Graph Theory – Lecture Notes

EPFL, Fall Semester 2024/2025

LECTURER: Mathieu Vérité
TEACHING ASSISTANT: Jiaye Wei

Acknowledgements: These notes are largely based on the lecture notes of the *Graph Theory* courses given by Friedrich Eisenbrand and István Tomon at EPFL, relying itself on material from Frank de Zeeuw, Andrey Kupavskii and Dániel Korándi.

Table of Content

Lecture 1
Introduction.

1.1 Definitions

Definition 1 (Graph)**.** *A* graph $G = (V, E)$ *consists of a finite set* V *and a set* E *of two-element subsets of* V*. The elements of* V *are called* vertices *and the elements of* E *are called* edges.

For instance, very formally we can introduce a graph like this:

$$V = \{1, 2, 3, 4\}, \quad E = \{\{1, 2\}, \{3, 4\}, \{2, 3\}, \{2, 4\}\}.$$

In practice we more often think of a drawing like this:

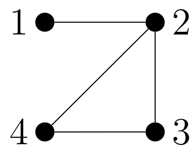

Technically, this is what is called a *labelled graph*, but we often omit the labels. When we say something about an unlabelled graph like , we mean that the statement holds for any labelling of the vertices.

Here are two examples of related objects that we do not consider graphs in this course:

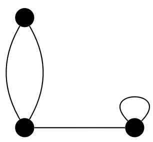

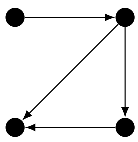

The first is a *multigraph*, which can have multiple edges and loops; the corresponding definition would allow the edge set and the edges to be multisets. The second is a *directed graph*, in which every edge has a direction; in the corresponding definition the edges would be ordered pairs instead of two-element subsets.

Although this course is mostly not about these variants, in some cases it will be more natural to state our results for directed or multigraphs. In any case, we will not treat *infinite graphs* in this course.

Graphs (and their above-mentioned variants) are highly applicable in- and outside mathematics because they provide a simple way of modeling many concepts involving connections between objects. For example, graphs can model social networks (vertices=people & edges=friendships), computer networks (computers & links), molecules (atoms & bonds) and many other things. The aim of this course is to study graphs in the abstract sense, and to introduce the fundamental concepts, tools, tricks and results about them.

Some notation: Given a graph G, we write $V(G)$ for the vertex set, and $E(G)$ for the edge set. For an edge $\{x, y\} \in E(G)$, we usually write xy, and we consider yx to be the same edge. If $xy \in E(G)$, then we say that $x, y \in V(G)$ are *adjacent* or *connected* or that they are *neighbors*. If $x \in e$, then we say that $x \in V(G)$ and $e \in E(G)$ are *incident*.

Definition 2 (Subgraphs)**.** *Two graphs G, G' are* isomorphic *if there is a bijection $\varphi : V(G) \to V(G')$ such that $xy \in E(G)$ if and only if $\varphi(x)\varphi(y) \in E(G')$. A graph H is a* subgraph *of a graph G, denoted $H \subset G$, if there is a graph H' isomorphic to H such that $V(H') \subset V(G)$ and $E(H') \subset E(G)$.*

With this definition we can for instance say that  is a subgraph of . As mentioned above, when we talk about graphs we often omit the labels of the vertices. A more formal way of doing this is to define an *unlabelled graph* to be an isomorphism class of labelled graphs. We will be somewhat informal about this distinction, since it rarely leads to confusion.

Definition 3 (Degree)**.** *Fix a graph $G = (V, E)$. For $v \in V$, we write*

$$N(v) = \{w \in V : vw \in E\}$$

for the set of neighbors of v (which does not include v). Then $d(v) = |N(v)|$ is the degree *of v. We write $\delta(G)$ for the minimum degree of a vertex in G, and $\Delta(G)$ for the maximum degree.*

Definition 4 (Examples)**.** *The following are some of the most common types of graphs.*

- Paths *are the graphs P_n of the form* *. The graph P_n has $n-1$ edges and n different vertices; we say that P_n has* length *$n-1$.*
- Cycles *are the graphs C_n of the form* 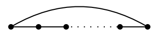 *. The graph C_n has n edges and n different vertices; the* length *of C_n is defined to be n.*
- Complete graphs *(or* cliques*) are the graphs K_n on n vertices in which all vertices are adjacent. The graph K_n has $\binom{n}{2}$ edges. For instance, K_4 is*  .
- *The* complete bipartite graphs *are the graphs $K_{s,t}$ with a partition $V(K_{s,t}) = X \cup Y$ with $|X| = s, |Y| = t$, such that every vertex of X is adjacent to every vertex of Y, and there are no edges inside X or Y. Then K_{st} has st edges. For example, $K_{2,3}$ is*  .

The following are the most common properties of graphs that we will consider.

Definition 5 (Bipartite)**.** *A graph G is* bipartite *if there is a partition $V(G) = X \cup Y$ such that every edge of G has one vertex in X and one in Y; we call such a partition a* bipartition.

Definition 6 (Connected)**.** *A graph G is* connected *if for all $x, y \in V(G)$ there is a path in G from x to y (more formally, there is a path P_k which is a subgraph of G and whose endpoints are x and y).*

A connected component *of G is a maximal connected subgraph of G (i.e., a connected subgraph that is not contained in any larger connected subgraph). The connected components of G form a partition of $V(G)$.*

1.2 Basic Results

In this section we prove some basic facts about graphs. It is a somewhat arbitrary collection of statements, but we introduce them here to get used to the terminology and to see some typical proof techniques.

Proposition 1.1. *In any graph G we have* $\sum_{v \in V(G)} d(v) = 2|E(G)|$.

Proof. We count the number of pairs $(v, e) \in V(G) \times E(G)$ such that $v \in e$, in two different ways. On the one hand, a vertex v is involved in $d(v)$ such pairs, so the total number of such pairs is $\sum_{v \in V(G)} d(v)$. On the other hand, every edge is involved in two such pairs, so the number of pairs must equal $2|E(G)|$. □

What we used here is a very powerful proof technique in combinatorics, called *double counting.* The lemma itself is sometimes called the "handshake lemma" because it says that at a party the number of shaken hands is twice the number of handshakes. It has useful corollaries, such as the fact that the number of odd-degree vertices in a graph must be even.

Next, we will describe a very important characterization of bipartite graphs. But first, we need two more definitions.

Definition 7 (Walk)**.** *A* walk *is a sequence $v_1e_1v_2e_2 \dots v_k$ of (not necessarily distinct) vertices v_i and edges e_i such that $e_i = v_iv_{i+1}$. A* closed walk *is a walk with $v_1 = v_k$. The length of this walk is the number of edges, $k-1$.*

It is easy to see that paths are exactly walks with no repeating vertices, and cycles are exactly closed walks with no repeating vertices apart from $v_1 = v_k$.

Definition 8 (Distance)**.** *The* distance *$d(u, v)$ of two vertices $u, v \in V(G)$ is the length of the shortest path (or walk) in G from u to v. (If there is no u-v path in G then $d(u, v) = \infty$.)*

Now we are ready to prove the characterization. Note that "contains a cycle" means that the graph has a subgraph that is isomorphic to some C_n, and similarly for paths. An "odd cycle" is just a cycle whose length is odd.

Theorem 1.2. *A graph is bipartite if and only if it contains no odd cycle.*

To prove the reciprocal of this statement, we need the result from the following lemma that we prove using an *inductive argument.*

Lemma 1.3. *Every closed walk of odd length contains an odd cycle.*

Proof. We apply induction on the length k of the walk. Since there is no closed walk of length 1, the statement is vacuously true for $k = 1$. We could use this as the base case, but it is also easy to see that the only closed walk of length 3 is the triangle (K_3), which itself is an odd cycle.

Now suppose the statement is true for every odd length $< k$, and let $W = v_1e_1v_2 \dots v_{k+1}$ be a closed walk of odd length k. Let j be the smallest index such that $v_i = v_j$ for some $i < j$. We have two cases. If $j - i$ is even, then deleting the $j - i$ edges $e_i, \dots, e_{j-1}$ from W yields another closed walk $W' \subseteq W$ of odd length. Applying induction on W' then gives an odd cycle in W' and hence in W.

On the other hand, if $j - i$ is odd (and it cannot be 1, so $j - i \geq 3$), then the $j - i$ edges $e_i, \dots, e_{j-1}$ form an odd cycle. Indeed, they form an odd walk without repeated vertices by the choice of v_j. This is what we were looking for. □

We can now provide the full proof of the theorem:

Proof of Theorem 1.2. To prove the easy direction of the statement, the implication, suppose that G is bipartite with bipartition $V(G) = X \cup Y$, and let $v_1 \cdots v_kv_1$ be a cycle in G with, say, $v_1 \in X$. We must have $v_i \in X$ for all odd i and $v_i \in Y$ for all even i. Since v_k is adjacent to v_1, it must be in Y, so k must be even and the cycle is not odd.

Now for the other direction, the reciprocal, suppose G has no odd cycles. We may assume that G is connected. Indeed, otherwise we can apply the same argument to each connected component G_i of G to get a bipartition $X_i \cup Y_i$ of G_i. Choosing $X = \cup_i X_i$ and $Y = \cup_i Y_i$ will then give a bipartition of G.

So if v is a fixed vertex, then every other vertex $u \in V(G)$ has finite distance from v. Let

$$X = \{u : \text{distance of } v \text{ and } u \text{ is } \textit{even}\}$$
$$Y = \{u : \text{distance of } v \text{ and } u \text{ is } \textit{odd}\}.$$

Our aim is to prove that this is a bipartition of G. For this, we need to check that no two vertices in X are adjacent and no two vertices in Y are adjacent.

Suppose for contradiction that some two vertices $u_1, u_2 \in X$ are adjacent, and let e be the edge u_1u_2. By construction, there are paths P_1 from v to u_1 and P_2 from u_2 to v that both have even lengths. But then joining P_1, P_2 and the edge e gives a closed walk P_1eP_2 of odd length, so by Lemma 1.3, G contains an odd cycle, as well, contradiction the assumption. (Note that P_1eP_2 is not necessarily a cycle because P_1 and P_2 might intersect!)

We can do the same to show that no two vertices $u_1, u_2 \in Y$ are adjacent: here the paths P_1, P_2 will both have odd lengths, so again P_1eP_2 is a closed odd walk. So $X \cup Y$ is indeed a bipartition of G. □

The proof above is a *constructive argument*, where we explicitly constructed the object we were looking for (and then proved that it satisfies the required properties)

We close up with a simple but important property regarding walks. Its proof is extremal.

