either x or y (w.l.o.g. from x). Thus $X \cup \{y\}$ separates v from x in G , but $|X \cup \{y\}| \leq k - 1$, contradicting the assumption that G is k -connected.

Conversely, if there are k independent paths between any two vertices, then this in particular holds for any pair of non-adjacent vertices. If there is no such pair, then G is a complete graph on at least $k + 1$ vertices, hence $\kappa(G) \geq k$. Otherwise we have $\kappa(x, y) \geq k$ by Menger's theorem for any pair of non-adjacent x, y , and hence $\kappa(G) \geq k$ by Proposition 96. \square

Sometimes, the following version is more useful:

Theorem 99 (Menger, v2). *Let $G = (V, E)$ be a finite simple graph and $A, B \subseteq V$. Then the minimum size of a set X that separates A from B is equal to the maximum number of disjoint paths with one end in A and the other end in B .*

Proof. To G , add two vertices s, t with arrows pointing from s to every vertex in A and from every vertex in B to t . Endow these arrows with infinite capacity, then proceed to split all of the original vertices as in the proof of Theorem 97. The remaining argument is the same. \square

We have seen already how 1-connected graphs look like (they're connected, duh!). What's special about 2-connected graphs?

Proposition 100. *A finite simple graph is 2-connected iff it can be constructed from a cycle graph by successively adding paths to it, both of whose endpoints lie in the graph already constructed.*

Proof. Every graph constructed in this fashion is 2-connected. Let's call a graph constructible if it can be obtained from the described procedure. If $G = (V, E)$ is a finite simple 2-connected graph, then G contains a cycle, which is constructible. Hence, there is a maximal constructible subgraph H in G . Then, H is induced: For $x, y \in V(H)$ and an edge $\{x, y\}$ in G but not in H , this edge would constitute a path with both its endpoints in H , contradicting the maximality of H . If H is not spanning, there is an edge $\{x, y\}$ with $x \in V(H)$ and $y \notin V(H)$. Since G is 2-connected, removing x from G still leaves a path P connecting y to $V(H)$. But that means that in G , going from x to y and then back along P to $V(H)$ is a path with both endpoints in $V(H)$, again contradicting the maximality of H . Hence H must be both induced and spanning, thus $H = G$. \square

Definition 101. Let $G = (V, E)$ be a finite simple graph. A *cutvertex* is a vertex whose removal increases the number of connected components of G . A *bridge* is an edge whose removal increases the number of connected components of G . A *block* is a maximal (with respect to inclusion) connected subgraph H without a cutvertex in H .

The idea here is that blocks are for 2-connectedness what connected components are for connectedness. Notice that any isolated vertex and any bridge (with its endpoints) form a block, as these subgraphs are simply too small to have cutvertices. Any larger block is guaranteed to be 2-connected, by the global version of Menger's theorem. Observe further that two blocks intersect in at most one vertex, which is then a cutvertex in G (otherwise, the union of the two blocks would be a larger connected subgraph without cutvertex). In particular, any edge of G belongs to exactly one block.


Let G be a finite simple graph, let A be its collection of cutvertices, and let \mathcal{B} be its collection of blocks. The *block graph* of G is a bipartite graph with vertices $A \uplus \mathcal{B}$, where an edge is drawn from $a \in A$ to $B \in \mathcal{B}$ if a is a vertex in B .

Lemma 102. *Let $G = (V, E)$ be a finite simple connected graph. Then the block graph of G is a tree.*

Proof. If the block graph of G contains a cycle, then by its bipartiteness it must traverse at least two blocks of G . This means that G must have contained a cycle traversing inside two different blocks. However, any cycle C of G , is a connected subgraph without cutvertex, and hence contained in a single block B . \square

Let's now think about 3-connected graphs (because we'll need them next time).

Definition 103. Let $G = (V, E)$ be a finite simple graph and $e \in E$. The *contraction* of e is the graph G/e , constructed in the following way: For $e = \{x, y\}$, introduce a new vertex v_{xy} and set $V(G/e) = V \setminus \{x, y\} \uplus \{v_{xy}\}$. Draw an edge in G/e between v_{xy} and w whenever w was a neighbour to x or y . Keep all edges of G that were not incident to x or y .



Lemma 104. *If the finite simple graph $G = (V, E)$ is 3-connected and has more than 4 vertices, then there exists an edge $e \in E$ such that G/e is again 3-connected.*

Proof. Suppose no such edge exists. In particular, for all edges $e = \{x, y\}$, G/e contains a separating set X on 2 or less vertices. Since G is 3-connected,

this is only possible if v_{xy} belongs to X and if X contains another element, say z (In particular, $z \neq x, y$). Then, $\{x, y, z\}$ must be a minimal separating set for G , hence each of the three vertices must have a neighbour in every component of $G - \{x, y, z\}$ (Why?). Denote by C the smallest such component, and suppose we've chosen $\{x, y\}$ and z in such a way, that $|V(C)|$ is minimal. Choose a neighbour v of z in C . Then again, $G/\{v, z\}$ is not 3-connected, so just as earlier for $\{x, y, z\}$, there is a vertex w s.t. $\{z, v, w\}$ separated G . And once more, z, v and w each have a neighbour in every component of $G - \{z, v, w\}$. One component, call it D , of $G - \{z, v, w\}$ must contain neither x nor y . Since $v \in C$, every neighbour of v that is in D must also be in C . Thus $D \cap C \neq \emptyset$, hence D is a strict subset of C . However, this contradicts our choice of $\{x, y\}$ and z and our construction of C . \square

We quote without proof the following theorem (again by Tutte):

Theorem 105 (Tutte 1961). *A finite simple graph G is 3-connected iff there is a sequence G_0, \dots, G_n of finite simple graphs, where $G_0 = K_4$, $G_n = G$, and for every $i = 0, \dots, n - 1$ the graph G_{i+1} has an edge $\{x, y\}$ with $\deg_{G_{i+1}}(x), \deg_{G_{i+1}}(y) \geq 3$ and $G_i = G_{i+1}/\{x, y\}$.*


Lecture 11 Planarity

For small graphs, we often prefer a picture of the graph over specifying the vertex and edge sets. This raises the question how “clean” of a picture we can draw for a given graph – for example, can we draw the graph without intersecting edges?

More specifically, we draw a graph by placing the vertices at certain points in \mathbb{R}^2 , and draw an edge by connecting its endpoints with an *arc*, that is, a continuous/piecewise smooth/piecewise linear curve from one endpoint to the other.

Definition 106. Let $G = (V, E)$ be a finite simple graph. A graph is *planar* if it can be drawn (*embedded*) in \mathbb{R}^2 such that the vertices are placed at different points, and such that no edges intersect.

For example, K_4 is planar as there is the following embedding to the plane:



The following definition relies on the very intuitive (but difficult to prove) *Jordan curve theorem*, stating that a closed continuous simple curve divides the plane into exactly two open regions (inside and outside of the curve), one of which is bounded, and one of which is unbounded. For the purpose of this lecture, we generously ignore the topological ramifications, and rely on our intuition.

Definition 107. Let $G = (V, E)$ be a planar graph. Any planar embedding of G divides the plane into several connected components, called *faces*, all but one of which are bounded.


Observe that an embedded tree only yields one face (the unbounded one), and that the boundary of any face contains a cycle in G if G contains a cycle at all.

Theorem 108. Let $G = (V, E)$ be a connected planar graph. Denote by f the number of faces for some planar embedding of G . Then $|V| - |E| + f = 2$. In particular, any two planar embeddings have the same number of faces.

Our proof uses the following concept, which is important enough to warrant its own definition:

Definition 109. Let $G = (V, E)$ be a planar graph and fix a planar embedding of it. The *planar dual* G^* is the multigraph whose vertices are the faces of the embedding of G and with edge set E . An edge $e \in E$ connects vertices f_1, f_2 in G^* iff e was part of the boundary of the faces f_1, f_2 .

Example 110. Here's an embedded graph together with its planar dual:



Note that the planar dual really depends on the embedding, i.e. there are graphs with different planar embeddings, that produce non-isomorphic planar duals!

Proof of Theorem 108. Let $G = (V, E)$ be a connected planar graph. Fix a certain planar embedding of it, and construct its planar dual G^* . Also fix a spanning tree $T = (V, E_T)$ of G , and consider the spanning subgraph $T^* = (V(G^*), E \setminus E_T)$ of G^* .

If T^* was disconnected, the embedding of T (i.e. the embedding of G , but restricted to the subgraph T) would contain a set of edges that completely encloses a face of G , hence T would contain a cycle. This is impossible due to T being a tree. If T^* contained a cycle, the edges of this cycle would not be in T , therefore separating at least one vertex in T , again contradicting T being a tree.

Now T is a tree on $|V|$ vertices and thus has $|V| - 1$ edges, and T^* is a tree on f vertices and thus has $f - 1$ edges. Since every edge in E is either in T or in T^* , we get $|V| - 1 + f - 1 = |E|$, which proves the claim. \square

Euler's formula has far-reaching consequences, but most importantly, we can use it to show that certain graphs are not planar:

Corollary 111. *If $G = (V, E)$ is a planar graph on at least 3 vertices, then $|E| \leq 3|V| - 6$. If $G = (V, E)$ is a planar graph on at least 3 vertices that does not contain a triangle, then $|E| \leq 2|V| - 4$.*

Proof. Place coins to both sides of each edge in a planar embedding of G . Then you placed $2|E|$ coins, and at the same time, you placed at least $3f$ coins, as every face contains at least 3 coins from the minimum 3 edges necessary to surround it. Hence $3f \leq 2|E|$. Together with Euler's formula $|V| - |E| + f = 2$, this yields $|E| \leq 3|V| - 6$ for connected planar graphs G , which can be extended to disconnected planar graphs by considering each component separately.

The second claim follows analogously, only that here, every face is bounded by at least 4 edges. \square


Corollary 112. *The graphs K_5 and $K_{3,3}$ are non-planar.*

Proof. K_5 contains 5 vertices and 10 edges, so cannot be planar by Corollary 111. $K_{3,3}$ contains 6 vertices and 9 edges, but because it is bipartite it also cannot contain a triangle. Therefore it is non-planar, again by Corollary 111. \square

As it turns out, these are, in a sense, the only two minimal non-planar graphs: Every graph that contains one of them is non-planar as well. But in this context, we can broaden our definition of containment a bit:

Definition 113. Let G, H be finite simple graphs. A *subdivision* of H is a graph where some of the edges of H were replaced by vertex-independent paths. If a graph G contains a subdivision of H , then H is a *topological minor* of G .

The following pictures show a graph H on four vertices together with a subdivision of H on the right. If a graph G contains such a subdivision, it is said to have H as a topological minor.



Definition 114. Let G, H be finite simple graphs. The graph H is a *minor* of G , if there is a subgraph G' of G such that H can be obtained from G' by a finite sequence of edge-contractions (cf. Definition 103).

Clearly, if H is a (topological) minor of G , and H is non-planar, then so is G . Hence, any graph with K_5 or $K_{3,3}$ as a (topological) minor cannot be planar. The converse, however, is also true:

Theorem 115 (Kuratowski, 1930). *A finite simple graph is planar if and only if it contains neither a K_5 nor a $K_{3,3}$ as a topological minor.*

Theorem 116 (Wagner, 1937). *A finite simple graph is planar if and only if it contains neither a K_5 nor a $K_{3,3}$ as a minor.*

The proof of the remaining implications happens in three steps: First, one shows that a graph contains K_5 or $K_{3,3}$ as a topological minor if and only if it contains one of those as a minor. Second, one proves the statement for 3-connected graphs, and then one shows that any edge-maximal graph without such a minor is 3-connected. We will only focus on the second step here.