Proposition 1.4. *Every u-v walk W contains a u-v path.*

Proof. Let $v_1v_2 \dots v_k$ be a shortest u-v walk in W (more precisely, in the graph defined by the edges of W), so $u = v_1$ and $v = v_k$. We claim that this walk is in fact a path. Indeed, if $v_i = v_j$ for some $i < j$, then $v_1v_2 \dots v_iv_{j+1} \dots v_k$ is also a u-v walk, and it is shorter (has fewer edges), which is not possible. So the shortest walk has no repeated vertices, i.e., it is a path. □

This fact has the useful corollaries that we can replace paths with walks in some of our definitions:

- The distance $d(u, v)$ is equal to the length of the shortest u-v walk.
- A graph is connected if and only if every pair of vertices u, v is connected by a walk.

The latter connectivity property is sometimes easier to check. Also, it clearly implies that connectivity is an equivalence relation.

Lecture 2

Basic results. Trees.

2.1 More basic results

The next theorem shows that if a graph has many edges, then it must contain a long path.

Theorem 2.1. *Let $k \geq 2$ be an integer and let G be a graph on n vertices with at least $(k-1)n$ edges. Then G contains a path of length k.*

To prove this theorem, we use the combination of the following two lemmas, which are interesting on their own.

First, we consider the special case when every degree of G is large.

Lemma 2.2. *If the minimum degree of G is k, then G contains a path of length k.*

Proof. Let $v_1 \cdots v_l$ be a maximal path in G, i.e., a path that cannot be extended. Then any neighbor of v_1 must be on the path, since otherwise we could extend it. Since v_1 has at least k neighbors, the set $\{v_2, \ldots, v_l\}$ must contain at least k elements. Hence $l \geq k+1$, so the path has length at least k. □

Note that in general this bound cannot be improved, because the complete graph K_{k+1} has minimum degree k, but its longest path has length k. An analogous statement exists for cycles.

The next statement shows that every graph with sufficiently many edges must contain a subgraph with large minimum degree. We give an inductive proof, although it could also be proved using an algorithmic argument.

Lemma 2.3. *Let G be graph on n vertices with at least $(k-1)n$ edges. Then G contains a subgraph H with $\delta(H) \geq k$.*

Proof. Let $k \geq 2$, we proceed by induction on n. Note that $n \leq 2(k-1)$ is impossible because such a graph has at most $n(n-1)/2 < n(k-1)$ edges. Also, for $n = 2k-1$ the only graph with $n(k-1)$ edges is $K_n = K_{2k-1}$, which has minimum degree $2k-2$, so we can take $H = G$.

Now suppose $n > 2k-1$. If $\delta(G) \geq k$, then we can take $H = G$. Otherwise, there is a vertex v of degree $d(v) \leq k-1$. Let $G' = G - v$ be the subgraph we obtain from G by deleting v and all the edges touching it. Then G' has $n-1$ vertices and at least $n(k-1)-(k-1) = (n-1)(k-1)$ edges, so by induction, there is a subgraph $H \subseteq G' \subseteq G$ such that $\delta(H) \geq k$. □

Proof of Theorem 2.1. We can now prove the original theorem. Let G a graph on n vertices with at least $(k-1)n$ edges, with $k \geq 2$. By Lemma 2.3, G contains a subgraph H with minimum degree at least k, and H contains a path of length k by Lemma 2.2. Therefore G contains a path of length k as well. □

The following lemma can be helpful when trying to prove certain statements for general graphs that are easier to prove for bipartite graphs. The lemma says that you don't have to remove more than half the edges of a graph to make it bipartite. The proof is an example of an *algorithmic proof*, where we prove the existence of an object by giving an algorithm that constructs such an object.

Proposition 2.4. *Any graph G contains a bipartite subgraph H with $|E(H)| \geq |E(G)|/2$.*

Proof. We prove the stronger claim that G has a bipartite subgraph H with $V(H) = V(G)$ and $d_H(v) \geq d_G(v)/2$ for all $v \in V(G)$. Starting with an arbitrary partition $V(G) = X \cup Y$ (which need not be a bipartition for G), we apply the following procedure. We refer to X and Y as "parts". For any $v \in V(G)$, we see if it has more edges to X or to Y; if it has more edges that connect it to the part it is in than it has edges to the other part, then we move it to the other part. We repeat this until there are no more vertices v that should be moved.

There are at most $|V(G)|$ consecutive steps in which no vertex is moved, since if none of the vertices can be moved, then we are done. When we move a vertex from one part to the other, we increase the number of edges between X and Y (note that a vertex may move back and forth between X and Y, but still the total number of edges between X and Y increases in every step). It follows that this procedure terminates, since there are only finitely many edges in the graph. When it has terminated, every vertex in X has at least half its edges going to Y, and similarly every vertex in Y has at least half its edges going to X. Thus the graph H with $V(H) = V(G)$ and $E(H) = \{xy \in E(G) : x \in X, y \in Y\}$ has the claimed property that $d_H(v) \geq d_G(v)/2$ for all $v \in V(G)$. □

Our last basic result gives a connection between the number of edges and the number of connected components in a graph.

Proposition 2.5. *If a graph G has n vertices and k edges, then it has at least $n-k$ components.*

Proof. Let us start with the empty graph and add the edges of G to it one-by-one. At the beginning there are n vertices and no edges, so we have n components. Each added edge touches at most 2 of the components, and joins these components if they are different (an edge within a component does not affect any components). This means that adding an edge decreases the number of components by at most 1. Adding k edges therefore decreases the number of components by at most k, so after adding all k edges of G, we are left with at least $n-k$ of them. □

Corollary 2.6. *Every connected graph on n vertices has at least $n-1$ edges.*

2.2 Trees

Definition 9. *A* tree *is a connected graph without cycles. A* forest *is a graph without cycles. In a tree or a forest, a vertex of degree one is called a* leaf.

Lemma 2.7. *Every tree with at least two vertices has two leaves.*

Proof. Consider a longest path $v_0v_1 \ldots v_k$ in the tree (so $k \geq 1$, since the tree has at least two vertices). A neighbor of v_0 cannot be outside the path, since then the path could be extended. But if v_0 were adjacent to v_i for some $i > 1$, then $v_0v_1 \cdots v_iv_0$ would be a cycle. So the only neighbor of v_0 is v_1, and thus v_0 is a leaf. The same argument shows that v_k is also a leaf. □

Theorem 2.8. *Any tree T on n vertices has $n-1$ edges.*

Proof. We use induction on the number of vertices. If $n = 1$, then we have 0 edges. Otherwise, by Lemma 2.7 there exists a leaf v in T. Let $T' = T - v$ be the graph obtained by removing v and its only edge. Then T' is connected, since for any $x, y \in V(T')$ there is a path from x to y in T, and this path cannot pass through v, so it is also a path in T'. Since T has no cycles, neither does T', so T' is a tree, on $n-1$ vertices. By induction T' has $n-2$ edges. Besides $|E(T')| = |E(T)| - 1$ Therefore T has $n-1$ edges. □

Theorem 2.9. *A graph G is a tree if and only if for all $u, v \in V(G)$ there is a unique path from u to v.*

Proof. First suppose we have a graph G in which any two vertices are connected by a unique path. Then G is certainly connected. Moreover, if G contained a cycle $v_1 \cdots v_k v_1$, then $v_1 v_k$ and $v_1 v_2 \cdots v_k$ would be two distinct paths between v_1 and v_k. Hence G is a tree.

Suppose now that G is a tree. Let $u, v \in V(G)$. Since G is connected, there is at least one path from u to v. Suppose there are two distinct paths P, P' from u to v. If these paths only intersect at u and v, we can immediately combine them into a cycle, but in general the paths could intersect in a complicated way, so we have to be careful. The paths P and P' could start out from u being the same; let x be the first vertex that they leave at different edges (so their next vertices are different). Let y be the first vertex of P after x that is also contained in P'. Then there is a cycle in G that goes along P from x to y, and then back along P' from y to x. This is a contradiction, so there is a unique path from u to v in G. □

Lecture 3

BFS. Euler tours. Hamilton cycles.

3.1 Breadth-first search

Definition 10. *A* spanning tree *of a graph G is a subgraph $T \subseteq G$, which is a tree with $V(T) = V(G)$.*

Theorem 3.1. *Every connected graph has a spanning tree.*

To prove the theorem, we provide an algorithm called *Breadth First Search (BFS)* that finds a spanning tree, the breadth first search tree, with some special properties. This algorithm will even let us answer various natural questions, like: what is the distance between two vertices, how many connected components exist, etc..

The algorithm starts with a vertex r that we refer to as the *root* of the spanning tree, and it gradually extends the tree T by examining the edges of G leaving the current T. We use $\partial(T)$ to denote this set of "leaving" edges. More precisely, given a graph G and a subgraph $H \subset G$, we define $\partial(H)$ as the set of edges connecting vertices in H to vertices not in H:

$$\partial(H) = \{xy \in E(G) : x \in V(H), y \notin V(H)\}$$

Additionally, for $x \in H$ and $y \notin H$, we use the notation $H + xy$ to denote the graph of H combined with the edge xy: $V(H + xy) = V(H) \cup \{y\}$ and $E(H + xy) = E(H) \cup \{xy\}$.

Input: a connected graph $G = (V, E)$, a root $r \in V(G)$
Output: the spanning tree T of G
1 initialize the tree, $T = (\{r\}, \emptyset)$, and the "leaving" edges, $\partial(T) = \{ry \mid y \in E(G)\}$
2 **while** $V(T) \neq V(G)$ **do**
3   **for** *each edge* $xy \in \partial(T)$ **do**
4   **if** $T + xy$ *has no cycle* **then**
5   add the edge xy to the current spanning tree: $T \leftarrow T + xy$
6   update the set of "leaving" edges $\partial(T)$

Algorithm 1: Bread-First-Search (BFS)

Proof of Theorem 3.1. To prove the theorem, we will prove that the BFS algorithm provides a spanning tree for any input connected graph G.

Termination: The `for` loop starting on line 3 has to terminate because it considers all edges in $\partial(T)$, a subset of $E(G)$ which is finite. The `while` loop starting on line 2 terminates when all vertices are in T. Let us consider an iteration k of this loop; we denote using the superscript (k) the values at the beginning of each iteration. If the condition on line 4 is never met, it means that each edge $xy \in \partial(T)^{(k)}$ is such that $T^{(k)} + xy$ has a cycle, while $T^{(k)}$ has no cycle. Therefore there are two distinct x-y paths in $T^{(k)} + xy$ and hence there is a x-y path in $T^{(k)}$. For all $xy \in \partial(T)^{(k)}$ we thus have $y \in T^{(k)}$. Then all vertices are in $T^{(k)}$, $V(T^{(k)}) = V(G)$, and the algorithm terminates. On the other hand, if the condition is met, i.e. there exists $xy \in \partial(T)^{(k)}$ such that $T^{(k)} + xy$ has no cycle, then the number of vertices in T at the end of the iteration is larger than at the beginning, because at least y has been added: $|V(T^{(k+1)})| > |V(T^{(k)}|$. Since the number of vertices in T is bounded by the number of vertices in G which is finite, it cannot strictly increase indefinitely. Therefore the condition can only be met for a finite number of iterations, then the algorithm terminates.

Correctness: Let us now prove that the algorithm actually provides a spanning tree as output. We prove that the output subgraph T is a tree by induction on the iterations of the `while` loop. At the first iteration, $T = (\{r\}, \emptyset)$, hence T is connected and acyclic, it is therefore a tree. Let us now consider iteration $k > 1$ and assume that $T^{(k)}$ is a tree. At iteration k, there exists an edge such that the condition on line 4 is met, otherwise the algorithm would have stopped at the previous iteration. For each such edge $xy \in \partial(T^{(k)})$, $T^{(k)} + xy$ is connected because $x \in T^{(k)}$, and it is acyclic according to the condition in the `for` loop. Therefore $T^{(k)} + xy$ is a tree and consequently so is $T^{(k+1)}$ at the end of iteration k. Besides, the algorithm terminates once $V(T) = V(G)$, hence the output tree is a spanning tree of G.

□

The BFS algorithm not only finds a spanning tree of G but also finds a shortest path from its root r to any other vertex.

Proposition 3.2. *Let G a connected graph, then the BFS algorithm applied from root $r \in V(G)$ gives a spanning tree T that contains an r-v path of length $d(r,v)$ for every vertex $v \in V(G)$.*

Proof. The proof is a combination of the following two claims.

Claim 1: if $d(r,v) = k$, then v is added to the BFS tree T in the first k iterations of the `while` loop (line 2).

We prove the claim by induction on k. If $k = 0$, then $r = v$ and the statement is trivial. Otherwise, let $d(r,v) = k \geq 1$ and let us consider a shortest r-v path $rv_1 \dots v_{k-1}v$. Then $d(r, v_{k-1}) \leq k-1$ because of the existence of the path $rv_1 \dots v_{k-1}$, so v_{k-1} is added during the first $k-1$ iterations. If v is also added during these iterations, then we are done. Otherwise, in the k^{th} iteration $v_{k-1}v \in \partial(T)$. Then v is added to T during the k^{th} iteration (via $v_{k-1}v$ or some other edge belonging to another shortest r-v path).

Claim 2: if v is added to the BFS T during the k^{th} iteration of the `while` loop 2, then T contains an r-v path of length at most k.

Again, we prove the claim by induction on k. If $k = 0$, then $r = v$ and the statement is trivial. Otherwise, suppose v is added during the k^{th} iteration via the edge uv. Then u is added during the $(k-1)^{\text{th}}$ iteration at the latest (actually, exactly then), so by induction, T contains an r-u path of length at most $k-1$. Adding uv to this path gives an r-v walk (actually path) of length at most k in T. By Proposition 1.4, this walk contains a r-v path, also of length at most k.

Given these claims, we see that if $d(v,r) = k$, then T contains an r-v path of length at most k. Of course there is no shorter path in G, so this path has length exactly k, as needed.

□

Definition 11. *The* diameter *of a graph G is the largest distance among any pairs of vertices: $diam(G) = \max_{x,y \in V(G)} d(x,y)$. If G is disconnected, then $diam(G) = +\infty$.*

BFS algorithm has a number of applications; it provides a simple solution to perform certain tasks on graphs. For example:

- *Find a shortest path from u to v in G:* Run the algorithm with root u to get a tree T. The unique path from u to v in T (which exists by Theorem 2.9) is a shortest path.
- *Find the connected components of G:* Run the algorithm with some root r. The vertices explored by BFS are exactly the connected component of r. If there is an unexplored vertex r', run BFS again from r' as a root. Repeat until all vertices are visited. This actually gives a spanning forest of G.
- *Compute diam(G):* For each pair of vertices, we can find a shortest path and thus the distance. Do this for all pairs and take the largest distance.
- *Find a shortest cycle in G:* For every edge uv, find a shortest path between u and v in $G - uv$ (if it exists), then combine this path with uv to get a cycle. This will be a shortest cycle through uv. Compare all these cycles to find the shortest.
- *Determine if G is bipartite:* Determine the connected components of G. In every component H, select a root r, and partition the vertices into $X = \{x \in V(H) : d(r, x) \text{ is even}\}$ and $Y = \{y \in V(H) : d(r, y) \text{ is odd}\}$. Then H is bipartite if and only if X and Y have no internal edges (see the proof of Theorem 1.2), and G is bipartite if and only if every component is bipartite.

Not all of these algorithms are the most efficient, there are all kinds of algorithms that do these tasks faster. However they are already much better than brute force approaches that go over all possible answers.

We should emphasize in particular, that for finding shortest paths between two vertices, BFS is the best general algorithm. Indeed, *Dijkstra's algorithm*, a variant of BFS that works for graphs with positive edge weights (representing, e.g. the lengths of streets) is directly used by routing softwares.

3.2 Euler tours

Definition 12. *A* trail *is a walk with no repeated edges. A* tour *is a closed trail, that is one that starts and ends at the same vertex.*

Reminder: multigraphs are graphs where there may be several edges connecting the same pair of vertices or edges with endpoints at the same vertex.

Definition 13. *An* Euler (or Eulerian) trail *in a (multi)graph $G = (V, E)$ is a trail in G passing through every edge (exactly once). An* Euler tour *is a tour in G passing through every edge.*

This notion originates from the "seven bridges of Königsberg" problem – the oldest problem in graph theory, originally solved by Euler in 1736 – that asked if it was possible to walk through all the seven bridges of Königsberg in one go without crossing any of them twice. This question can be turned into a graph problem asking for an Euler trail. Euler solved the problem by noticing that the existence of Euler trails is closely related to the degree parities.

Theorem 3.3 (Euler - 1736)**.** *A connected (multi)graph has an Eulerian tour if and only if each vertex has even degree.*

The proof of this theorem is based on the following simple lemma.

Lemma 3.4. *In a graph where all vertices have even degree, every maximal trail is a closed trail.*

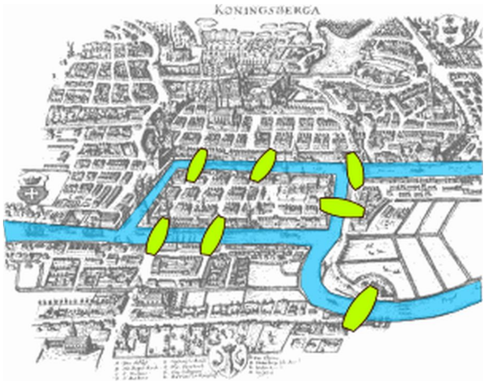

(a) Representation of Königsberg old town. The seven bridges are highlighted in green.

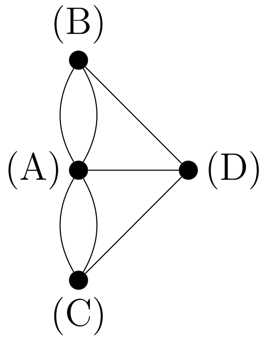

(b) Graph modeling of the problem: each vertex is a part of the town, each edge is a bridge.

Figure 1: The seven bridges of Königsberg problem

Proof. Let T be a maximal trail. If T is not closed, then T has an odd number of edges incident to the final vertex v. However, as v has even degree, there is an edge touching v that is not contained in T. This edge can be used to extend T to a longer trail, contradicting the maximality of T. □

Proof of Theorem 3.3. To see that the condition is necessary, suppose G has an Eulerian tour C. If a vertex v was visited k times in the tour C, then each visit used 2 edges incident to v (one incoming edge and one outgoing edge). Thus, $d(v) = 2k$, which is even.

To see that the condition is sufficient, let G be a connected graph with even degrees. Let $T = e_1e_2 \dots e_\ell$, where $e_i = (v_{i-1}, v_i)$ be a longest trail in G, then it is maximal. According to Lemma 3.4, T is closed, which means $v_0 = v_\ell$. G is connected, so if T does not include all the edges of G then there is an edge e outside of T that touches it, i.e. $e = uv_i$ for some vertex v_i in T, with $u \in V(G)$. Since T is closed, we can start walking through it at any vertex. But if we start at v_i then we can append the edge e at the end. More precisely $T' = e_{i+1} \dots e_\ell e_1 e_2 \dots e_i e$ is a trail in G which is longer than T. It is a contradiction to the fact that T is a longest trail in G. Thus, T must include all the edges of G and so it is an Eulerian tour. □

Corollary 3.5. *A connected (multi)graph G has an Euler trail if and only if it has either* 0 *or* 2 *vertices of odd degree.*

Proof. Suppose T is an Euler trail from vertex u to vertex v. If $u = v$ then T is an Eulerian tour and so by Theorem 3.3, it follows that all the vertices in G have even degree. If $u \neq v$ then let us add a new edge $e = uv$ to G. In this new graph $G \cup \{e\}$, $T \cup \{e\}$ is an Euler tour. By Theorem 3.3 we see that all the degrees in $G \cup \{e\}$ are even. This means that in the original graph G, the vertices u, v are the only ones that have odd degree.

Now we prove the other direction of the corollary. If G has no vertices of odd degree then by Theorem 3.3 it contains an Eulerian tour which is also an Eulerian trail. Suppose now that G has 2 vertices u, v of odd degree. Then add a new edge e to G. Now all vertices of the resulting graph $G \cup \{e\}$ have even degree, so, by Theorem 3.3, it has an Eulerian tour C. Removing the edge e from C gives an Eulerian trail of G from u to v. □

3.3 Hamilton cycles

Euler trails are walks that use each edge exactly once. But what if we want to use each *vertex* exactly once?

Definition 14. *A* Hamilton (or Hamiltonian) cycle *in a graph G is a cycle that contains all vertices of G. A* Hamilton path *in a graph G is a path that contains all vertices of G. A graph G is* Hamiltonian *if it contains a Hamilton cycle.*

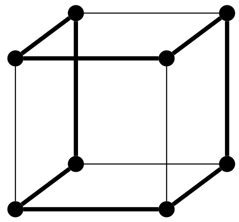

Figure 2: A Hamilton cycle on the graph representing the corners and edges of a cube.