The first step is the following lemma, which we will use without proof:

Lemma 117. *Let G, H be finite simple graphs.*

- (i) *If H is a topological minor of G , then H is also a minor of G .*
- (ii) *If $K_{3,3}$ is a minor of G , then it is also a topological minor of G .*
- (iii) *If K_5 is a minor of G , then K_5 or $K_{3,3}$ are a topological minor of G .*

(For reference, these statements are a combination of Proposition 1.7.3 and Lemma 4.4.2 in Diestel's book)

Here comes the second step:

Lemma 118. *Let $G = (V, E)$ be a finite simple 3-connected graph without a K_5 or a $K_{3,3}$ as minor. Then G is planar.*

Proof. Use induction on $|V|$. For $|V| = 4$ we have $G = K_4$, and the claim is evident. So, assume that $|V| > 4$, and that it holds for all graphs on fewer vertices. Since G is 3-connected, by Lemma 104 there exists an edge $e = \{x, y\}$ in G such that the contraction G/e is again 3-connected. Moreover, G/e is a minor of G and hence can't contain a K_5 or $K_{3,3}$ minor. Thus, by induction hypothesis, G/e is planar.

Consider a planar embedding of G/e . Removing the vertex v_{xy} (the one obtained from contracting $e = \{x, y\}$) is then again a planar graph, and in that drawing of $(G/e) - \{v_{xy}\}$ there is one face f that contained v_{xy} . Since G/e is 3-connected, the boundary of f is a cycle C . Denote by X the set of neighbours of x in G except for y . Then $X \subseteq V(C)$. Symmetrically, denote by Y the set of neighbours of y in G except for x . Then $Y \subseteq V(C)$. Taking our embedding of G/e and removing all the edges $\{v_{xy}, w\}$ for $w \in Y \setminus X$, we get an embedding of $G - \{y\}$, with the vertex v_{xy} replacing x . Our goal is to show that to this embedding of $G - \{y\}$, we can add back y and all its edges without destroying planarity.

Observe that the vertices in X (call them x_1, \dots, x_r in counterclockwise fashion around C) partition the cycle C into paths P_i that go from x_i to x_{i+1} ($i = 1, \dots, r$ with the convention that $x_{r+1} = x_1$). We need to show that all vertices in Y lie on the same path P_i , since then we can insert y in the face bounded by P_i , $\{x_i, x\}$ and $\{x, x_{i+1}\}$.

Suppose for a contradiction that not all neighbours of y lie on the same path P_i . We distinguish three cases:

- Suppose $y_1 \in Y \setminus X$ lies on the interior of some P_i , and there is another neighbour y_2 of Y that does not lie on P_i . Then $\{x, y_1, y_2\}$ and $\{y, x_i, x_{i+1}\}$ form the vertices of a topological $K_{3,3}$, a contradiction.
- Suppose that y has three neighbours that are also neighbours to x , say x_i, x_j, x_k . But now $\{x, y, x_i, x_j, x_k\}$ form the vertex set of a topological K_5 , again a contradiction.
- In the last remaining case, $\deg(y) = 3$ (because G is 3-connected, it cannot have lower degree) and two of those neighbours are shared with x , say x_i and x_k – the third neighbour of y being x itself. If x_i and x_k don't lie on a common path P , then there must be vertices x_j and x_ℓ in between them, and then $\{x_i, x, x_k\}$ and $\{x_j, y, x_\ell\}$ form once again the vertices of a topological $K_{3,3}$.

Observe that these are really all the cases that can occur, but they are not mutually exclusive. □

The third step requires the following two lemmas:

Lemma 119. *Let $G = (V, E)$ be a finite simple graph with $\kappa(G) \leq 2$. Let $V_1, V_2 \subseteq V$ such that $V_1 \cap V_2$ is a separating set of G with $|V_1 \cap V_2| = \kappa(G)$. Set $G_i = G[V_i]$ for $i = 1, 2$. If G is edge-maximal without a K_5 or a $K_{3,3}$ as a topological minor, then so are G_1 and G_2 , and then $G[V_1 \cap V_2] = K_2$.*

Lemma 120. *Let G be a finite simple graph on at least 4 vertices. If G is edge-maximal without K_5 or $K_{3,3}$ as topological minors, then G is 3-connected.*

Those two lemmas (4.4.4 and 4.4.5 in Diestel) together say that even if our graph is not 3-connected, but without a K_5 or a $K_{3,3}$ as a minor, you can add edges until it becomes edge-maximal, in which case it is 3-connected. Then, Lemma 118 states that the graph is planar.

Lecture 12 Vertex colourings

Definition 121. Let $G = (V, E)$ be a finite simple graph. A (*proper*) k -colouring of G is a map $V \rightarrow \{1, \dots, k\}$, where we think of the numbers $1, \dots, k$ as colours, such that no two adjacent vertices are mapped to the same colour. The *chromatic number* of G , denoted $\chi(G)$, is the smallest integer k such that G admits a proper k -colouring.

Unless otherwise stated, proper colourings will be the only ones we consider, and we will therefore drop the adjective “proper”.

Example 122. We have the following values for chromatic numbers: $\chi(K_n) = n$, $\chi(G) = 1$ iff G doesn't contain any edges, $\chi(G) = 2$ iff G is a bipartite graphs with at least one edge. Moreover,

$$\chi(C_n) = \begin{cases} 2 & \text{if } n \text{ is even} \\ 3 & \text{if } n \text{ is odd} \end{cases}.$$

Definition 123. Let $G = (V, E)$ be a finite simple graph. A set $C \subseteq V$ is a *clique* if the induced subgraph $V[C]$ is complete. The *clique number* of G , denoted $\omega(G)$ is the largest integer k such that there exists a clique of size k in G .

Definition 124. Let $G = (V, E)$ be a finite simple graph. A set $S \subseteq V$ is *independent* if there are no edges between any two vertices of S . The *independence number* of G , denoted $\alpha(G)$, is the maximum size of an independent set in G .

Large cliques mean we need many colours to colour a graph, whereas large independent sets mean we can colour many vertices with the same colour. More precisely:

Lemma 125. For a finite simple graph $G = (V, E)$, we have $\chi(G)\alpha(G) \geq |V|$ and $\chi(G) \geq \omega(G)$.

Proof. If G contains a clique on k vertices, then at least k colours are required to colour this clique, hence $\chi(G) \geq k$ and therefore $\chi(G) \geq \omega(G)$. On the other hand, any $\chi(G)$ -colouring of G gives a partition of V into (not necessarily non-empty) independent sets $V_1, \dots, V_{\chi(G)}$, where V_i is the set of all vertices having colour i , for $i = 1, \dots, \chi(G)$. Thus

$$|V| = \sum_{i=1}^{\chi(G)} |V_i| \leq \chi(G)\alpha(G),$$

since $|V_i| \leq \alpha(G)$ for all i . □

Lemma 125 gives lower bounds for the chromatic number. Upper bounds can be achieved by constructing colourings. One useful tool is the following algorithm:

Greedy colouring algorithm: Let $G = (V, E)$ be a finite simple graph, with an ordering on their vertices, say $V = \{v_1 \prec v_2 \prec \dots\}$. On the set of colours, we consider the usual ordering $1 < 2 < 3 < \dots$. As long as not all vertices are coloured, take the first vertex v with respect to the ordering \prec that is not yet coloured, and colour it with the smallest (w.r.t. $<$) colour that has not yet been used on any neighbours of v . Then continue with the next vertex.

The number of colours required by the greedy colouring algorithm strongly depends on the ordering used, and can become almost arbitrarily bad. However, we get the following estimate:

Lemma 126. *For a finite simple graph $G = (V, E)$ with maximal degree Δ , the chromatic number satisfies $\chi(G) \leq \Delta + 1$.*

Proof. Take any ordering on the vertices and run the greedy colouring algorithm. Then, when a vertex v is about to be coloured, it has at most Δ already coloured neighbours, that have at most Δ different colours. Thus, even in the worst case, v can be coloured with the $(\Delta + 1)$ -th colour, hence $\chi(G) \leq \Delta + 1$. \square

Moreover, the greedy algorithm can be optimal if the ordering is chosen correctly:

Lemma 127. *Let $G = (V, E)$ be a finite simple graph with chromatic number $\chi(G)$. Then there exists an ordering on the vertices such that the greedy colouring algorithm with respect to this ordering requires exactly $\chi(G)$ colours.*

Proof. Fix a $\chi(G)$ -colouring of G , and order the vertices according to their colours: The vertices with colour 1 go first, then the vertices with colour 2, etc. Because every colour class is an independent set, the greedy colouring algorithm will not require more than $\chi(G)$ colours. \square

In Example 122, we have seen two examples where equality in Lemma 126 holds: The complete graphs and the cycle graphs of odd length. As it turns out, those are the only (connected) graphs for which this is true:

Theorem 128 (Brooks, 1941). *Let $G = (V, E)$ be a finite simple connected graph with maximum degree Δ . Then $\chi(G) \leq \Delta$ unless G is complete or a cycle of odd length.*

The first part of the proof is the following lemma:

Lemma 129. *If $G = (V, E)$ is a finite simple connected graph with maximum degree Δ . If there is a vertex v with $\deg(v) < \Delta$, then $\chi(G) \leq \Delta$.*

Proof. Explore the graph using *breadth-first search*, starting from v : List the neighbours of v in some order, and mark them as discovered. Then, go to the first vertex in this list, append all its yet undiscovered neighbours to the back of the list and mark them as discovered. Then delete the current vertex from the list and repeat the procedure until the list is empty. The order in which the vertices are visited is called the breadth-first-ordering of V starting from v .

Let \prec be the *inverse* ordering of the breadth-first-ordering, say $v_n \prec v_{n-1} \prec \dots \prec v_2 \prec v_1 = v$, and use the greedy colouring algorithm with respect to this ordering. Then, for $i \geq 2$, the vertex v_i has at most $\Delta - 1$ neighbours in $\{v_n, v_{n-1}, \dots, v_{i+1}\}$. Hence, for all of these vertices Δ colours suffice. Finally, since $\deg(v) < \Delta$, the same holds true for $v_1 = v$. \square

So, for Brooks' theorem it suffices to show the statement for regular graphs.

Proof of Brooks' theorem. Let $G = (V, E)$ be a finite simple connected and k -regular graph, which is neither complete nor a cycle of odd length. This is impossible if $k = 1$ or $k = 2$, and for $k = 2$ the only graphs satisfying these conditions are the cycles of even length, for which we know $\chi(G) = 2$. Hence, we may assume $k \geq 3$.

Case 1: Assume G has connectivity $\kappa(G) = 1$. Then there exists some vertex v as well as non-empty vertex sets $V_1, V_2 \subseteq V$ with $V = \{v\} \uplus V_1 \uplus V_2$ such that removing v splits the graph into two components with vertex sets V_1 and V_2 . Write $G_i = G[V_i \cup \{v\}]$ for $i = 1, 2$. Observe that $\Delta(G_1) = \Delta(G_2) = k$, but that $\deg_{G_1}(v), \deg_{G_2}(v) < k$. Hence, by Lemma 129 there exist k -colourings of G_1 and G_2 . Since the colours in a colouring can be arbitrarily permuted, both colourings give colour 1 to the vertex v w.l.o.g. But then the colourings can be combined to give a k -colouring of G , hence $k = \Delta(G) = \chi(G)$.

Case 2: Assume $\kappa(G) = 2$, and that there exists a minimal separating set $\{u, v\} \subseteq V$ such that u, v are non-adjacent. Analogously to case 1, let V_1, V_2 be the vertex sets of the components after removing u and v , and set $G_i = G[V_i \cup \{u, v\}]$ for $i = 1, 2$. If one of the vertices u, v (w.l.o.g. u) has $\deg_{G_i}(u) \leq \Delta - 2$ for both $i = 1, 2$, then we can k -colour both G_i in such a way that u and v have different colours from each other, but the same colours in G_1 and G_2 . Once again, gluing together the colourings gives a k -colouring in G .

Otherwise, both u, v have degree 1 in one of G_1 or G_2 . If both u, v have degree 1 in G_1 (or, analogously in G_2) then they cannot both have the same neighbour $x \in G_1$, as removing x would disconnect G , contradicting $\kappa(G) = 2$. Hence v has a unique neighbour v' in G_1 that is not a neighbour to u . But

then u and v' form another minimal separating set of non-adjacent vertices, where (after redefining G_1, G_2 accordingly) $\deg_{G_1}(u) = \deg_{G_2}(v') = 1$. In this case, however, it is once again possible to permute/manipulate the k -colourings in G_1 and G_2 in such a way that they can be glued together to a k -colouring in G .

Case 3: In any other case, let $w \in V$. There are neighbours u, v of w such that $\{u, v\} \notin E$, as otherwise G would be complete. Then G with u and v removed is still connected, otherwise case 2 would apply. In the graph with u and v removed, construct the breadth-first-ordering starting from $w = v_1$, and append it with the vertices $v_{n-1} = v, v_n = u$. Then run the greedy colouring algorithm on the inverted ordering $u = v_n \prec v = v_{n-1} \prec \cdots \prec v_2 \prec v_1 = w$. Then, u, v will be assigned colour 1, for all v_i with $i \geq 2$ there are strictly less than k neighbours already coloured, and for v_1 two of the k neighbours (u and v) share a colour. Hence, k colours suffice. \square

Definition 130. A finite simple graph $G = (V, E)$ is *perfect* if $\chi(H) = \omega(H)$ holds for all induced subgraphs H of G .

In other words, a graph is perfect if the trivial lower bound to the chromatic number is exact, not only to the graph itself, but also for all its induced subgraphs. As an example, consider a bipartite G . Then, an induced subgraph H that contains an edge (those are the only interesting ones) has $\chi(H) = 2 = \omega(H)$, and so G is perfect. For a non-example, consider C_{2n+1} with $2n + 1 \geq 5$. Then $\chi(C_{2n+1}) = 3 > 2 = \omega(C_{2n+1})$, so the odd cycles of length at least 5 are not perfect.

Definition 131. Let $G = (V, E)$ be a finite simple graph. The complement \bar{G} of G is the simple graph $(V, \mathcal{P}_2(V) \setminus E)$. Thus, two vertices are adjacent in \bar{G} if and only if they are non-adjacent in G .

Theorem 132 ((weak) perfect graph theorem; Lovász, 1972). *A finite simple graph G is perfect if and only if its complement \bar{G} is.*

It follows from Definition 130 that induced subgraphs of perfect graphs are perfect. Thus, any imperfect graph needs to have minimal induced subgraphs that are imperfect - therefore, one could hope to describe perfect graphs in terms of forbidden induced subgraphs, similar to how planar graphs can be characterized by forbidden (topological) minors. We have already seen that odd cycles of length at least 5 are imperfect, and by Lovász's theorem, it follows that their complements are also imperfect. It was a conjecture from Berge (1961/63) that those would be all forbidden induced subgraphs, and indeed:

Theorem 133 ((strong) perfect graph theorem; Chudnovsky, Robertson, Seymour, Thomas, 2006). *A finite simple graph is perfect if and only if it*

contains neither an odd cycle of length at least 5 nor the complement thereof as an induced subgraph.

Observe that the strong perfect graph theorem implies the weak, since the condition in Theorem 133 is closed under taking complements.

Lecture 13 More on colourings

Last time, we have shown in Lemma 125 that $\chi(G) \geq \omega(G)$, so that the existence of cliques in G provides a lower bound for the chromatic number of G . A natural question to ask is how bad this bound can become. That is, if we require $\omega(G) \leq k$, how large can $\chi(G)$ become? Sadly, the following holds:

Theorem 134. *For any integer n , there exists a finite simple graph $G_n = (V_n, E_n)$ that does not contain a triangle (and hence has $\omega(G_n) \leq 2$) such that $\chi(G_n) = n$.*

Proof. For $n = 1, 2$ we find K_1 and K_2 , respectively. For bigger n , we use the following construction:


Let $G = (V, E)$ be a finite simple graph. Denote its vertices by $V = \{v_1, \dots, v_m\}$. The *Mycielskian* $M(G)$ of G is the finite simple graph on the $2m + 1$ vertices $\{v_1, \dots, v_m, w_1, \dots, w_m, x\}$ constructed in the following way: On the vertices v_1, \dots, v_m , we have a copy of G , that is $M[\{v_1, \dots, v_m\}] = G$. For each vertex $w_i, i = 1, \dots, m$, draw an edge to a vertex v_j iff v_j is a neighbour of v_i . Finally, draw an edge between x and w_i for every $i = 1, \dots, m$.

Lemma 135. *Let G be a finite simple graph.*

- (i) *If G is triangle-free, then so is $M(G)$.*
- (ii) *If $\chi(G) = k$, then $\chi(M(G)) = k + 1$.*

This lemma allows us to conclude the proof: Since K_2 is triangle-free with chromatic number 2, $M(K_2) \cong C_5$ is triangle-free with chromatic number 3, and by iterating Mycielski's construction, we obtain triangle-free graphs of arbitrarily high chromatic number. \square

The following picture shows $M(K_2)$ and $M(C_5)$.



Proof of Lemma 135. (i) Let $G = (V, E)$ be a triangle free graph on vertices v_1, \dots, v_m . Using the notation from the proof of Theorem 134, we note that $\{w_1, \dots, w_m\}$ is an independent set in $M(G)$. Thus, if there is a triangle in $M(G)$, it cannot contain more than one vertex from $\{w_1, \dots, w_m\}$ (and thus also cannot contain the vertex x). On the other hand, G is triangle-free. Hence, the only remaining possibility for a triangle is to be spanned by vertices v_i, v_j, w_k for some indices i, j, k . Then $\{v_i, v_j\} \in E$, but by construction of $M(G)$, v_i and v_j also have the common neighbour v_k , so $G[\{v_i, v_j, v_k\}]$ is a triangle in G , a contradiction. Therefore, $M(G)$ is triangle-free.

(ii) Assume $G = (V, E)$ has chromatic number k , and fix a k -colouring of G . The construction of $M(G)$ enables us to extend this colouring to $M(G)$ by colouring w_i with the same colour as v_i . Then x has neighbours in k different colours, and $M(G)$ is $k + 1$ -colourable. It remains to show that $M(G)$ is not also k -colourable.

To this end, assume that $M(G)$ is k -colourable. W.l.o.g. the colour k is assigned to the vertex x . Then the vertices w_1, \dots, w_m are coloured with $1, \dots, k - 1$. Let $A \subseteq \{v_1, \dots, v_m\}$ be the vertices that are assigned colour k . For each $v_i \in A$, change its colour to the colour of w_i . Does this produce a proper colouring? The set A is independent as it was a subset of a colour class in $M(G)$. Moreover, any neighbour of v_i is a neighbour of w_i , hence any neighbour to v_i must have a colour different from the colour of w_i ; so, we obtain indeed a proper colouring. Restricting this colouring to the vertices $\{v_1, \dots, v_m\}$ gives a $k - 1$ -colouring of G , which is a contradiction due to $\chi(G) = k$. \square

So, Theorem 134 states that a graph might have arbitrarily large chromatic number, even if it does not contain any large cliques. Similarly, one could try to outlaw the existence of short cycles: Given $k, \ell \in \mathbb{N}$, is there a graph with $\chi(G) > k$ and no subgraph C_3, C_4, \dots, C_ℓ ? Surprisingly:

Theorem 136 (Erdős, 1959). *For any integer k , there exists a finite simple graph G with $\chi(G) > k$ and such that G contains no cycles of length $\leq k$.*

We will give a proof of this result later, in Lecture 19, as an application of the *probabilistic method*.

In the 1850s, the question was raised about how many colours are required to paint a map, such that every connected region on the map gets a different colour. In particular, the conjecture was raised that 4 colours might suffice. Of course, in graph theoretic language, this conjecture says nothing more than “Every planar graph is 4-colourable”. Despite early “proofs” in the 1870s and ’80s (that later turned out to be false), the first bit of progress made was by Heawood in 1891, after finding a mistake in one of those earlier arguments:

Theorem 137. *If $G = (V, E)$ is a planar graph, then $\chi(G) \leq 5$.*

Proof. We use induction on $|V|$. For $|V| = 1$ we have $\chi(G) = 1$, so assume that for some n , we have $\chi(G) \leq 5$ for all planar graphs with $|V| = n$. Let G be a planar graph with $|V| = n + 1$. We use the following *Fact*: There exists a vertex $v \in V$ with $\deg(v) \leq 5$.

We choose such a v with minimal degree. Let G' denote the graph $G - \{v\}$. Then G' is planar and by the induction hypothesis has $\chi(G') \leq 5$.

Case 1: If $\deg(v) \leq 4$, then any 5-colouring of G' can be extended to a 5-colouring of G by giving v a colour that is not used among its neighbours.

Case 2: If the minimal degree of G is 5, and any two neighbours of v have the same colour, then we can repeat the argument from case 1. So assume that all 5 neighbours of v have different colours. Fix a planar embedding of G and call the neighbours v_1, \dots, v_5 in a counterclockwise fashion s.t. vertex v_i has (w.l.o.g.) colour i for $i = 1, \dots, 5$. Let $V_{1,3}$ be the set of all vertices in G' that are coloured either 1 or 3, and let $C_{1,3}^{(1)}$ be the connected component of $G'[V_{1,3}]$ that contains v_1 . Similarly, let $C_{1,3}^{(3)}$ be the connected component of $G'[V_{1,3}]$ that contains v_3 . If $C_{1,3}^{(1)} \neq C_{1,3}^{(3)}$ then the two components have to be disjoint, and we can swap the colours 1 and 3 in one of them, to the effect that v_1 and v_3 end up with the same colour. Then v is only adjacent to vertices in 4 different colours, and we can again argue like in case 1.

So, let's assume $C_{1,3}^{(1)} = C_{1,3}^{(3)}$, in which case there is a path from v_1 to v_3 using only vertices with colours 1 and 3. This path, together with the edges $\{v_1, v\}, \{v, v_3\}$ and the vertex v form a cycle C , which in the embedding of G divides \mathbb{R}^2 into an inside and an outside. However, by our choice of embedding, v_2 and v_4 lie in different components of $\mathbb{R}^2 \setminus C$. Hence, performing the same construction as with colours 1 and 3, the components $C_{2,4}^{(2)}$ and $C_{2,4}^{(4)}$ have to be disjoint. Thus, we can recolour one of them by swapping colours 2 and 4. Then, v is once again only adjacent to vertices of 4 different colours. \square

Sadly, the existence of planar graphs of minimum degree 5 prevents a modification of this argument into a proof of the 4-colour-conjecture. The 4-colour-conjecture was eventually proven in the 1970s by Appel and Haken, who showed that any minimal counterexample can be reduced to one of 1834 configurations, and then used computer search to verify that none of those provided a counterexample. The proof was since further improved, reducing the number of configurations and removing inaccuracies from the original proof, and is now widely accepted. Hence, we have:

Theorem 138 (4-colour-theorem). *Any planar graph is 4-colourable.*

We also quote the following without proof:

Theorem 139 (Grötzsch). *Any planar graph without triangles is 3-colourable.*

Let us conclude our discussion of vertex-colourings by looking at some estimates coming from *spectral graph theory*. For this, recall the definition of the adjacency matrix A of G from Lecture 4: If G is a finite simple graph on vertices labelled $1, \dots, n$, then A is the $n \times n$ -matrix with entry 1 in position (i, j) if i and j are adjacent, and 0 otherwise.

Since the adjacency matrix is symmetric, the spectral theorem from linear algebra implies that it has n real eigenvalues, which can be arranged in increasing size, say $\lambda_{\min}(A) \leq \dots \leq \lambda_{\max}(A)$. Moreover, the adjacency matrix encodes all the structure of the graph, so it should not be too surprising that e.g. some graph properties can be read off from the characteristic polynomial of A . And the connections don't stop there. We also have the following lemma:

Lemma 140. *Let $G = (V, E)$ be a finite simple graph, and H an induced subgraph of G . Denote by A_G and A_H the adjacency matrices of G and H , respectively. Then*

$$\lambda_{\min}(A_G) \leq \lambda_{\min}(A_H) \leq \lambda_{\max}(A_H) \leq \lambda_{\max}(A_G).$$

and $\delta(G) \leq \lambda_{\max}(A_G) \leq \Delta(G)$, where $\delta(G)$ and $\Delta(G)$ denote the minimal and maximal degree of G , respectively.

We shall use this lemma (which relies on the technique of Rayleigh-quotients from linear algebra) to prove the following theorem, providing yet another estimate for $\chi(G)$.