Although the proof of Theorem 3.3 is not really described in an algorithmic way, it can actually be turned into an efficient algorithm. However, no "good" algorithm is known that would find a Hamilton cycle in a graph.

Definition 15. *The* girth *of a graph G is the length of the shortest cycle contained in G. The* circumference *is the length of the longest cycle contained in G.*

For example, the complete graph K_n for $n \geq 3$ has girth 3 and circumference n.

We have seen that computing the girth of a graph can be done using BFS. But finding the circumference is even more difficult than deciding if a graph is Hamiltonian, because a graph has a Hamilton cycle if and only if its circumference is n.

Another related problem is to find a shortest Hamilton cycle in a graph with weighted edges; this is called the *travelling salesman problem (TSP)* and is one of the most famous computationally hard problems, with many real-life applications.

Although we have no general algorithm or recipe for finding Hamilton cycles, we can still prove some theorems that are useful in certain situations. Our first example is a necessary condition for Hamiltonicity.

Given a graph G and a set $S \subset V(G)$ of vertices, we write $G - S$ for the graph obtained by removing the vertices of S from G, along with all the edges touching the vertices in S.

Lemma 3.6. *If G has a Hamilton cycle, then for all $S \subset V(G)$, $G - S$ has at most $|S|$ connected components.*

Proof. The Hamilton cycle must visit all the components of $G - S$ (viewed as subgraphs of G), and to get from one component to another the cycle must pass through a vertex of S. Thus every component is connected to S by two edges of the cycle (and possibly by other edges not in the cycle). Since every vertex is incident to two edges of the cycle, we have that twice the number of components is at most twice the number of vertices of S. □

This lemma can be useful to show that a graph does not have a Hamilton cycle. For example, consider the graph G presented Figure 3a and select S as the set of two vertices in the middle and their connecting edge. Then $G - S$ has three connected components, so by Lemma 3.6 the graph has no Hamilton cycle.

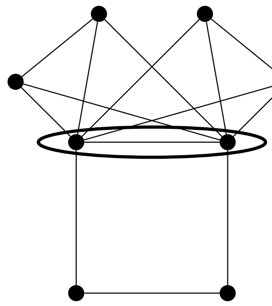

(a) The condition is not met, hence the graph is not Hamiltonian

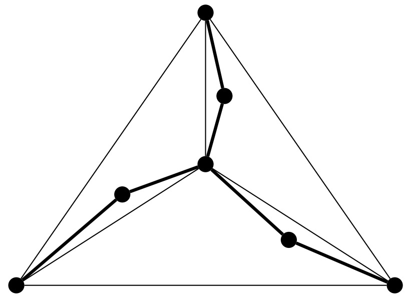

(b) The condition is met but the graph is not Hamiltonian

Figure 3: Illustration of the result of Lemma 3.6

On the other hand, it is important to emphasize that the condition in Lemma 3.6 is not sufficient. For example, one can check that the graph G on Figure 3b satisfies the condition that for all $S \subset V(G)$, $G - S$ has at most $|S|$ components, yet it has no Hamilton cycle. To see the latter, observe that for each vertex of degree 2, both incident edges would need to be in the cycle; but then the middle vertex would be incident to at least three edges of the cycle, which is impossible.

Lecture 4

Hamilton cycles.

In this chapter, we are interested in stating sufficient conditions for Hamiltonicity. Historically several approaches have been used to tackle the problem.

4.1 Number of edges and minimum degree

A natural idea that comes to mind is that a graph with many edges must have a Hamilton cycle. Unfortunately, there is no chance to get a very good bound on the number of edges sufficient to guarantee a Hamilton cycle, because there exist "almost complete" graphs that are not Hamiltonian. Still, the following result holds.

Theorem 4.1. *If G is an n-vertex graph with $|E(G)| > \binom{n-1}{2} + 1$ edges, then G contains a Hamilton cycle.*

Proof. We apply induction on n. The statement is clearly true for $n = 1, 2, 3$, so we assume $n > 3$.

We claim that G has a vertex v of degree at least $n - 2$. Indeed, otherwise every vertex has degree at most $n - 3$, so:

$$|E(G)| = \frac{\sum_{v \in V(G)} d(v)}{2} \leq \frac{n(n-3)}{2} < \frac{(n-1)(n-2)}{2} = \binom{n-1}{2}$$

This contradicts our assumption.

Now let $G' = G - v$, so G' has $n - 1$ vertices. We distinguish the two cases $d(v) = n - 2$ and $d(v) = n - 1$.

<u>Case 1:</u> suppose that $d(v) = n - 2$. Then:

$$|E(G')| = |E(G)| - (n-2) > \binom{n-1}{2} + 1 - (n-2) = \binom{n-2}{2} + 1$$

Hence G' has a Hamilton cycle $\mathcal{C}$ by induction. Since $d(v) = n - 2$ and $n > 3$, v must be adjacent to two consecutive vertices of this cycle, x and y. Then we can remove the edge xy from $\mathcal{C}$ and replace it by xv and vy to get a Hamilton cycle in G.

<u>Case 2:</u> now suppose $d(v) = n - 1$. In this case we only have $|E(G')| > \binom{n-2}{2}$, so we cannot apply induction right away. If G' is complete, then G' has a Hamilton cycle, and we can add v as in the previous case (in fact, as $d(v) = n - 1$, v is now adjacent to every vertex of the cycle).

Otherwise, there is an edge xy missing from G'. Let us look at the graph $G' + xy$ that we get by adding xy to G'. Now this graph has more than $\binom{n-2}{2} + 1$ edges, so we can apply induction to find a Hamilton cycle $\mathcal{C}$ in it. If $\mathcal{C}$ does not contain xy, then we can again add v as in the previous case. If $\mathcal{C}$ does contain xy, then replacing xy with the path xvy in $\mathcal{C}$ gives a Hamilton cycle in G. □

As mentioned before, the condition in the theorem above is somewhat weak in the sense that many graphs that have a Hamilton cycle do not satisfy the condition.

The following sufficient condition does better by looking at the minimum degree instead of the total number of edges. Note that we have previously seen that every graph G has a cycle of length at least $\delta(G) + 1$. The following theorem says that something much stronger is true when the minimum degree is at least $|V(G)|/2$.

Theorem 4.2 (Dirac - 1952)**.** *Let G be a graph on $n \geq 3$ vertices. If $\delta(G) \geq \frac{n}{2}$, then G contains a Hamilton cycle.*

Proof. First observe that G must be connected, because each component contains at least $\delta(G) + 1 > n/2$ vertices, so G cannot have more than one component.

Consider a longest path $\mathcal{P} = v_1 v_2 \dots v_k$ in G. By maximality, all neighbors of v_1 and v_k are in the path. Let us say that an edge $v_i v_{i+1}$ is type-1 if $v_{i+1} \in N(v_1)$, and let us say that it is type-2 if $v_i \in N(v_k)$. As $\delta(G) \geq n/2$, we have at least $n/2$ type-1 and $n/2$ type-2 edges in $\mathcal{P}$. But $\mathcal{P}$ has at most $n-1$ edges, so some edge $v_j v_{j+1}$ is both type-1 and type-2, i.e., $v_1 v_{j+1}$ and $v_j v_k$ are edges of G. Then $\mathcal{C} = \mathcal{P} - v_j v_{j+1} + v_1 v_{j+1} + v_j v_k = v_j \dots v_1 v_{j+1} \dots v_k v_i$ is a cycle.

In fact, $\mathcal{C}$ is a Hamilton cycle. Indeed, suppose not all vertices are contained in $\mathcal{C}$. Since G is connected, there must be an edge uv_i where $u \notin \mathcal{C}$. Then there is a path that goes from u to v_i and then all around the cycle $\mathcal{C}$ to a neighbor of v_i. This path contains $k+1$ vertices, contradicting the maximality of $\mathcal{P}$. □

This theorem is again best possible, in the sense that a weaker bound on the minimum degree would not imply a Hamilton cycle.

It is easy to check that the proof of Dirac's theorem works for the following strengthening, as well:

Theorem 4.3 (Ore - 1960)**.** *Let G be a graph on $n \geq 3$ vertices. If $d(u) + d(v) \geq n$ for any non-adjacent vertices u and v, then G contains a Hamilton cycle.*

Using this theorem, one can also obtain a short proof of Theorem 4.1, so one can think of Ore's theorem as a common generalization of Theorems 4.1 and 4.2.

4.2 Closure and degree sequence

We can observe that it is possible to modify the proof of Dirac's theorem to yield a stronger sufficient condition. The basis of the approach is the following lemma.

Lemma 4.4. *Let G a graph with n vertices and u, v two non-adjacent vertices in G such that $d(u) + d(v) \geq n$. Then G is Hamiltonian if and only if $G + uv$ is Hamiltonian.*

Proof. If G is Hamiltonian, then obviously so is $G + uv$. Let us prove the other implication by contradiction and assume that $G + uv$ is Hamiltonian but G is not. Then there exists a Hamilton cycle in $G + uv$ and therefore a Hamilton path $\mathcal{P}$ in G from u to v. Similarly to the proof of Theorem 4.2, we define E_1 and E_2 the sets of respectively type-1 and type-2 edges.

Since $\mathcal{P}$ contains $n-1$ edges, $|E_1 \cup E_2| < n$. Furthemore we have $E_1 \cap E_2 = \emptyset$, otherwise we could define a Hamilton cycle in G similarly to the proof of Theorem 4.2, which is not possible because G is assumed to be non-Hamiltonian.

Besides, by maximality of the Hamilton path in G, all neighbors of u and v are on $\mathcal{P}$. Therefore $d(u) + d(v) = |E_1| + |E_2| = |E_1 \cup E_2| + |E_1 \cap E_2| < n$, which contradicts the original assumption on G. □

Now this lemma motivates the definition of the *closure* of a graph.

Definition 16. *The* closure *of a graph G is the graph obtained from G by recursively joining pairs of non-adjacent vertices whose degree sum is at least $|V(G)|$ until no such pair remains. The closure of G is denoted $c(G)$.*

Lemma 4.5. *The closure $c(G)$ of a graph G is well defined.*

Proof. We have to prove that the operation "take the closure of", $c(\cdot)$, is a well defined function, i.e. for a given input graph G it yields the same output graph $c(G)$ regardless in which order edges are added to G to form its closure.

Let $G^{(1)}$ and $G^{(2)}$ two distinct graphs obtained from a graph G with n vertices by applying the closure procedure. We will prove that $E(G^{(1)}) = E(G^{(2)})$. We denote $e_1^{(1)}, \ldots, e_k^{(1)}$ and $e_1^{(2)}, \ldots, e_\ell^{(2)}$ the sequences of edges added to G to obtain $G^{(1)}$ and $G^{(2)}$ respectively. Let us assume that there exists $e_{i+1}^{(1)} = uv$ defined as the first edge in the sequence $e_1^{(1)}, \ldots, e_k^{(1)}$ that is not in $G^{(2)}$.

We define $H = G + \{e_1^{(1)}, \ldots, e_i^{(1)}\}$. Then from the definition of $G^{(1)}$ it follows that $d_H(u) + d_H(v) \geq n$. Now H is a subgraph of $G^{(2)}$ and $e_{i+1}^{(1)} \notin E(G^{(2)})$. Therefore $d_{G^{(2)}}(u) + d_{G^{(2)}}(v) \geq n$. It is a contradiction because u and v are not adjacent in $G^{(2)}$ and should therefore have been added during the procedure.

Hence $E(G^{(1)}) \subset E(G^{(2)})$. The inclusion in the other direction can be proven similarly. Therefore $G^{(1)} = G^{(2)}$ and thus $c(G)$ is well defined. □

Using the observation from Lemma 4.4, the following theorem from Bondy and Chvátal establishes the equivalence of the Hamiltonicity between a graph and its closure.

Theorem 4.6 (Bondy, Chvátal - 1974). *A graph G is Hamiltonian if and only if its closure $c(G)$ is Hamiltonian.*

Proof. The theorem is proved directly by induction using Lemma 4.4. □

Corollary 4.7. *Let G be a graph on $n \geq 3$ vertices. If $c(G)$ is complete, then G is Hamiltonian.*

The theorem from Bondy and Chvátal has many interesting consequences. In particular, it can be used to derived various sufficient conditions of Hamiltonicity expressed in terms of the degree sequence of the vertices. One example is the following theorem by Chvátal.

Theorem 4.8 (Chvátal - 1972). *Let G be a graph on $n \geq 3$ vertices such that the degree sequence of its vertices is $d_1 \leq d_2 \leq \cdots \leq d_n$. If for all $1 \leq k < \frac{n}{2}$, $d_k \leq k \Rightarrow d_{n-k} \geq n-k$, then G is Hamiltonian.*

4.3 Independent sets and connectivity

Another approach to the characterization of Hamiltonian graphs is to try link it to global properties of the graph, in particular its *independence number* and its *vertex connectivity.*

Definition 17 (Independent/stable set). *Let G be a graph. A vertex set $I \subseteq V(G)$ is said to be* independent *in G if no two vertices of I are connected by an edge of G. I is called an* independent set *or a* stable. *The* independence number $\alpha(G)$ *is the size of a largest independent set in G.*

Examples

- for the clique on n vertices, the largest independent set is a single vertex: $\alpha(K_n) = 1$;
- the largest independent set of the fully connected bipartite graph $K_{n,m}$ is the largest of its two parts: $\alpha(K_{n,m}) = \max(n, m)$;
- the largest independent set of the empty graph on n vertices is itself: $\alpha(G) = n$.

Theorem 4.9. *Let G be a graph on $n \geq 3$ vertices. If G has at least $\alpha(G)$ vertices of degree $n-1$, then it contains a Hamilton cycle.*

Proof. Let $k = \alpha(G)$, and let us take a longest cycle $\mathcal{C}$ in G. We will prove that $\mathcal{C}$ is a Hamilton cycle. So suppose that $\mathcal{C}$ is not Hamiltonian, and let w be any vertex not contained in the cycle. Then w is not adjacent to both of any two consecutive vertices x, y in $\mathcal{C}$. Indeed, otherwise we could replace the edge xy in $\mathcal{C}$ with the path xwy to get a longer cycle (see Figure 4a). In particular, this observation means that every vertex of degree $n-1$ is in the cycle, and no two of them are adjacent in $\mathcal{C}$.

So let $v_1, \ldots, v_k \in \mathcal{C}$ be k of the vertices of degree $n-1$, and for each $i = 1, \ldots, k$, let u_i be the vertex immediately following v_i on $\mathcal{C}$ in the clockwise direction. The observation above implies that no u_i is a neighbor of w, in particular, the u_i are all different from the v_i.

Now the set $\{w, u_1, \ldots, u_k\}$ has size $k+1 > \alpha(G)$, so by assumption, it cannot be independent. G therefore contains an edge $u_i u_j$. But then we can remove the edges $v_i u_i$ and $v_j u_j$ from $\mathcal{C}$ and replace them by the edges $u_i u_j$, $v_i w$ and $w v_j$ (see Figure 4b). Again, we get a longer cycle, and this contradiction shows that $\mathcal{C}$ must contain all vertices of G. □

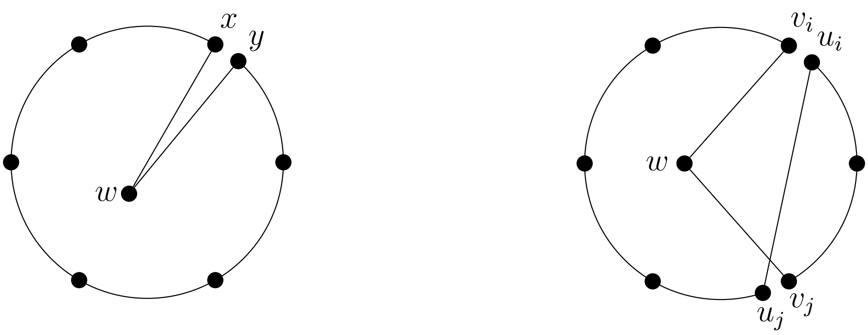

(a) (b)

Definition 18 (Connectivity)**.** *A graph G on n vertices is said to be k-connected, for $k \leq n$, if there exists no subset of $k-1$ vertices which removal disconnects G. The* connectivity *of G, denoted $\kappa(G)$, is the largest k such that G is k-connected.*

Loosely speaking, saying that a graph G is k-connected means that we need to remove at least k vertices from it to create a disconnected subgraph. We do not go into further details about connectivity here because it is an important topic in itself that will be covered in a future lecture. However, we can mention this important theorem stating a sufficient condition of Hamiltonicity in terms of the independence number and connectivity of a graph.

Theorem 4.10 (Chvátal, Erdös - 1972)**.** *Let G be a graph on $n \geq 3$ vertices. If G is k-connected with $\alpha(G) \leq k$, in particular if $\alpha(G) \leq \kappa(G)$, then G is Hamiltonian.*

Proof. The proof of this theorem is very similar to the previous one but require the use of Menger's theorem that we will cover later.

Let us consider $\mathcal{C}$ a longest cycle in a k-connected graph G on $n \geq 3$ vertices and assume that $\alpha(G) \leq k$. We will prove that $\mathcal{C}$ is a Hamilton cycle by contradiction. Now suppose that $\mathcal{C}$ is not a Hamilton cycle, then there exists $w \in V$ not contained in the cycle.

Since G is k-connected, by Menger's theorem there exist k pairwise edge-disjoint paths starting at w and terminating in $\mathcal{C}$ which share with $\mathcal{C}$ only their terminal vertices denoted $v_1, \ldots, v_k$. We define the set of vertices $u_1, \ldots, u_k$ such that each u_i is the direct successor of v_i on $\mathcal{C}$ considered in a specified ordering. Then the remainder of the proof is identical to Theorem 4.9. □

4.4 Covering the vertices with multiple paths

Not every graph has a Hamilton path. For example, if G has k components, then we surely need at least k paths to visit all the vertices of G, and even that is not necessarily possible. The next result shows that $\alpha(G)$ paths are always enough to cover all vertices of G

Theorem 4.11. *Every graph G contains a set of at most $\alpha(G)$ vertex-disjoint paths (i.e. paths that share no vertex in common) that together contain all vertices of G.*

Proof. We proceed by induction on $\alpha(G)$. If $\alpha(G) = 1$ then G is a complete graph, so it has a Hamilton path. So suppose $\alpha(G) > 1$, and let $\mathcal{P} = v_1 \dots v_k$ be a maximal path in G. Our plan is to delete this path and its vertices from G and apply the induction hypothesis to the remaining graph $G' = G - \{v_1, \dots, v_k\}$.

For this we just need to observe that $\alpha(G') < \alpha(G)$. Indeed, let X be an independent set in G' of size $\alpha(G')$. As $\mathcal{P}$ is maximal, all neighbors of v_1 are in $\mathcal{P}$, so in particular, v_1 is not adjacent to any vertex of G'. But then $X \cup \{v_1\}$ is an independent set of size $\alpha(G') + 1$ in G.

Hence we can apply induction to G', and get a set of at most $\alpha(G')$ disjoint paths that cover all vertices of G'. Together with $\mathcal{P}$, this gives a collection of at most $\alpha(G)$ disjoint paths that cover all vertices of G. □

Lecture 5

Planar graphs.

5.1 Drawings of graphs

We have so far only considered *abstract graphs*, although we often used pictures to illustrate the graph. In this lecture, we prove some facts about pictures of graphs and their properties. To set this on a firm footing, we give a formal definition of what we mean by "picture", although in most of the lecture we will be less formal.

Definition 19. *A* drawing *of a graph G consists of an injective map $f : V(G) \to \mathbb{R}^2$, and a curve γ_{xy} (the image of an injective continuous map $[0,1] \to \mathbb{R}^2$) from x to y for every $xy \in E(G)$, such that $f(z) \notin \gamma_{xy}$ for any vertex $z \neq x, y$. A drawing is* planar *if the curves γ_{xy} do not intersect each other except possibly at endpoints.*

Definition 20. *A graph is* planar *if it has a planar drawing.*

To show that a graph is planar, we only have to supply a planar drawing. For example, the graphs K_4 and $K_{2,3}$ are planar graphs. It is often a little harder to show that a graph is not planar.