Theorem 141 (Wilf, 1967). *For any finite simple graph G with adjacency matrix A_G we have $\chi(G) \leq 1 + \lambda_{\max}(A_G)$.*

Proof. Among all induced subgraphs of G , there exists a minimal subgraph H (w.r.t. inclusion) with $\chi(H) = \chi(G)$. Let v be a vertex of H . Then $H - \{v\}$ admits a $\chi(G) - 1$ -colouring, and if $\deg_H(v) < \chi(G) - 1$, then this colouring could be extended to a $\chi(G) - 1$ -colouring of H , contradicting our choice of H . Hence, the minimum degree in H is at least $\chi(G) - 1$. Denote by A_H the adjacency matrix of H . Then,

$$\chi(G) \leq 1 + \delta(H) \leq 1 + \lambda_{\max}(A_H) \leq 1 + \lambda_{\max}(A_G),$$

where we used the inequalities from Lemma 140. □

Wilf's theorem tends to be much better than the estimate $\chi(G) \leq 1 + \Delta$.

There also exists a lower spectral estimate for the chromatic number:

Theorem 142 (Hoffman, 1970). *For any finite simple graph G with at least one edge we have $\chi(G) \geq 1 + \frac{\lambda_{\max}(A_G)}{-\lambda_{\min}(A_G)}$.*

Note that $\lambda_{\min}(A_G) < 0$ unless $A = 0$, since the sum of eigenvalues equals zero.

Lecture 14 Edge-colourings and Ramsey Theory

Let us look at another variation of colourings: This time, we're colouring the edges in a graph.

Definition 143. Let $G = (V, E)$ be a finite simple graph. A (*proper*) k -edge-colouring is a map $E \rightarrow \{1, \dots, k\}$ such that no two edges with a common endpoint are assigned the same colour. The *chromatic index* or *edge-chromatic number*, denoted $\chi'(G)$, is the smallest integer k for which a proper k -edge-colouring exists.

Any proper edge-colouring partitions the edge set into sets of edges with the same colour, and by definition, those sets are matchings. So, similar to how vertex-colourings were closely related to independent sets, edge-colourings are related to matchings.

Moreover, if v is a vertex of maximum degree, then the Δ edges incident to this vertex must be of distinct colours, hence $\chi'(G) \geq \Delta$. It is a classical result that this bound is attained for bipartite graphs:

Theorem 144 (König's line-colouring theorem, 1916). *For every finite simple bipartite graph G with maximal degree Δ , we have $\chi'(G) = \Delta$.*

Proof. Let $G = (V, E)$ be bipartite. We use induction on the number $m := |E|$ of edges. For $m = 0$ there is nothing to show. So, assume that the claim is true for all graphs on $m \geq 0$ edges, and consider a graph G with $m + 1$ edges.

Denote by Δ the maximal degree of this G , and fix an arbitrary edge $\{v, w\}$ from G . Let G' be the graph obtained from G by deleting the edge $\{v, w\}$. By the induction hypothesis, there is a proper Δ -edge-colouring of G' . Moreover, each of v, w are incident to at most $\Delta - 1$ edges in G' . Hence, there are colours i, j such that v is not incident to an edge of colour i and w is not incident to an edge of colour j . If $i = j$, then we can colour $\{v, w\}$ by i and in this way obtain a Δ -edge-colouring of G .

So assume $i \neq j$. Let A, B be the vertex bi-partition of G , wlog $v \in A, w \in B$. Consider a trail in G' starting at v using alternately edges of colour j and edges of colour i (Recall that a trail cannot repeat edges). Since v is not incident to an edge of colour i , the first edge on any such trail has to be of colour j . Moving along this trail will alternate between vertices from A and B , where an edge of colour j will be occur exactly when moving from A to B . Hence, no such trail can contain the vertex w , since w is not incident to a vertex of colour j . Let $C_{i,j}$ denote the component that contains v of the subgraph of G' that is induced by all edges of colour i or j . We just showed that $w \notin C_{i,j}$, so we can swap colours i and j in $C_{i,j}$ without affecting w . After swapping, we have a colouring of G' in which both v, w are not incident

to an edge of colour j , so we can extend the colouring as we did in the first case. \square

Besides having $\chi'(G) = \Delta$, what else can happen? Not much:


Theorem 145 (Vizing, 1964). *Let $G = (V, E)$ be a finite simple graph with maximum degree Δ . Then either $\chi'(G) = \Delta$ (in which case G is said to be class 1) or $\chi'(G) = \Delta + 1$ (in which case G is said to be class 2).*

Curiously enough, there is no known characterisation that distinguishes between class 1 and class 2 graphs, but many partial results are known. Let's look at an example instead:

Example 146. The complete graphs K_n are of class 1 if n is even, and of class 2 if n is odd. This is because the edge set of a K_{2n} is a disjoint union of $n - 1$ perfect matchings, and we can assign one colour to each matching. For K_{2n-1} , after choosing which edges to colour with 1, there will always be at least one vertex remaining with all incident edges being uncoloured. Hence, Δ colours cannot suffice.

Definition 147. Let $G = (V, E)$ be a finite simple graph. Its *line graph* $L(G)$ has vertex set E , where two vertices e_1, e_2 are adjacent if and only if the edges e_1 and e_2 shared an endpoint in G .

As an example, here are $K_{2,3}$ and its line graph:



Directly from the definition we can translate concepts from G and $L(G)$, for example:

- If $H = G\langle S \rangle$ for $S \subseteq E(G)$ is an edge-induced subgraph of G , then $L(H)$ is exactly the by $S \subseteq E(G) = V(L(G))$ vertex-induced subgraph of $L(G)$. In other words, edge-induced subgraphs of G correspond precisely to vertex-induced subgraphs of $L(G)$.
- The (proper) edge-colourings of G are exactly the (proper) vertex colourings of $L(G)$. In particular, $\chi'(G) = \chi(L(G))$.

- Vertices with degree d in G correspond to cliques of size d in $L(G)$. In particular, $\Delta(G) = \omega(L(G))$.
- The edges of a matching in G don't share a endpoints, hence they are mutually non-adjacent vertices in $L(G)$. Thus matchings in G are independent sets in $L(G)$.

Hence, Theorems 144 and 145 can be re-interpreted as statements about line graphs. Specifically, since any edge-induced subgraph of a bipartite G is again bipartite, Theorem 144 then states that $\chi(H) = \omega(H)$ for any vertex-induced subgraph H of $L(G)$. Hence, we could rephrase the theorem as “The line graphs of bipartite graphs are perfect”.

A minimum of Ramsey theory. Consider now a graph K_n , and colour its edges with two colours, say red and blue. (Obviously, this requires us to drop the condition that edges with a common endpoint need to have different colours).

Definition 148. Let $k \in \mathbb{N}$. The k -th Ramsey number $R(k)$ is defined to be the smallest integer n such that every red-/blue-colouring of K_n contains a monochromatic K_k .

There is an alternative formulation: Given any red-/blue-colouring of K_n , we can construct any arbitrary graph on n -vertices by including exactly the blue edges. The Ramsey number $R(k)$ is then the smallest n such that every finite simple graph on n vertices contains a clique on k vertices (if a monochromatic K_k was blue) or an independent set on k vertices (if a monochromatic K_k was red).

Notice that the definition does not guarantee the existence of any of the $R(k)$.

Theorem 149 (Ramsey,1930). *For any $k \geq 2$, we have $R(k) \leq 2^{2k-3}$. In particular, Ramsey-numbers exist.*

Proof. For brevity, set $n = 2^{2k-3}$. In K_n , we construct sets V_1, \dots, V_{2k-2} of vertices, together with distinguished vertices $v_i \in V_i$ such that

- (i) For $i = 1, \dots, 2k - 2$, we have $|V_i| = 2^{2k-2-i}$.
- (ii) For $i = 2, \dots, 2k - 2$, we have $V_i \subseteq V_{i-1} \setminus \{v_{i-1}\}$.
- (iii) The vertex v_{i-1} only has edges of one colour to all vertices in V_i for $i = 2, \dots, 2k - 2$.

For V_1 , we have to take the entire vertex set of K_n , and we pick $v_1 \in V_1$ arbitrarily. Assume we have chosen V_{i-1} and v_{i-1} according to (i)-(iii). Then,

$V_{i-1} \setminus \{v_{i-1}\}$ can be partitioned into two sets: The vertices that are adjacent to v_{i-1} with a red edge, and those that are adjacent to v_{i-1} with a blue edge. One of those sets contains at least 2^{2k-2-i} vertices. Choose V_i from this set according to (i)-(iii), and $v_i \in V_i$ arbitrarily.

This gives $2k-3$ vertices v_1, \dots, v_{2k-3} . Hence, there are at least $k-1$ vertices v_{i-1} among them that connect to the respective V_i with a common colour. Then these $k-1$ vertices together with v_{2k-2} induces a monochromatic K_k . \square

Example 150. We have $R(1) = 1$ and $R(2) = 2$. We'll show next that $R(3) = 6$. For higher k , it is known that $R(4) = 18$, and above that only bounds are known: $43 \leq R(5) \leq 49$ and $102 \leq R(6) \leq 165$, and determining the exact values seems to be completely out of reach of current methods.

Proposition 151. *We have $R(3) = 6$, that is, every red-/blue-colouring of K_6 contains a monochromatic triangle, but there is a red-/blue-colouring of K_5 without a monochromatic triangle.*

Proof. The part about K_5 is left as an exercise. For K_6 , observe that there are $\binom{6}{3} = 20$ triangles in K_6 . Consider triplets (x, y, z) of vertices such that the edge $\{x, y\}$ is red and the edge $\{y, z\}$ is blue.

Case 1: If all edges incident to y are of the same colour, then y can not be the middle vertex of such a triplet.

Case 2: If all but one edge incident to y are of the same colour, then y can be the middle vertex of at most 4 such triplets.

Case 3: If two edges incident to y are of one colour, and the other three are of the other colour, then y can be the middle vertex of at most 6 such triplets.

Hence, there are at most 36 such triplets. It is clear that such a triplet can only occur in a non-chromatic triangle, but every non-chromatic triangle yields two different triplets. Hence, there are at least $20 - \frac{36}{2} = 2$ monochromatic triangles. \square

Lecture 16 Szemerédi's regularity lemma

Large graphs and networks tend to occur in real-life applications in abundance: There's the internet (both in its physical appearance as a network of devices, routers and servers, and in its virtual structure: webpages and hyperlinks connecting them), social relations and the brain (neurons and how they connect to each other), just to name a few. Analyzing the exact structure of those graphs with often billions of vertices is almost always impossible, after all, the number of potential edges grows like the square of the number of vertices.

It is all the more important to find approximations and limiting behaviours of large graphs. Of course, both of these notions bring up many immediate follow-up questions: What does it mean for a graph to be an approximation of another graph? How exactly should we understand a limit of a (growing) sequence of graphs, and what kind of object should that be? What graph parameters behave well under approximations and/or limits? And while questions of this kind are often subject to current research, there are some established tools, one of them being Szemerédi's regularity lemma.

Recall that a finite simple graph $G = (V, E)$ on n vertices has at most $\binom{n}{2}$ edges, and that a bipartite graph with $V = A \uplus B$ has at most $|A||B|$ edges.

Definition 152. Let $G = (V, E)$ be a finite simple graph on n vertices. The quantity $d(G) := |E| \binom{n}{2}^{-1}$ is the *density* of G . Given $X, Y \subseteq V$ disjoint, denote by $e(X, Y)$ the number of edges in G with one endpoint in X and the other in Y . The quantity $d(X, Y) := \frac{e(X, Y)}{|X||Y|}$ is the *density* of the pair (X, Y) .

By construction, both $d(G)$ and $d(X, Y)$ are numbers between 0 and 1.

Definition 153. Let $\varepsilon > 0$. A pair $A, B \subseteq V$ of disjoint vertex sets is called ε -*regular* if $|d(X, Y) - d(A, B)| \leq \varepsilon$ for all $X \subseteq A$, $Y \subseteq B$ with $|X| \geq \varepsilon|A|$ and $|Y| \geq \varepsilon|B|$.