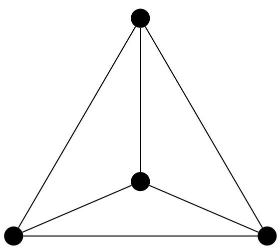

Figure 5: Drawing of K_4

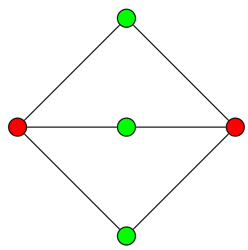

Figure 6: Drawing of $K_{2,3}$

Proposition 5.1. *The graph K_5 is not planar.*

Informal proof. The graph contains a K_3, which can basically be drawn in only one way. If in a drawing the fourth vertex is inside this K_3 and the fifth is outside, then the edge between them must cross the K_3, which means that the drawing is not planar. If both vertices are inside the K_3, then the three edges of one vertex divide the inside of K_3 into three regions. The other vertex must then be in one of these regions, and one vertex of K_3 is outside this region, so again, the edge between these two vertices crosses the boundary of the region. A similar argument applies when both vertices are outside the K_3. □

Note that this proof heavily relied on the fact that an edge connecting a vertex inside of K_3 (or any other cycle) to a vertex outside must cross the K_3 (or cycle). This is a highly non-trivial statement that depends on a strong topological theorem called the **Jordan Curve Theorem**, that states that every non-self-intersecting closed curve in $\mathbb{R}^2$ divides $\mathbb{R}^2$ into an *inside* and an *outside*, and any path between a point inside and a point outside must pass through the closed curve. We remark that this theorem is specific to the plane: it is not always true in some other surfaces, for example, on a torus.

The Jordan Curve Theorem is well beyond the scope of this course. To avoid using such heavy machinery, we could replace continuous curves with *polygonal paths* in the definition of drawings, where a polygonal path is a special curve, whose image is a union of finitely many segments. The Jordan Curve Theorem is much easier to prove for such paths.

Nevertheless, many technicalities arise when adding edges to drawings, even if we use polygonal paths, but we are going to ignore these for clarity. Therefore, **our proofs** concerning planarity will sometimes appeal to intuition, hence they **will not always be completely rigorous**.

5.2 Euler's formula

One important property that distinguishes planar drawings from general graph drawings is that plane drawings have faces.

Definition 21. *The* faces *of a planar drawing D of a graph G are the maximal connected regions of $\mathbb{R}^2 \setminus D$. It is a set of open subsets of $\mathbb{R}^2$ denoted $F_D(G)$.*

Since D is bounded, it has exactly one unbounded face called the *outer face*; all other faces are the *inner* faces. We think of the *boundary* of a face as the closed walk around it. If an edge only bounds one face, then it appears in the boundary walk twice. We can observe that if the graph is disconnected, then the boundary of some faces will not be one closed walk, but a union of closed walks.

Note that the faces of a drawing are usually not determined by the graph, i.e. different drawings might have different boundary walks. However, the following theorem shows that the *number of faces* is always the same.

Theorem 5.2 ("Euler's formula", Euler - 1752)**.** *Let D be a plane drawing of a connected planar graph with n vertices, e edges and f faces. Then:*

$$n - e + f = 2.$$

Proof. We use induction on e. Since the graph is connected, we have $e \geq n - 1$, so we can start with the base case $e = n - 1$. In that case, G contains no cycles (it is a tree), so any planar drawing has only one face and we have $f = 1$. Then: $n - e + f = n - (n - 1) + 1 = 2$.

Now suppose $e \geq n$. Then G contains a cycle $\mathcal{C}$. Let xy be an edge of this cycle. The graph $G \setminus xy$ is connected, and erasing this edge from D gives a planar drawing D' of $G - e$ with n' vertices, e' edges and f' faces. Clearly, $n' = n$ and $e' = e - 1$.

Moreover, by the Jordan Curve Theorem, there is no face both inside and outside $\mathcal{C}$, so xy is on the boundary of two different faces. Removing xy merges these two faces, and does not affect any other faces, so $f' = f - 1$. Now applying the induction assumption to D', we get: $n - e + f = n - (e - 1) + (f - 1) = n' - e' + f' = 2$. □

As we mentioned before, this means that the number of faces is the same in any plane drawing of a connected planar graph, namely $f = 2 + e - n$. Euler's formula allows to prove many properties about planar graphs. In particular one important consequence is that a planar graph cannot be too dense.

Proposition 5.3. *Every planar graph G on $n \geq 3$ vertices has at most $3n - 6$ edges.*

Proof. We may assume that G is connected: otherwise we repeatedly add edges between different components. This does not change planarity, so if the bound holds for this graph, it will also hold for the original with fewer edges.

Now in a drawing of a connected graph, every face F is bounded by a closed walk; let ℓ_F be the length of this walk. As we have observed before, each edge of G is either on the boundary of two faces, or it appears twice in the boundary walk of the one face. Either way, each edge is counted twice in the sum of the ℓ_F. On the other hand, we have $\ell_F \geq 3$ for every face F because G is connected on at least 3 vertices. Hence:

$$3f \leq \sum_{F \text{ face}} \ell_F = 2e$$

G is connected, so we can apply Euler's formula. Together with the inequality $f \leq 2e/3$, it gives $2 = n - e + f \leq n - e + \frac{2e}{3} = n - \frac{e}{3}$. which rearranges to $e \leq 3n - 6$, as needed. □

Proposition 5.4. *Let G be a planar graph on $n \geq 3$ vertices without triangle, then it has at most $2n - 4$ edges. It is the case in particular for bipartite graphs.*

Proof. Since G does not contain any triangle, the length of a closed walk boundary of face F is such that $\ell_F \geq 4$. Therefore, using a similar argument as in the previous proof, we have $4f \leq \sum_{F \text{ face}} \ell_F = 2e$. Using Euler's formula, we thus have $2 = n - e + f \leq n - e + \frac{e}{2} = n - \frac{e}{2}$ which yields the expected inequality: $e \leq 2n - 4$. □

Proposition 5.5. *If G is planar, then it has a vertex of degree at most five.*

Proof. Suppose not. Then every vertex has degree at least 6, so the number of edges in G is at least $e = \frac{1}{2}\sum d(v) \geq 3n$. Now according to the reasoning in the proof of Proposition 5.3, we have $f \leq 2e/3$. Pluging those inequalities into Euler's formula yields: $2 = n - e + f \leq \frac{e}{3} - e + \frac{2e}{3} = 0$, which is a contradiction. □

Those results gives us easy to check necessary conditions for a graph to be planar, for example we now have a much simpler proof of Proposition 5.1.

Corollary 5.6. *K_5 and $K_{3,3}$ are not planar.*

Proof.

- K_5 has 5 vertices and 10 edges, but $10 > 3 \cdot 5 - 6 = 9$, so it cannot be planar by application of Proposition 5.3.
- Similarly, $K_{3,3}$ is bipartite, with 6 vertices and 9 edges, but $9 > 2 \cdot 6 - 4 = 8$, so Proposition 5.4 shows it is not planar.

□

If a graph H is a subgraph of a graph G and H is not planar, then G is also not planar, since a planar drawing of G would give a planar drawing of H. A stronger version of this observation can be stated using graph *subdivisions*.

Definition 22. *The graph H is a* subdivision *of graph G if it can be obtained from G by repeatedly replacing an edge $xy \in E(G)$ by a path xzy for some new vertex z. Informally, we simply place the new vertex z somewhere on top of the edge xy, subdividing it into two edges xz and zy.*

It is easy to see that a subdividion of G is planar if and only if G is planar because the new vertices do not have any effect on whether ot not we can draw the graph without crossing edges. So if a graph contains a subdivision of K_5 or $K_{3,3}$, then it is not planar. Amazingly, the converse is also true. If the graph is not planar, then it contains a K_5 or $K_{3,3}$-subdivision.

Theorem 5.7 (Kuratowski - 1930)**.** *A graph is planar if and only if it does not contain a subdivision of K_5 or $K_{3,3}$.*

So in some sense K_5- and $K_{3,3}$-subdivisions are the only "obstructions" for planarity. Note that it states an equivalence between a topological statement (the graph having a planar drawing) and a combinatorial statement (the graph not containing a certain subdivision).

5.3 Coloring planar graphs

In 1852, Francis Guthrie asked the following question: "How many colors are needed to paint the countries of a political map so that adjacent countries get different colors?" He conjectured that four colors are always enough.

This problem can be modeled as the coloring of a planar graph if we place a vertex inside each country and connect two vertices by an edge if the associated countries have a common border.

Definition 23. *A* proper vertex coloring *of a graph G is a map $c : V(G) \to X$ for some set of colors X, such that $c(u) \neq c(v)$ whenever $uv \in E(G)$. The* chromatic number $\chi(G)$ *is the minimum size of such a set X, i.e. it is the minimum number of colors that the vertices of G can be properly colored with.*

Guthrie's conjecture can then be stated in terms of the chromatic number of planar graphs. After many failed and partially successful attempts, the conjecture was proved 120 years later by Appel and Haken.

Theorem 5.8. *For every planar graph G, $\chi(G) \leq 6$.*

Proof. We proceed by induction on the number of vertices n. For $n \leq 6$, there is nothing to do, we can always color the vertices with different colors. Let $n \geq 7$.

According to Proposition 5.5 there exists a vertex $v \in V(G)$ degree at most 5. Then by induction, $G - v$ can be (properly) 6-colored. In this coloring, the neighbors of v use at most 5 colors, so we can extend this to a 6-coloring of G by giving v a color not used by its neighbors. □

Theorem 5.9 ("Five color theorem", Heawood - 1890)**.** *If G is planar, then $\chi(G) \leq 5$.*

Proof. We use induction on the number of vertices n. As before, $n \leq 5$ is clear. By Proposition 5.5, there is a vertex $v \in V(G)$ of degree at most five, so by induction, we can color $G - v$ with five colors. If this coloring uses at most four colors on $N(v)$, then we can color v with the fifth color and we are done.

Now assume that v has five neighbors and all have different colors in the $G - v$ coloring. We label v neighbors $x_1, \ldots, x_5$ according to the clockwise order of the edges $vx_1, \ldots, vx_5$ leaving v. We also label the colors $i \in \{1, \ldots, 5\}$ such that each x_i has color i. Now consider a planar drawing of G. We call a path in G an (i, j)-path if all its vertices have color i or j.

Suppose that there is no $(1, 3)$-path from x_1 to x_3. Let R be the set of vertices that are reachable from x_1 by a $(1, 3)$-path. By assumption, x_3 is not in R. Then we can swap the colors 1 and 3 among the vertices in R: it gives a proper coloring of $G - v$, and it makes both x_1 and x_3 colored with 3. Then coloring v with 1 gives a 5-coloring of G.

Now suppose that there is a $(1, 3)$-path from x_1 to x_3. Together with v this path forms a cycle $\mathcal{C}$, all of whose vertices are colored with 1 or 3 (or uncolored in the case of v). Let S be the set of vertices reachable by $(2, 4)$-paths from x_2. Then the cycle $\mathcal{C}$ separates S from x_4 (here we use the Jordan Curve Theorem), so x_4 is not in S. Thus we can swap colors 2 and 4 in S, and then color v with 2. □

The (i, j)-paths used in this proof are called *Kempe chains*, named after Alfred Kempe who proposed a proof of Guthrie's conjecture in 1879. Although his proof was erroneous, the chain structures that he introduced were necessary for Heawood to prove his "five color theorem".