In other words, a pair (A, B) being regular means that as soon as we sample sets X, Y that are not too small, the density of (X, Y) will not be too far away from the density of (A, B) . This is also what would happen if the graph was random, with every edge being present independently from each other with some fixed probability p .

Theorem 154 (Szemerédi's regularity lemma, 1978). *For every $\varepsilon > 0$ and every $m \in \mathbb{N}$ there is an $M \in \mathbb{N}$ such that every finite simple graph $G = (V, E)$ on at least M vertices has a partition $V = V_0 \uplus V_1 \uplus \dots \uplus V_k$ such that*

1. $m \leq k \leq M$
2. $|V_0| \leq \varepsilon|V|$ (and V_0 may be empty)

3. $|V_1| = \dots = |V_k|$
4. all but at most εk^2 of the pairs (V_i, V_j) are ε -regular for $1 \leq i, j \leq k$.

The idea of the regularity lemma is as follows: We have a small garbage bin V_0 to ensure that the other sets V_1, \dots, V_k are of the same size and to collect special vertices. Letting m grow large, we can moreover ensure that the size of V_1, \dots, V_k becomes small enough so that most edges are in between different sets. Then, up to any given tolerance ε , any sufficiently large graph looks mostly like a random graph, except for some small proportion of pairs, and on the exceptional set V_0 .

When applying the regularity lemma, it is often useful to discard unnecessary edges and vertices. In particular, the exceptional set V_0 , any internal edges to the V_i and pairs (V_i, V_j) with too low of an edge density tend to be annoying. However, discarding too many edges might destroy important features of the graph. Luckily, there is the following version of the regularity lemma:

Lemma 155 (Degree version of the regularity lemma). *For every $\varepsilon > 0$ and every $m \in \mathbb{N}$ there is an $M \in \mathbb{N}$ such that for every finite simple graph $G = (V, E)$ on at least M vertices and for any $\delta \in [0, 1]$, there exists a partition $V = V_0 \uplus V_1 \uplus \dots \uplus V_k$ and a spanning subgraph $G' \subseteq G$ such that*

1. $m \leq k \leq M$
2. $|V_0| \leq \varepsilon|V|$ (and V_0 may be empty)
3. $|V_1| = \dots = |V_k|$
4. $\deg_{G'}(v) > \deg_G(v) - (\delta + \varepsilon)|V|$ for all $v \in V$
5. The sets V_i for $i \geq 1$ are independent in G'
6. all of the pairs (V_i, V_j) for $1 \leq i, j \leq k$ are ε -regular, either with density $\geq \delta$ or with density 0.

Having such a G' , we can construct the graph $G'' = G' - V_0$ for even more comfort. This graph still has most of the edges of G (assuming $d(G) > 0$), because we have

$$\deg_{G''}(v) > \deg_G(v) - (\delta + \varepsilon)|V| - |V_0| \geq \deg_G(v) - (\delta + 2\varepsilon)|V|$$

for all vertices $v \in V(G'')$. It then follows from the handshaking lemma that

$$|E(G'')| > |E(G)| - (\delta + 3\varepsilon)\frac{|V|^2}{2}. \quad (2)$$

Given $\varepsilon > 0$ and a certain threshold $\delta \in [0, 1]$, we can *reduce* a partitioned graph with $V = \{V_1, \dots, V_k\}$ in the following way: Let R have vertices

$1, \dots, k$ such that two vertices i and j are adjacent iff (V_i, V_j) is an ε -regular pair with density $\geq \delta$. So, the reduced graph keeps track of which partition sets are densely connected, and which ones are not.

Conversely, given any finite simple graph R and a positive integer t , we can “blow up” R by replacing every vertex x of R with an independent set V_x of t vertices, and by replacing an edge $\{x, y\}$ in R with all possible edges between the independent sets V_x and V_y . In other words, $V_x \cup V_y$ induces a complete bipartite graph in this new graph. We denote the graph obtained in such a way by $R(t)$.

Lemma 156 (“Key lemma”). *Let $\delta > \varepsilon > 0$, and let G, G'' , and R be as above (meaning G'' is obtained from Lemma 155 after deleting V_0 , and R is the reduced graph of G''). Fix $t \in \mathbb{N}$. Let H be a subgraph of $R(t)$ with maximum degree $\Delta(H)$. If*

$$\varepsilon \leq \frac{(\delta - \varepsilon)^{\Delta(H)}}{2 + \Delta(H)} \quad \text{and} \quad t - 1 \leq \frac{(\delta - \varepsilon)^{\Delta(H)}}{2 + \Delta(H)} m$$

then H is a subgraph of G'' , and therefore also of G .

Why is this lemma “key”? Imagine we obtained G' with some δ from the degree formulation of the regularity lemma, lemma 155. Then, after throwing away V_0 we obtain some G'' which will have some reduced graph R for the threshold δ . Then in $R(t)$ we replace vertices by independent sets and edges by complete bipartite graphs between them; and we can use this rigid structure to exhibit a subgraph H in $R(t)$. This lemma then states that if H satisfies some technical condition, then H was already present in G .

In the rest of the lecture, we want to look at the Erdős-Stone theorem and how to prove it with the help of Szemerédi’s regularity lemma. However, we need a bit of set-up for this.

An application. Let us consider the following, innocent-looking question: How does a graph on n vertices with the maximal number of edges but without a K_r -subgraph look like? The idea is to consider a complete $(r - 1)$ -partite graph on n vertices, where every set of the partition is as close to $n/(r - 1)$ as possible.

Definition 157. Let $n, r \in \mathbb{N}$. Then there exist unique non-negative integers a, b such that $n = a(r - 1) + b$ with $0 \leq b < r - 1$. The *Turán graph* $T_{r-1}(n)$ is the complete $(r - 1)$ -partite graph on n vertices such that b sets of the vertex-partition have $a + 1$ elements, and the other $r - 1 - b$ sets have a elements. In particular, for $n \leq r - 1$, we have $T_{r-1}(n) = K_n$. The number of edges in $T_{r-1}(n)$ are the Turán numbers $t_{r-1}(n)$.

Theorem 158 (Turán, 1941). *For $n, r \in \mathbb{N}$, any finite simple graph on n vertices that does not contain a K_r subgraph and has the maximum number of edges is a $T_{r-1}(n)$.*

Proof. It is easy to see that among complete multipartite graphs not containing a K_r , the Turán graphs contain the maximum number of edges. Thus it suffices to show that any such extremal graph $G = (V, E)$ is complete multipartite.

Consider the following notion of *vertex duplication*: For some $v \in V$, add a “copy” v' to the vertex set, and draw edges such that the neighbours of v' are exactly the neighbours of v . Now, if G is not complete multipartite, then there exist vertices x, y_1, y_2 such that x and y_1 are non-adjacent, x and y_2 are non-adjacent, but y_1 and y_2 are adjacent (Why?). In the case that $\deg(y_1) > \deg(x)$, deleting x and duplicating y_1 yields a graph without K_r and more edges than G . Analogously, $\deg(y_2) > \deg(x)$ is impossible. Hence $\deg(y_1), \deg(y_2) \leq \deg(x)$. But now, deleting both y_1 and y_2 and duplicating x twice yields a graph without K_r and more edges than G . \square

Observe that the Turán graphs $T_{r-1}(n)$ have edge density $d(T_{r-1}(n)) \rightarrow \frac{r-2}{r-1}$, which follows from the formula for the number of edges in complete multipartite graphs, Proposition 19. Also note that, while K_r -free, $T_{r-1}(n)$ contains a lot of K_{r-1} 's (unless $n < r - 1$): Any choice of picking one vertex from each set of the partition yields a K_{r-1} . Similarly, the following theorem states that if a graph on a large enough number of vertices fails to have less than $t_{r-1}(n)$ edges, then it contains a large number of K_r -subgraphs.

Theorem 159 (Erdős-Stone, 1946). *For all integers $r \geq 2$ and $s \geq 1$, and all $\beta > 0$, there exists an N such that any finite simple graph with $n \geq N$ vertices and more than $\left(\frac{r-2}{r-1} + \beta\right) \binom{n}{2}$ edges contains the complete r -partite graph $K_{s, \dots, s}$.*

Observe that necessarily $\beta \leq \frac{1}{r-1}$ if such a graph is to exist.

Proof. Let G_n be finite simple graphs on n vertices having more than $\left(\frac{r-2}{r-1} + \beta\right) \binom{n}{2}$ edges. Set $\delta = \beta/2$ and let $\varepsilon = (\beta/6)^{rs}$. Then, by the degree formulation of the regularity lemma there is a (spanning) subgraph G' together with a partition of $V(G)$ into sets V_0, V_1, \dots, V_k satisfying the properties in Lemma 155. Let $G'' = G' - V_0$ and let R be the reduced graph of G'' , with parameter ε and threshold δ . Counting the edges in R and comparing it with the edge densities in G'' yields

$$\frac{|E(R)|}{|V(R)|^2/2} \geq \frac{|E(G'')|}{n^2/2} > \frac{r-2}{r-1},$$

where for the last inequality, we relied on the estimate (2). Therefore, R has higher edge density than the Turán graphs $T_{r-1}(n)$, hence R contains an r -clique, and so $R(s)$ contains an r -partite $K_{s, \dots, s}$.

It remains to check that for $H = K_{s,\dots,s}$ and our choice of parameters, the conditions of the key lemma (Lemma 156) are satisfied. Indeed, we have $\Delta(H) = (r-1)s$, and then obtain:

$$\frac{(\delta - \varepsilon)}{(2 + \Delta(H))^{1/\Delta(H)}} > \frac{\beta}{2} - \left(\frac{\beta}{6}\right)^{rs} \geq \frac{\beta}{2} - \frac{\beta^2}{36} \geq \left(\frac{\beta}{6}\right)^r$$

for all $\beta > 0$, which implies the first technical condition. For the second one, observe that $m = |V_1| \geq \frac{n(1-\varepsilon)}{k}$. In particular, choosing n large enough, we get

$$s - 1 \leq \frac{(\delta - \varepsilon)\Delta(H)}{2 + \Delta(H)}m$$

for free. Thus, for n large enough, G_n contains $K_{s,\dots,s}$. □

One can ask similar questions for any fixed subgraph H : Given n , what is the maximum number of edges a graph on n vertices can have without containing a copy of H as a subgraph? (Such graphs are called *extremal*). This number is denoted $\text{ex}(n, H)$. By Turán's theorem, we have $\text{ex}(n, K_r) = t_{r-1}(n)$. The following corollary is a cute consequence of Turán's theorem and the Erdős-Stone theorem:


Corollary 160. *Let H be a finite simple graph containing at least one edge and having chromatic number $\chi(H)$. Then*

$$\lim_{n \rightarrow \infty} \binom{n}{2}^{-1} \text{ex}(n, H) = \frac{\chi(H) - 2}{\chi(H) - 1}.$$

Lecture 17 The Rado graph

Construct an infinite simple graph as follows: Let $V = \mathbb{N}_0$, and represent each vertex in binary. For $i < j$, draw an edge between the vertices i and j if the digit representing 2^i in the binary expression for j is a 1, otherwise don't. Here's an example: $5 = 2^2 + 2^0 = (101)_2$, so vertex 5 has neighbours 0 and 2, as well as $2^5, 2^5 + 1, 2^5 + 2, \dots$ and many more. The graph obtained in this way is known as *Rado graph*.

Another way to understand this construction is to consider the vertices neighbouring 0, 1, 2, ...: Neighbours of 0 must have a 1 in their last binary position, so 0 has an edge to all odd numbers. Similarly, neighbours of 1 have a 1 in their second-to-last binary position, and those are exactly the numbers 2, 3, 6, 7, 10, 11, ..., i.e. the numbers congruent to 2 or 3 mod 4. Compare also to the following picture showing the subgraph induced by the vertices 0, ..., 9



The Rado graph is known to have several astonishing properties, and we will look into some of them in this lecture. Our motivating goal will be to show the following theorem:

Theorem 161. *Let R be the Rado graph, and let G be any finite simple graph. Then there is an induced subgraph H of R such that $H \cong G$.*

In other words, R contains all finite simple graphs as induced subgraphs. The proof relies on the following, innocent looking property a graph can have:

Definition 162. A simple graph $G = (V, E)$ is said to have the *extension property* (P) if, for all finite disjoint subsets $U, W \subseteq V$, there exists a vertex x such that x is a neighbour to every vertex in U , but not a neighbour to any vertex in W .

It follows immediately that any graph $G = (V, E)$ with property (P) must be infinite, since for a finite graph, we could set $U = V, W = \emptyset$, and property (P) would yield the existence of a vertex not in V , which is absurd.

Lemma 163. *The Rado graph R satisfies the extension property.*

Proof. Consider two finite, disjoint subsets $U, W \subseteq \mathbb{N}_0$, say $U = \{u_1, \dots, u_m\}$ and $W = \{w_1, \dots, w_n\}$ for $m, n \geq 0$. Our goal is to construct an integer r such that r is a neighbour to u_1, \dots, u_m according to the construction of Rado's graph, but not a neighbour to w_1, \dots, w_n . So, set $\ell := 1 + \max(U \cup W)$, and define

$$r := 2^\ell + \sum_{i=1}^m 2^{u_i}.$$

With this definition, we ensure that r is larger than any number in either U or W . Moreover, all numbers in U are neighbours to r , as their binary digits occur in the binary expansion of r , by construction. Similarly, since W is disjoint to U , none of the digits of r belonging to w_1, \dots, w_n can be 1; hence r is not a neighbour to any of those vertices. \square

Proof of Theorem 161. Let $G = (V, E)$ be any finite simple graph, and assume w.l.o.g. that there is some ordering on its vertex set, say $V = \{v_1, \dots, v_n\}$ with the ordering given by the indices. We will inductively construct a sequence of induced subgraphs R_1, \dots, R_n of R such that for each k , $R_k \cong V[\{v_1, \dots, v_k\}] =: G_k$.

Begin by embedding vertex v_1 to any vertex w_1 of R , and set $R_1 = R[\{w_1\}]$. Now, assume that for $k = 1, \dots, n - 1$ we have already constructed $R_k = R[\{w_1, \dots, w_k\}]$ in such a way that the isomorphism between G_k and R_k maps v_j to w_j for all $j = 1, \dots, k$. Now, consider v_{k+1} , and set

$$\begin{aligned} U &:= \{w_j : 1 \leq j \leq k \text{ s.t. } v_j \text{ is adjacent to } v_{k+1}\} \\ W &:= \{w_j : 1 \leq j \leq k \text{ s.t. } v_j \text{ is not adjacent to } v_{k+1}\} \end{aligned}$$

Since R satisfies the extension property, there exists a vertex w_{k+1} that is a neighbour to all $w_j \in U$ and not a neighbour to each $w_j \in W$. Hence mapping v_{k+1} to w_{k+1} extends the isomorphism between G_k and R_k to an isomorphism of the induced subgraphs G_{k+1} and R_{k+1} .

Thus, we can repeatedly extend this construction until we map $G = G_n$ to R_n . \square

Observe how we did not require anything about the Rado graph in the proof of Theorem 161 other than the fact that it satisfies property (P). We thus get immediately:

Corollary 164. *Any graph with the extension property contains all finite simple graphs as induced subgraphs.*

As it turns out, having property (P) has many more interesting consequences, perhaps most importantly:

Theorem 165. *Let $G = (V, E)$ be a simple graph on a countable vertex set V , and assume G satisfies (P). Then $G \cong R$.*

Thus, up to isomorphism R is the only graph on a countable vertex set that satisfies (P).

Proof. Similarly to the proof of Theorem 161, we use property (P) to extend isomorphisms between smaller subgraphs, but require a bootstrap argument to make sure we extend to the entire infinite graph in the end.

Start by labelling the vertices of G by v_0, v_1, v_2, \dots (thus giving an ordering of the vertices), and let $G_0 := G[\{v_0\}]$ and $R_0 := R[\{0\}]$. Clearly, $G_0 \cong R_0$. Let us now assume that we have extended this isomorphism to two induced subgraphs G_n and R_n , having vertices a_0, a_1, \dots, a_n and b_0, b_1, \dots, b_n respectively, with the isomorphism mapping a_j to b_j for $0 \leq j \leq n$. Now proceed along the following two steps:

1. Let a_{n+1} be the earliest vertex in G not yet included in G_n . Set $G_{n+1} := G[\{a_0, a_1, \dots, a_{n+1}\}]$. By the argument in the proof of Theorem 161, there exists a vertex b_{n+1} in R that extends the isomorphism to $G_{n+1} \cong R_{n+1} := R[\{b_0, b_1, \dots, b_{n+1}\}]$.

2. Let b_{n+2} be the earliest vertex in R not yet included in R_{n+1} . Analogously to step 1, there exists a vertex a_{n+2} in G that extends the isomorphism to $R_{n+2} := R[\{b_0, b_1, \dots, b_{n+2}\}]$ and $G_{n+2} := G[\{a_0, a_1, \dots, a_{n+2}\}]$.

Now, starting from $G_0 \cong R_0$, we can extend the isomorphism by repeating steps 1 and 2 over and over again. It only remains to show that in this way, every vertex in either graph gets mapped to some vertex in the other graph. However, from alternating between the two steps, it follows that v_n is contained in $\{a_0, a_1, \dots, a_{2n}\}$, and a vertex n in R is contained in $\{b_0, b_1, \dots, b_{2n}\}$. Therefore, this procedure indeed extends to an isomorphism $G \cong R$. \square

The extension property is a very strong condition, and it has a number of consequences, among them:

- $\text{diam}(R) = 2$
- To all finite disjoint vertex sets U, W exist infinitely many vertices that are neighbours to every $u \in U$, and non-adjacent to every $w \in W$.
- If a finite number of edges is either added to or deleted from R , then the resulting graph is still isomorphic to R .
- For $\mathbb{N}_0 = V_1 \uplus V_2$, at least one of $R[V_1]$ and $R[V_2]$ is isomorphic to R .

Let us now consider a different construction of infinite graphs, yielding the *infinite random graph* $G(p)$:

Fix $0 \leq p \leq 1$. Let $V = \mathbb{N}$, and for any pair of different integers i, j , toss a weighted coin: In this way, the graph contains the edge ij with probability

p , and doesn't contain this edge with probability $1 - p$, independently of all other edges.

The cases $p = 0$ and $p = 1$ are not particularly interesting: For $p = 0$, $G(p)$ will contain no edges at all (with probability 1), whereas for $p = 1$, $G(p)$ turns out to be the complete graph on countably infinite many vertices (again, with probability 1). For all values in between, one would expect $G(p)$ depend both on the parameter p as well as on the randomness involved in the construction. However, this is not the case:

Theorem 166. *For $0 < p < 1$, the graph $G(p)$ satisfies the extension property with probability 1.*

By Theorem 165, this immediately implies that $G(p) \cong R$ almost surely.

As first tool for the proof, we require the following fact: If A_1, A_2, \dots is a sequence of events such that $\mathbf{P}[A_n] = 1$ for all n , then

$$\mathbf{P} \left[\bigcap_{n=1}^{\infty} A_n \right] = 1 - \mathbf{P} \left[\bigcup_{n=1}^{\infty} A_n^c \right] \geq 1 - \sum_{n=1}^{\infty} \mathbf{P}[A_n^c] = 1.$$

The second useful fact is the next lemma, which is a famous theorem in probability theory:

Lemma 167 (Borell-Cantelli). *Let A_1, A_2, \dots be a sequence of events. Denote by A the event that infinitely many of these events occur.*

- (a) *If $\sum_{n=1}^{\infty} \mathbf{P}[A_n] < \infty$, then $\mathbf{P}[A] = 0$ (i.e. with probability 1, only finitely many of the events A_1, A_2, \dots occur).*
- (b) *If $\sum_{n=1}^{\infty} \mathbf{P}[A_n] = \infty$ and the events A_1, A_2, \dots are independent, then $\mathbf{P}[A] = 1$. In particular, at least one of A_1, A_2, \dots occur with probability 1.*

Proof of Theorem 166. Let $U, W \subseteq \mathbb{N}_0$ be finite, disjoint vertex-subsets of $G(p)$. Let $r := \max U \cup W$, and define events A_n to be

$$A_n := \{(r+n) \sim u, (r+n) \not\sim w \text{ for all } u \in U, w \in W\}.$$

In other words, A_n describes the event that the vertex $r+n$ is the vertex required for property (P). For A_n to occur, we thus need $|U|$ -many edges to be in $G(p)$ (each of them is present with probability p), and $|W|$ -many edges to not be in $G(p)$ (which happens with probability $1-p$ each). We thus have

$$\mathbf{P}[A_n] = p^{|U|}(1-p)^{|W|}$$

and therefore

$$\sum_{n=1}^{\infty} \mathbf{P}[A_n] = \sum_{n=1}^{\infty} p^{|U|}(1-p)^{|W|} = \infty,$$

so by the Borel-Cantelli lemma, with probability 1 at least one of the events A_n occurs.

For given U, W , denote by $B_{U,W}$ the event that there exists a vertex verifying property (P) with probability 1. Since there are only countably many configurations for U and W (as both are assumed to be finite), we obtain

$$\mathbf{P}[G(p) \text{ satisfies } (P)] = \mathbf{P} \left[\bigcap_{U,W} B_{U,W} \right] = 1$$

by the earlier comment on the intersection of countably many events with probability 1. □

Lecture 18 The Erdős-Rényi random graph

In this lecture, we look at the most classic of all random graph models, the *Erdős-Rényi random graph* $G(n, p)$: Consider a graph on n labelled vertices. For each of the $\binom{n}{2}$ potential edges, draw the edge with probability p and don't draw it with probability $1 - p$, independently of all other edges.²

The resulting graph will depend on parameters and on the randomness of which edges to include, and will therefore have a probability distribution on the set of all labelled graphs on n vertices. In fact, if $G = (V, E)$ is a fixed labelled graph on n vertices, then the probability that a $G(n, p)$ coincides with G is

$$p^{|E|}(1 - p)^{\binom{n}{2} - |E|}$$

because the edges are chosen independently of each other. In particular, for $p = \frac{1}{2}$, we obtain a uniformly chosen random graph on n vertices. Some other properties follow in a similar fashion: The number of edges in $G(n, p)$ is a binomially-distributed random variable with parameters $\binom{n}{2}$ and p . Moreover, it is relatively easy to obtain estimates such as the following

Lemma 168. $\mathbf{P}[\alpha(G(n, p)) \geq k] \leq \binom{n}{k}(1 - p)^{\binom{k}{2}}$

Proof. The probability that any fixed set of k vertices is independent is $(1 - p)^{\binom{k}{2}}$. There are $\binom{n}{k}$ such sets, so we can number them from 1 to $\binom{n}{k}$. Denote by A_i the event that the i -th set is an independent set. Then,

$$\mathbf{P}[\alpha(G(n, p)) \geq k] = \mathbf{P}\left[\bigcup_i A_i\right] \leq \sum_i \mathbf{P}[A_i] = \binom{n}{k}(1 - p)^{\binom{k}{2}},$$

as desired. □

For fixed p and $n \rightarrow \infty$, the graph $G(n, p)$ will “tend to” $G(p) \cong R$ (in some sense). In order to obtain a richer model, however, we will from now on assume that $p = p(n)$, typically in such a way that $p(n) \rightarrow 0$ as $n \rightarrow \infty$. A typical statement for the Erdős-Rényi random graphs will then be of the form: “For a certain $p(n)$, $G(n, p)$ has property ... *asymptotically almost surely* or *with high probability* (a.a.s. or w.h.p. for short)”. By this we mean that the probability of $G(n, p)$ having the desired property tends to 1 as $n \rightarrow \infty$. Here are some examples:

Proposition 169. *For fixed $p \in (0, 1)$, we have $\text{diam}(G(n, p)) = 2$ a.a.s.*

²Technically, the original paper by Erdős and Rényi from 1959 deals with the slightly different model $G(n, m)$ where among n labelled vertices, a random set of m edges is chosen. The random graph $G(n, p)$ was originally introduced by Gilbert also in 1959.