The conjecture is finally confirmed in 1976 by the famous "four colors theorem". The bound of four colors is of course optimal since K_4 is a planar graph.

Theorem 5.10 ("Four color theorem", Appel, Haken - 1976)**.** *For every planar G, $\chi(G) \leq 4$.*

The proof of this theorem required heavy computer assistance. Even today, there is no proof known for this theorem that would be feasible by hand (the simplest proofs still involve checking over 600 cases).

For specific families of planar graphs, four colors are not even necessary. In particular three colors are enough for graphs without triangles.

Theorem 5.11 (Grötschz - 1959)**.** *Let G be a planar graph on $n \geq 3$ vertices without triangle, then $\chi(G) \leq 3$.*

Theorem 5.12 (Aksionov - 1974)**.** *Let G be a planar graph on $n \geq 3$ vertices such that G contains at most three triangles, then $\chi(G) \leq 3$.*

Lecture 6

Coloring.

6.1 Vertex coloring

Recall the we call a graph (properly) *k-colorable* if its vertices can be colored with k colors in such a way that no two adjacent vertices get the same color, and that the chromatic number $\chi(G)$ is the minimum k such that G is k-colorable.

So far we have looked at the chromatic number of planar graphs, but it also makes sense to study the concept for general graphs. Actually, many problems can be interpreted as graph coloring problems, such as assignment or scheduling problems.

Example: The students at a certain university have annual examinations in all the courses they take. Naturally, examinations in different courses cannot be held concurrently if the courses have students in common. How can all the examinations be organized in as few parallel sessions as possible? To find a schedule, consider the graph G whose vertex set is the set of all courses, two courses being joined by an edge if they give rise to a conflict. Clearly, independent sets of G correspond to conflict-free groups of courses. Thus, the required minimum number of parallel sessions is the chromatic number of G.

Definition 24. *The* complement *of a graph G is the graph $\overline{G}$ with vertex set $V(\overline{G}) = V(G)$ and edge set $E(\overline{G}) = \{xy : x, y \in V(G), xy \notin E(G)\}$. The* clique number *$\omega(G)$ is the size of the largest complete subgraph (clique) of G. So $\omega(G) = \alpha(\overline{G})$.*

Proposition 6.1. *Below is a list of basic properties about the chromatic number.*

1. $\chi(K_s) = s$
2. *G is bipartite if and only if $\chi(G) \leq 2$*
3. *If H is a subgraph of G, then $\chi(H) \leq \chi(G)$*
4. *$\chi(G) \geq \omega(G)$ for every graph G.*

Proof.

- 1. and 2. The properties follow from the definition.
- 3. We simply need to notice that every proper coloring of G provides a coloring of H.
- 4. It can be proven by a combination of properties 1. and 3.

□

So $\omega(G)$ is a trivial lower bound on the chromatic number, but it is not necessarily tight. Indeed, the following graph G satisfies $\chi(G) = 4$ and $\omega(G) = 3$.

$G =$ 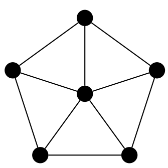

As a follow up question, we can wonder if large cliques are somehow a necessary conditions of graph with large chromatic number. It turns out that it is not: in general, the chromatic number cannot be bounded by any function of the clique number. The first result in this direction is the following theorem for graphs without triangle.

Theorem 6.2 (Zykov - 1949). *For every positive integer k, there exists a graph G with $\omega(G) = 2$ and $\chi(G) \geq k$.*

Proof. We prove this by induction on k. For $k = 2$, we can take G_2 be a single edge.

Now suppose that $k \geq 3$. Let $H_1, \ldots, H_{k-1}$ be $k-1$ disjoint copies of G_{k-1}. For every $(v_1, \ldots, v_{k-1}) \in V(H_1) \times \cdots \times V(H_{k-1})$, add a vertex $X(v_1, ..., v_{k-1})$ that is only connected to $v_1, ..., v_{k-1}$. Define G_k as the resulting graph.

Then $\omega(G_k) = 2$. Indeed the subgraphs $H_1, \ldots, H_{k-1}$ do not contain a triangle, so if there exists a triangle in G_k, then one of its vertices is $X(v_1, \ldots, v_{k-1})$ for some $(v_1, \ldots, v_{k-1}) \in H_1 \times \cdots \times H_{k-1}$. But for $i \neq j$, v_i and v_j are not neighbors. Thus $X(v_1, \ldots, v_{k-1})$ cannot be contained in a triangle either.

Now we show that $\chi(G_k) \geq k$. Let $c : V(G_k) \to \{1, \ldots, \ell\}$ be a proper ℓ-coloring of G_k. Since c restricted to each H_i, $i \in \{1, \ldots, k-1\}$ is also a proper ℓ-coloring, by the induction assumption $\ell \geq k-1$. Furthermore there exist $v_1 \in H_1, \ldots, v_{k-1} \in H_{k-1}$ which are colored with exactly $k-1$ different colors in c. Now those vertices are the neighbors of vertex $X(v_1, \ldots, v_{k-1})$ in G_k. Therefore $X(v_1, \ldots, v_{k-1})$ is colored in c with a k^{th} color, and thus $\ell \geq k$.

□

Other graphs can be constructed that achieve arbitrary large chromatic number while maintaining the triangle free property.

Shift graphs: Let $n \geq 2^k$. The shift graph $S_{n,2}$ is the graph with vertex set $\{(i,j) : 1 \leq i < j \leq n\}$, where (i,j) and (k,l) are joined by an edge if $j = k$. Then for all $k \geq 2$, $\chi(S_{n,2}) \geq k$ while $\omega(S_{n,2}) = 2$.

Mycielsky graphs: For a graph G on n vertices $v_1, \ldots, v_n$, the Mycielskian of G is denoted $\mu(G)$. It contains G as a subgraph and $n+1$ additional vertices $u_1, \ldots, u_n$ and w. All the u_i, $i \in \{1, \ldots, n\}$ are connected to w and for all $v_iv_j \in E(G)$, $\mu(G)$ contains the edges u_iv_j and v_iu_j. Some important properties of the Mycielskian are: (i) if G is triangle free, so is $\mu(G)$; (ii) $\chi(\mu(G)) = \chi(G) + 1$.

The Mycielsky graphs are a sequence of graphs defined by: $M_2 = K_2$ (one edge graph) and for $k > 2$, $M_k = \mu(M_{k-1})$. Then for all $k \geq 2$, $\chi(M_k) = k$ while $\omega(M_k) = 2$.

A generalization of Theorem 6.2 was provided by Erdös. The following theorem states that even graphs containing only large cycles (which can be thought as the "opposite" of cliques) can have a large chromatic number.

Theorem 6.3 (Erdös - 1959). *For every positive integers k and g, there exists a graph of girth larger that g and such that $\chi(G) \geq k$.*

Now we investigate another direction for bounding the chromatic number of a graph in terms of its largest independent set and largest degree. A k-coloring of a graph can be thought as the splitting of its vertices into k independent sets. This readily implies the following lower bound on $\chi(G)$.

Proposition 6.4. *For every graph G, we have $\chi(G) \geq \frac{|V(G)|}{\alpha(G)}$.*

Proof. Given a proper coloring with $\chi(G)$ colors, the color classes, labelled $V_1, \ldots, V_{\chi(G)}$, are independent sets, and thus have size at most $\alpha(G)$. Hence we have:

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

□

Proposition 6.5. *For any two graphs G_1 and G_2 on the same set of vertices V, $\chi(G_1 \cup G_2) \leq \chi(G_1)\chi(G_2)$, where the set of edges of the graphs union is simply defined by $E(G_1 \cup G_2) = E(G_1) \cup E(G_2)$.*

Proof. Denote $k_i = \chi(G_i)$ for $i = 1, 2$ and let us consider $c_i : V \to \{1, \dots, k_i\}$ a proper k_i-coloring of G_i. We can color the vertices in V with elements of the set $\{1, \dots, k_1\} \times \{1, \dots, k_2\}$. Indeed, $c(v) = (c_1(v), c_2(v))$ gives a proper k_1k_2-coloring of $G = G_1 \cup G_2$, because if two vertices u, v are adjacent in G, then they are also adjacent in some G_i, hence the i^{th} coordinate of their colors will differ. So $\chi(G) \leq k_1k_2$. □

Definition 25. *Let $G = (V, E)$ a graph. We say that G is* k-degenerate *if every subgraph of G has a vertex of degree less than or equal to k.*

The following proposition connects the degeneracy of G and the chromatic number of G. The six color theorem for planar graphs is in fact a special case of it.

Theorem 6.6 (Halin - 1967)**.** *If G is k-degenerate, then $\chi(G) \leq k + 1$.*

Proof. We proceed by induction on the number of vertices. The statement is clearly true for all G with at most $k + 1$ vertices. For larger graphs, pick a vertex v of degree $d(v) \leq k$ in G. The graph $G - v$ is k-degenerate because every subgraph of $G - v$ is a subgraph of G, so it has a proper $(k + 1)$-coloring c. Then there are at most k different colors in c among the neighbors of v since $d(v) \leq k$. Thus we can assign to v a color that does not appear among its neighbors. This extends c into a proper $k + 1$-coloring of G. □

Corollary 6.7. $\chi(G) \leq \Delta(G) + 1$ *for every graph G.*

Proof. Recal that $\Delta(G)$ is the largest degree in G. Then G is clearly $\Delta(G)$-degenerate. The proof is therefore immediately concluded by application of Theorem 6.6. □

This inequality is tight, i.e. it is an equality, for cliques and odd cycles. Interestingly those two types of grahs are basically the only ones for which $\chi(G) = \Delta(G) + 1$, as shown by the following theorem.

Theorem 6.8 (Brooks, 1941)**.** *Let G be a connected graph which does not have a subgraph isomorphic to K_n or C_{2k+1} for any integer n or k. Then $\chi(G) \leq \Delta(G)$.*

6.2 Edge coloring

Definition 26. *A* proper edge coloring *of a graph G is a map $c : E(G) \to X$ for some set of colors X, such that $c(e) \neq c(e')$ whenever e, e' are distinct edges that share a vertex.*

The edge-chromatic number $\chi'(G)$ *of G is the minimum size of such a set X, i.e. it is the minimum number k such that the G can be properly k-edge-colored.*

Examples

- The edge chromatic number of the triangle is 3: $\chi'(K_3) = 3$.
- It is the same for K_4: $\chi'(K_4) = 3$.
- Even length cycles have edge chromatic number 2: $\chi'(C_{2k}) = 2$

The picture on the right is an edge coloring of the Petersen graph with four colors. It is not difficult to see that its edge chromatic number is equal to 4.