Proof. We first show that $\text{diam}(G(n, p)) > 1$ a.a.s. – to this end, note that the diameter of a graph is 1 iff the graph is complete. This happens exactly with probability $p^{\binom{n}{2}} \rightarrow 0$, hence it only suffices to show $\text{diam}(G(n, p)) \leq 2$.

We call a pair of vertices i, j bad if i and j are non-adjacent and do not have a common neighbour. So, we have to show that the probability of there being a bad pair in $G(n, p)$ tends to zero. For any pair of vertices, define the indicator random variable

$$X_{ij} = \begin{cases} 1 & \text{if } i, j \text{ is bad} \\ 0 & \text{otherwise} \end{cases}.$$

Observe that $\mathbf{E}[X_{ij}] = \mathbf{P}[X_{ij} = 1] = (1-p)(1-p^2)^{n-2}$ in $G(n, p)$, and note that furthermore $X := \sum_{1 \leq i < j \leq n} X_{ij}$ counts the number of bad pairs. Since the expectation is linear, we obtain

$$\mathbf{E}[X] = \sum_{1 \leq i < j \leq n} \mathbf{E}[X_{ij}] = \binom{n}{2} (1-p)(1-p^2)^{n-2}$$

and therefore $\mathbf{E}[X] \rightarrow 0$ as $n \rightarrow \infty$. Moreover, X can only attain values in \mathbb{N}_0 , hence $\mathbf{P}[X > 0] \leq \mathbf{E}[X] \rightarrow 0$. \square

What happened in the last line of the proof? We used the following lemma from probability theory:

Lemma 170 (First-moment-method). *If X is an integer-valued non-negative random variable, then $\mathbf{P}[X > 0] \leq \mathbf{E}[X]$. In particular, if $\mathbf{E}[X] = 0$, then also $\mathbf{P}[X > 0] = 0$.*

This lemma is a direct consequence of the Markov inequality from probability theory, stating that a non-negative random variable X satisfies $\mathbf{P}[X \geq a] \leq a^{-1}\mathbf{E}[X]$ for any $a > 0$.

There is also a method involving the second moment of a random variable, and it is named accordingly (this is a consequence of Chebyshev's inequality):

Lemma 171 (Second-moment-method). *If X is an integer-valued non-negative random variable with finite expectation, then*

$$\mathbf{P}[X = 0] \leq \frac{\mathbf{Var}[X]}{\mathbf{E}[X]^2}.$$

These lemmas are useful for a proof strategy called the *probabilistic method*. The idea is as follows: Suppose you want to show the existence of an element with a certain property in a set A . Then it is enough to show that if you sample elements from A randomly (according to a distribution of your choice),

then you obtain an element with the desired property with positive probability. Note that the property in question can be entirely deterministic (in the sense that it is not relying on randomness) and we will see examples of this later. For now, let's return to $G(n, p)$.

Proposition 172. *As $n \rightarrow \infty$, we have³*

$$\mathbf{P}[G(n, p) \text{ contains a } K_3] \rightarrow \begin{cases} 0 & \text{if } p(n) \ll \frac{1}{n} \\ 1 & \text{if } p(n) \gg \frac{1}{n} \end{cases}$$

Proof. Let's number the $\binom{n}{3}$ potential triangles in $G(n, p)$ from 1 to $\binom{n}{3}$ in some arbitrary order, and define the indicator random variable

$$Y_i = \begin{cases} 1 & \text{if triangle } i \text{ is present} \\ 0 & \text{otherwise} \end{cases}$$

Denote by X the number of triangles in $G(n, p)$. In particular, $X = \sum_{i=1}^{\binom{n}{3}} Y_i$. Since any three vertices could potentially span a triangle in $G(n, p)$, each with probability p^3 , we have

$$\mathbf{E}[X] = \sum_{i=1}^{\binom{n}{3}} \mathbf{E}[Y_i] = \binom{n}{3} p^3 = \frac{n(n-1)(n-2)}{6} p^3 \sim \frac{n^3 p^3}{6}.$$

Therefore, for $p(n) \ll \frac{1}{n}$ we get $\mathbf{E}[X] \ll \frac{1}{6}$, so $\mathbf{E}[X] \rightarrow 0$ and hence $\mathbf{P}[X > 0] \rightarrow 0$ by Lemma 170.

For the other claim, let's calculate the variance of X , using the fact that $Y_i^2 = Y_i$:

$$\begin{aligned} \mathbf{E}[X^2] &= \mathbf{E}\left[\left(\sum_i Y_i\right)^2\right] = \sum_i \mathbf{E}[Y_i^2] + \sum_{\substack{i, j \\ i \neq j}} \mathbf{E}[Y_i Y_j] \\ &= \mathbf{E}[X] + \sum_{\substack{i, j \\ i \neq j}} \mathbf{P}[\text{Both triangles } i, j \text{ are present}] \end{aligned}$$

The triangles i and j in the last line could share 0, 1, or 2 vertices, depending on their position relative to each other:

- If they have no vertices in common, then $\mathbf{E}[Y_i Y_j] = p^6$, and this happens in $\binom{n}{3} \binom{n-3}{3}$ -many configurations.
- If they have one vertex in common, then $\mathbf{E}[Y_i Y_j] = p^6$, and this happens in $3 \binom{n}{3} \binom{n-3}{2}$ -many configurations.

³We write $f \ll g$ if $f(x)/g(x) \rightarrow 0$ as $x \rightarrow \infty$, and analogously for $f \gg g$. We also use $f \sim g$ to express $f(x)/g(x) \rightarrow 1$.

- If they have two vertices in common, then they share an edge as well, so $\mathbf{E}[Y_i Y_j] = p^5$, and this happens in $3\binom{n}{3}(n-3)$ -many configurations.

Therefore, continuing the calculation above:

$$\mathbf{E}[X^2] = \mathbf{E}[X] + \binom{n}{3} \left(\binom{n-3}{3} p^6 + 3 \binom{n-3}{2} p^6 + 3(n-3)p^5 \right)$$

Finally, for the variance, we obtain

$$\begin{aligned} \mathbf{Var}[X] &= \mathbf{E}[X^2] - \mathbf{E}[X]^2 \\ &= \binom{n}{3} p^3 + \binom{n}{3} \binom{n-3}{3} p^6 + 3 \binom{n}{3} \binom{n-3}{2} p^6 \\ &\quad + 3(n-3) \binom{n}{3} p^5 - \binom{n}{3}^2 p^6 \\ &\leq \binom{n}{3} p^3 + 3 \binom{n}{3} \binom{n-3}{2} p^6 + 3(n-3) \binom{n}{3} p^5 \\ &\sim \frac{(np)^3}{6} + \frac{(np)^6}{4n} + \frac{(np)^5}{2n}. \end{aligned}$$

Hence, the expression $\mathbf{Var}[X]/\mathbf{E}[X]^2$ from the right-hand side of Lemma 171 is asymptotically less or equal to

$$\frac{\frac{(np)^3}{6} + \frac{(np)^6}{4n} + \frac{(np)^5}{2n}}{\frac{(np)^6}{36}} = \frac{6}{(np)^3} + \frac{9}{n} + \frac{18}{n^2 p}.$$

Now, if $np(n) \rightarrow \infty$ then

$$\mathbf{P}[X = 0] \leq \frac{\mathbf{Var}[X]}{\mathbf{E}[X]^2} \rightarrow 0 \quad \text{for } n \rightarrow \infty$$

using the second-moment-method. □

Definition 173. Let (P) be some property. A function $f(n)$ is called a *critical function of threshold* for (P) if

$$\mathbf{P}[G(n, p) \text{ satisfies } (P)] \rightarrow \begin{cases} 0 & \text{if } p(n) \ll f(n) \\ 1 & \text{if } p(n) \gg f(n) \end{cases}$$

Thus, Proposition 172 can be rephrased as “The critical function for the occurrence of triangles is $1/n$ ”. As it can be shown, if (P) is a *monotone property*, i.e. such that adding edges to a graph with (P) does not yield a graph without (P) , then (P) admits a critical function. Thus, properties such as containing a fixed subgraph H , being connected, being non-planar, or being Hamiltonian all have critical functions (more on that next lecture).

Remark 174. What happens for $p = \frac{c}{n}$ for some $c > 0$? In this case, one needs to refine the argument a bit: If the random variables Y_i were independent, we could write

$$\mathbf{P}[X_0] = \mathbf{P}[Y_i = 0 \ \forall i] = \prod_i \mathbf{P}[Y_i = 0].$$

Sadly, they're not independent because some triangles share edges, so the second equality is wrong. Let $\Delta = \sum \mathbf{P}[Y_i Y_j = 1]$, where the sum runs over all i, j such that Y_i and Y_j are dependent. Then, using Janson's inequality for the upper estimate, one can obtain

$$\prod_i \mathbf{P}[Y_i = 0] \leq \mathbf{P}[X = 0] \leq e^{-\mathbf{E}[X] + \Delta}. \quad (3)$$

Based on our earlier calculations, we have

$$\Delta = 3(n-3) \binom{n}{3} p^5 \sim \frac{c^5}{2n} \rightarrow 0 \quad \text{and} \quad \mathbf{E}[X] \rightarrow \frac{c^3}{6}.$$

On the other hand, we have

$$\prod_i \mathbf{P}[Y_i = 0] = (1 - p^3)^{\binom{n}{3}} = \left(1 - \frac{c^3}{n^3}\right)^{\binom{n}{3}} \sim e^{-c^3/6}.$$

Plugging both results into (3) yields

$$\mathbf{P}[G(n, c/n) \text{ contains a } K_3] \rightarrow 1 - e^{-c^3/6}.$$

In fact, one can show that in $G(n, c/n)$, the number of triangles X is asymptotically a Poisson-distribution with parameter $c^3/6$.


Lecture 19 More on random graphs

Last time, we have seen that monotone properties of graphs exhibit thresholds $f(n)$ for $G(n, p)$, where the property is satisfied a.a.s. for $p(n) \gg f(n)$, and a.a.s. not satisfied for $p(n) \ll f(n)$; and we could determine the threshold function for the containment of triangles using the first- and second moment methods, together with a bit of combinatorics.

By the same methods (though the computations are a little bit more tedious) one can obtain the critical function for the containment of a K_4 :

$$\mathbf{P}[G(n, p) \supseteq K_4] \rightarrow \begin{cases} 0 & \text{if } p(n) \ll n^{-2/3} \\ 1 & \text{if } p(n) \gg n^{-2/3} \end{cases}$$

For comparison, let now H denote the following graph and consider a probability slightly below the threshold for K_4 , say $p(n) = n^{-7/10} \ll n^{-2/3}$:



In a labelled K_n , there are $c(n, H) := 20 \binom{n}{5}$ different copies of H – this is because there are $\binom{n}{5}$ -many ways to pick the 5 vertices, then there are 5 possibilities to select the vertex with degree 1 among the 5 vertices, and 4 possible choices for the vertex of degree 4. For all other vertices in H , the choice among the remaining three vertices doesn't matter. Enumerate these copies from 1 to $c(n, H)$. The probability of any individual copy of H appearing in $G(n, p)$ is p^7 , as H contains 7 edges. Denote by $Y_i, i = 1, \dots, c(n, H)$ the indicator random variable for the event that the i -th copy of H in K_n , and by X the number of copies of H in $G(n, p)$. Then $X = \sum Y_i$, and by linearity of expectation for our choice of p :

$$\mathbf{E}[X] = \sum_{i=1}^{c(n, H)} \mathbf{E}[Y_i] = \sum_{i=1}^{c(n, H)} p^7 = c(n, H)p^7 \sim \frac{1}{6}n^5 p^7 = \frac{1}{6}n^{1/100} \rightarrow \infty.$$

Yet any copy of H contains a copy of K_4 , so

$$\mathbf{P}[G(n, n^{-7/10}) \supseteq H] \leq \mathbf{P}[G(n, n^{-7/10}) \supseteq K_4] \rightarrow 0$$

since $n^{-7/10}$ is below the critical function for K_4 . So the expected number of copies of H goes to infinity, yet the probability for there being even one H tends to zero!

How to make sense of this? Note that $\mathbf{P}[G(n, n^{-7/10}) \supseteq K_4] > 0$ for any fixed n . But given a copy of K_4 in $G(n, p)$, there are $4(n-4)$ ways to connect

an extra edge to it to obtain a copy of H . Thus any copy of K_4 is contained in $\approx np = n^{3/10}$ different copies of H . So, as n increases, even though it becomes more and more unlikely for $G(n, p)$ to contain a K_4 , these fewer copies of K_4 give more and more copies of H , which explains the weird results above. It can be shown that the threshold for containing H is actually at $n^{-2/3}$, and thereby coincides with the threshold for K_4 . Intuitively, this means once again that if you have a high enough p to produce a K_4 , then the probability is also large enough to guarantee the additional edge in H (at least, as $n \rightarrow \infty$).

Let's now consider an arbitrary finite simple graph $G = (V, E)$ with $v_G = |V|$ and $e_G = |E|$. Denote by $\text{Aut}(G)$ the set of isomorphisms $G \rightarrow G$ (such isomorphisms are also called automorphisms, hence the notation). By the same approach as above, one can prove the next lemma.

Lemma 175. *With the notation just introduced, the following holds:*

(i) *There are exactly⁴ $c(n, G) = \binom{n}{v_G} \frac{v_G!}{|\text{Aut}(G)|}$ copies of G in a labelled K_n .*

(ii) *If X_G denotes the number of copies of G in $G(n, p)$, then*

$$\mathbf{E}[X_G] = c(n, G)p^{e_G} \sim \frac{1}{|\text{Aut}(G)|} n^{v_G} p^{e_G} \rightarrow \begin{cases} 0 & \text{if } p \ll n^{-v_G/e_G} \\ \infty & \text{if } p \gg n^{-v_G/e_G} \end{cases}$$

What about $\mathbf{Var}[X_G]$? Turns out, a similar approach as for triangles works for general subgraphs as well!

Lemma 176. *With notation as above, and assuming $e_G > 0$, we have⁵*

$$\mathbf{Var}[X_G] \asymp (1 - p) \max_{H \subseteq G, e_H > 0} \frac{\mathbf{E}[X_G]^2}{\mathbf{E}[X_H]}.$$

The following theorem gives thresholds for the containment of any finite subgraph. To formulate the theorem, set

$$m(G) := \max_{H \subseteq G, v_H > 0} \frac{e_H}{v_H}.$$

Theorem 177 (Bollobás, 1981). *Let G be a finite simple graph with at least one edge. Then*

$$\mathbf{P}[G(n, p) \supseteq G] \rightarrow \begin{cases} 0 & \text{if } p \ll n^{-1/m(G)} \\ 1 & \text{if } p \gg n^{-1/m(G)}. \end{cases}$$

⁴If you're familiar with abstract algebra, then note that $\text{Aut}(G)$ is a subgroup of the symmetric group on v_G elements, hence $|\text{Aut}(G)|$ is a divisor of $v_G!$.

⁵We write $a_n \asymp b_n$ to express that there exist constants $c, C > 0$ such that $cb_n \leq a_n \leq Cb_n$ for all sufficiently large n .

Proof. Assume $p \ll n^{-1/m(G)}$ and denote by H' a subgraph of G maximizing $\frac{e_{H'}}{v_{H'}}$, i.e. $m(G) = \frac{e_{H'}}{v_{H'}}$. Then

$$\mathbf{P}[X_G > 0] \leq \mathbf{P}[X_{H'} > 0] \leq \mathbf{E}[X_{H'}] \rightarrow 0$$

by the first moment method and part (ii) in Lemma 175.

Conversely, if $p \gg n^{-1/m(G)}$ then $p \gg n^{-v_H/e_H}$ for all subgraphs $H \subseteq G$ with $e_H > 0$. Consequently, we have $\mathbf{E}[X_H] \asymp n^{v_H} p^{e_H} \rightarrow \infty$, again by Lemma 175. Now, applying the second moment method and Lemma 176 yields

$$\begin{aligned} P[X_G = 0] &\leq \frac{\mathbf{Var}[X_G]}{\mathbf{E}[X_G]^2} \asymp \frac{1-p}{\mathbf{E}[X_G]^2} \left(\max_{H \subseteq G, e_H > 0} \frac{\mathbf{E}[X_G]^2}{\mathbf{E}[X_H]} \right) \\ &= O\left(\frac{1}{\min_{H \subseteq G, e_H > 0} \mathbf{E}[X_H]} \right), \end{aligned}$$

therefore $P[X_G = 0] \rightarrow 0$. □

The evolution of $G(n, p)$. Let us now shift our viewpoint on $G(n, p)$: If we fix the number of vertices, but increase the probability, it will look as if we would add randomly more edges into the graph. In fact, this can be made rigorous by using a *coupling*: Consider a labelled K_n where every edge e is equipped with an independent random variable X_e that is uniformly distributed on $[0, 1]$. Then for any fixed p , the Erdős-Rényi random graph $G(n, p)$ has the same distribution as the random spanning subgraph of K_n where we only keep edges with weight $\leq p$. Using this coupling, we can thus ensure a version of $G(n, p)$ that is monotonous in p . The same is true if we focus on the asymptotics of $G(n, p(n))$. Thus, for different asymptotic behaviour of $p(n)$ our $G(n, p)$ showcases different behaviour:

- The function n^{-2} is the critical function for the existence of an edge, and for fixed k the functions $n^{-k/(k-1)}$ is the threshold for the existence of a tree on k vertices.
- Similarly to the result from last lecture, for $p(n) \ll 1/n$ there are a.a.s. no cycles in $G(n, p)$, whereas for $p(n) \gg 1/n$ they are asymptotically guaranteed.
- For $p(n) = \frac{c}{n}$ with $0 < c < 1$, we observe many components that are either trees or contain a single cycle. If C_{\max} denotes the number of vertices in the largest component, then we have

$$\frac{C_{\max}}{\log n} \rightarrow \frac{1}{c - 1 - \log c}$$

in probability.⁶

⁶A sequence X_n converges to $a \in \mathbb{R}$ in probability if $\mathbf{P}[|X_n - a| > \varepsilon] \rightarrow 0$ as $n \rightarrow \infty$ pointwise for all $\varepsilon > 0$.

- For $p(n) = \frac{1}{n}$, the largest component has grown to contain about $n^{2/3}$ vertices.
- For $p(n) = \frac{c}{n}$ with $c > 1$, there exists a constant $\zeta_c > 0$ such that for all $\nu \in (1/2, 1)$ there is a $\delta = \delta(c, \nu) > 0$ such that

$$\mathbf{P}[|C_{\max} - \zeta_c n| \geq n^\nu] = O(n^{-\delta}).$$

In other words, the largest component contains a positive proportion of vertices. Moreover, the largest component is unique, as the probability is large enough to guarantee an edge between any two vertex sets of positive proportion. This is also called the *giant component*. The giant component contains enough edges to be non-planar. All other components are trees on $O(\log(n))$ vertices.

- For $p(n) = \frac{\log(n)+c}{n}$ where $c > 0$, all vertices are either isolated or in the giant component. The number of isolated vertices is Poisson-distributed with parameter e^{-c} .
- For $p(n) = c \frac{\log(n)}{n}$ with $c > 1$, the graph $G(n, p)$ is *a.a.s* connected and Hamiltonian.
- The functions $n^{-2/(k-1)}$ are the critical functions for the appearance of K_k .

Observe how in all these cases, the edge density of $G(n, p)$ still tends to zero (so, we say that those graphs are *sparse networks*, as opposed to the graphs we considered in lectures 16 and 17, where the edge density was positive. Many real-world networks tend to be sparse, and have rather specific degree sequences, that differ from the degree sequences obtained from the Erdős-Rényi random graph (where $\deg(v)$ is binomially distributed with parameters $n - 1$ and p).

Let us conclude this course with an application of the probabilistic method to prove a statement about deterministic graphs. Recall Erdős' theorem (Theorem 136) from Lecture 13:

For all positive integers k there exists a graph G with $\chi(G) > k$ and without a cycle of length $\leq k$.

The trick of the proof will be to look at random graphs, and show that both of these properties together hold with positive probability.

Proof. Fix a $k \geq 3$ and an $\varepsilon \in (0, 1/k)$. Set $p = p(n) = n^{\varepsilon-1}$ and denote by X the number of cycles of length at most k in $G(n, p)$. Similarly to how we

found the expected value to the number of triangles in $G(n, p)$ last lecture, one can find the expected value for the number of cycles of length k to be

$$\frac{n!}{(n-k)!} \frac{p^k}{2k}.$$

We therefore obtain

$$\mathbf{E}[X] = \sum_{i=3}^k \frac{n!}{(n-i)!} \frac{p^i}{2i} \leq \frac{1}{2} \sum_{i=3}^k (np)^i \leq \frac{1}{2} (k-2) n^k p^k,$$

because $np = n^\varepsilon > 1$, so $(np)^i \leq (np)^k$. Hence

$$\mathbf{P} \left[X \geq \frac{n}{2} \right] \leq \frac{\mathbf{E}[X]}{n/2} \leq (k-2) n^{k-1} p^k = (k-2) n^{k\varepsilon-1} \rightarrow 0$$

for $n \rightarrow \infty$ by Markov's inequality (which states $\mathbf{P}[X \geq a] \leq a^{-1} \mathbf{E}[X]$).

We also claim that for any p with $p(n) \geq \frac{16k^2}{n}$ for all n large enough, the independence number α satisfies

$$\mathbf{P} \left[\alpha(G(n, p)) \geq \frac{n}{2k} \right] \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

(The proof of this claim is an exercise for L20). In particular, this claim holds for our choice of $p(n) = n^{\varepsilon-1}$.

By what we've shown so far, for all n large enough we have both

$$\mathbf{P} \left[X < \frac{n}{2} \right] > \frac{1}{2} \quad \text{and} \quad \mathbf{P} \left[\alpha(G(n, p)) < \frac{n}{2k} \right] > \frac{1}{2}.$$

But this implies that both events cannot be disjoint; in particular, $G(n, p)$ contains neither more than $\frac{n}{2}$ cycles of length $\leq k$ nor an independent set of size $\geq \frac{n}{2k}$ with some positive probability. So, consider a graph for which both of these properties are true, and remove one vertex from each cycle of length $\leq k$, thus removing all these unwanted short cycles. The remaining graph H is induced on at least $\frac{n}{2}$ vertices, hence

$$\chi(H) \geq \frac{|V(H)|}{\alpha(H)} \geq \frac{n/2}{\alpha(G)} \geq k,$$

which finishes the proof. □

Some references

Most of this course followed

- Reinhard Diestel: *Graph Theory*. Graduate texts in mathematics 173. Springer, 5th ed., 2017.

with material also coming from the following other textbooks or lecture notes:

- Norman Biggs: *Algebraic Graph Theory*. Cambridge University Press, 2nd ed., 1993.
- Béla Bollobás: *Graph Theory. An Introductory Course*. Graduate texts in mathematics 63. Springer, 1979.
- Béla Bollobás: *Modern Graph Theory*. Graduate texts in mathematics 184. Springer, 1998.
- Lutz Volkmann: *Graphen an allen Ecken und Kanten*. Lecture notes RWTH Aachen, 2006. (In German)

For lectures 16-19, concerning the regularity lemma, the Rado graph and the results on random graphs, the following sources were additionally used:

- Noga Alon and Joel Spencer: *The Probabilistic Method*. Wiley, 3rd ed., 2008.
- Peter J. Cameron: The random graph. arXiv:1301:7544
- Svante Janson, Tomasz Łuczak, Andrzej Ruciński: *Random Graphs*. Wiley, 2000.
- János Komlós and Miklós Simonovits: *Szemerédi's regularity lemma and its applications in graph theory*. DIMACS Technical Report, 1996.
- László Lovász: *Large Networks and Graph Limits*. American Mathematical Soc., 2012.
- Joel Spencer: *The Strange Logic of Random Graphs*. Algorithms and Combinatorics 22. Springer, 2001.
- Remco van der Hofstadt: *Random Graphs and Complex Networks, Vol. I*. Cambridge University Press, 2016.
- Remco van der Hofstadt: *Random Graphs and Complex Networks, Vol. II*. www.win.tue.nl/~rhofstad/NotesRGCNII.pdf, November 16, 2020.