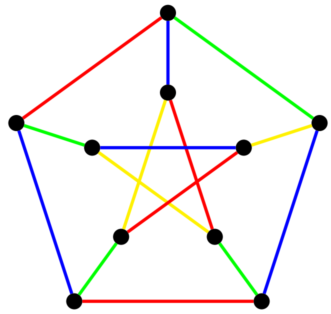

Definition 27. *A* matching $\mathcal{M}$ *is a set of vertex-disjoint edges, i.e. a set of edges such that no two of them share an end vertex.*

Each color class is a matching, so a proper edge-coloring is a partition of the edges into matchings.

Lemma 6.9. *For any graph G with at least one edge:* $\Delta(G) \leq \chi'(G) \leq 2\Delta(G) - 1$.

Proof. There is vertex of degree $\Delta(G)$, and an edge coloring must give different colors to each of the $\Delta(G)$ edges at that vertex. This implies $\chi'(G) \geq \Delta(G)$.

The upper bound follows by a proof similar to the proof of Theorem 6.6: we apply induction on the number of edges. So delete an edge e, and take a $(2\Delta(G)-1)$-edge-coloring of $G-e$. Since an edge shares a vertex with at most $2(\Delta(G)-1)$ other edges, there must be a free color left for e. □

This simple upper bound can however be significantly improved. Indeed, the following theorem states that any graph G has edge-chromatic number either $\Delta(G)$ or $\Delta(G)+1$. Both are possible, since even cycles have $\chi'(G) = \Delta(G)$ whereas odd cycles have $\chi'(G) = \Delta(G)+1$. However the algorithm in the proof still does not always give the exact number. In fact deciding which of the two values is the edge-chromatic number of a given graph is an NP-hard problem.

Theorem 6.10 (Vizing - 1964)**.** *For every graph G we have $\chi'(G) = \Delta(G)$ or $\Delta(G)+1$.*

Proof. We only need to show $\chi'(G) \leq \Delta(G)+1$. We prove this by induction on the number of edges. The statement clearly holds for a graph without edges.

Given an edge $xy \in E(G)$, we describe an algorithm that, given an edge coloring of $G-xy$ with at most $\Delta(G)+1$ colors, produces an edge coloring of G with this many colors. But to find a color for xy, the algorithm may have to change the colors of other edges.

So take a proper $\Delta(G)+1$-edge-coloring of $G-xy$. As there are more than $\Delta(G)$ available colors, every vertex has at least one color missing from the edges touching it. We say that a vertex v is c-free, if no edge incident to v has color c.

Now let us build a "fan" of edges around x as follows: take the edge xy, and let $y_0 = y$. We know that y_0 is c_1-free for some color c_1. Now let xy_1 be an edge of color c_1, if such an edge exists. Then y_1 is c_2-free for some color c_2.

Continue similarly: if there is an edge at x of color c_i that we have not looked at, let xy_i be this edge, and let c_{i+1} be a color missing at y_i. We repeat this process until it is no longer possible to add an edge. Let xy_k be the last edge added. There are two possible reasons for getting stuck: either there is no edge at x of color c_{k+1}, or this edge already appeared at some previous xy_i.

Case 1. x is c_{k+1}-free.

In this case we can just "rotate" the colors around x: define the new color of xy_i to be c_{i+1} for every $i = 0, \ldots, k$, keeping the colors of all other edges. This does not create any issue with any y_i because it had been c_{i+1}-free. There is no problem at x either, because the only new color introduced is c_{k+1}. Hence this is a proper $\Delta(G) + 1$-edge-coloring of G including the edge $xy = xy_0$.

Case 2. x is not c_{k+1}-free.

As noted above, the maximality of k implies that some edge xy_i is c_{k+1}-colored, so by definition $c_{k+1} = c_i$. Let us just call this color c, and let d be a color such that x is d-free.

Now let $\mathcal{P}$ be a maximal path starting at x that only uses the colors c and d. Since every vertex touches at most one c-colored and one d-colored edge, and x is d-free, there is a unique such maximal path $\mathcal{P} = xy_i \ldots z$ for some vertex z. Let us swap the colors c and d along the path. The same observation shows that this is still a proper edge-coloring of $G - e$. Now we again distinguish two sub-cases.

Sub-case 2.A. $z = y_{i-1}$.

Note that y_{i-1} was c-free, so the last edge of $\mathcal{P}$ must have had color d. After swapping, y_{i-1} is no longer c-free, but it is now d-free. However, the edge xy_i has color d, as well, so the edges xy_i satisfy the required properties with $c_i = d$ now. As x is now c-free, we arrive at a situation covered by Case 1, which gives a proper coloring of G, by rotating the colors, i.e. coloring each xy_j with color c_{j+1}.

Sub-case 2.B. $z \neq y_{i-1}$.

In this case, y_{i-1} remains c-free, and x becomes c-free as well. This means that we can rotate the colors on the first i edges only to get a proper edge coloring: recolor xy_j with the color c_{j+1} for every $j = 0, \ldots, i - 1$, and leave the color of other edges the same as what we had after swapping c and d along $\mathcal{P}$. By the same arguments as before, every edge is colored, but no two touching edges get the same color. □

Lecture 7

Matchings. Bipartite Graphs.

7.1 Stable matchings

Recall the following from previous chapters: (i) a *matching* $\mathcal{M}$ is a set of vertex-disjoint edges; (ii) a graph is k-regular if all degrees equal k.

Definition 28. *A* perfect matching $\mathcal{M}$ *in a graph* G *is a matching that touches every vertex of* G*. In other words, it is a 1-regular subgraph of* G*.*

Example: suppose that you have n students that want to do their internships in n companies. Each have sent their applications to all the companies. Each student and each company has their list of preferences, and we want to pair up the students with the companies so that this assignment is *stable* in the following sense: there is no student and company that would both prefer to work with each other rather than their assigned pairs. Such configuration can be formally defined as a *stable* matching.

Definition 29. *Let* G *be the complete bipartite graph* $K_{n,n}$ *with parts* A *and* B *(*$|A| = |B| = n$*). Assume that for each vertex* v *there exist a complete strict ordering* $\prec_v$ *in* $N(v)$*, called "set of preferences". We say that a perfect matching* $\mathcal{M}$ *is* stable *if there exists no vertices* $a, x \in A$ *and* $b, y \in B$*, such that:* $ay, xb \in \mathcal{M}$ *but* $y \prec_a b$ *and* $x \prec_b a$*.*

Theorem 7.1 ("Stable mariage theorem", Gale, Shapley - 1962)**.** *Let* $G = (V = A \cup B, E)$ *be the complete bipartite graph on* $2n$ *vertices such that* $|A| = |B| = n$*. Then for any sets of preferences* $\prec_v$*,* $v \in V$*,* G *has a stable matching.*

This theorem can be proven using the following algorithm which actually constructs such a stable matching.

Input: The complete bipartite graph $G = (A \cup B, E)$ such that $|A| = |B| = n$.
Input: A set of preferences $\prec_v$ associated to each vertex v of G.
Output: A stable perfect matching $\mathcal{M}$ of G.

Initialize the matching: $\mathcal{M} = \emptyset$.
For $a \in A$ consider $b_a^{(n)} \prec_a \cdots \prec_a b_a^{(1)}$ the ordered list of neighbors of a.
Initialize the index in the preference list: $i_a = 1$.
while $\exists a \in A$ *such that* $\nexists ay \in \mathcal{M},\ y \in B$ **do**
  Set $b = b_a^{(i_a)}$.
  /* a ``proposes'' to b */
  if $\nexists\, a'b \in \mathcal{M},\ a' \in A$ **or** $\exists\, a'b \in \mathcal{M},\ a' \prec_b a$ **then**
    Add ab to the matching: $\mathcal{M} \leftarrow \mathcal{M} \cup \{ab\}$.
    Remove $a'b$ from the matching: $\mathcal{M} \leftarrow \mathcal{M} \setminus a'b$.
    Update the next neighbor to try for a': $i_{a'} \leftarrow i_{a'} + 1$.
  else
    Update the next neighbor to try for a: $i_a \leftarrow i_a + 1$.

Algorithm 2: Gale-Shapley

Proof.

Termination:

At each iteration of the `while` loop, line 4, either the index in the preference list of a, i_a, or in the list of a', $i_{a'}$, is increased by 1. Since the preference list of each vertex v is finite, so is the maximum number of iterations of the loop. Therefore the algorithm terminates.

Correctness:

The output matching is perfect. Indeed, if there is a pair of vertices $a \in A, b \in B$, such that both are not in the matching, then a must have proposed to b at some point. However, if a vertex $b \in B$ is in $\mathcal{M}$ at some step of the algorithm, then it stays in $\mathcal{M}$.

The output matching is stable. Assume that $ab \notin \mathcal{M}$. Upon completion of the algorithm, it is not possible for both a and b to prefer each other over their current match. If a prefers b to its current match, then a must have proposed to b before its current match. If b accepted the proposal but is matched to another vertex at the end, then b prefers its current match over a. If b rejected the proposal, then b was already matched to a vertex that it prefers. □

7.2 Maximum matchings

The previous problem is about perfect matchings in a complete bipartite graph. But not every graph has a perfect matching. So how can we find a *largest* or *maximum* matching? As a reminder, a *maximum* matching must not be confused with a *maximal* matching, which just means that it cannot be extended, i.e. no larger matching contains it as a subgraph. Now, how can we even tell if a graph has a perfect matching? How can we check if a given matching is maximum?

Observe that maximum matching are necessarily maximal but the opposite is not true. For instance, in a path with three edges, the matching consisting of only the middle edge is maximal but not maximum. A greedy approach that keeps adding edges without removing any, like the one used for spanning trees, would thus probably not lead to a maximum matching. Instead, an algorithm to find maximum matchings would need some kind of backtracking where some edges previously selected can be thrown away latter in the execution. The following notion allows to elaborate upon this idea.

Definition 30. *Given a matching $\mathcal{M}$ in a graph G, a path is* alternating *if among any two consecutive edges on the path, exactly one is in $\mathcal{M}$. An alternating path with at least one edge is* augmenting *if its first and last vertices are not covered, also said to be "saturated", by $\mathcal{M}$. Note that an augmenting path always has odd length.*

Theorem 7.2 (Berge - 1957)**.** *A matching $\mathcal{M}$ is maximum if and only if there is no augmenting path for $\mathcal{M}$.*

Proof. We prove that $\mathcal{M}$ is not maximum if and only if there is an augmenting path for $\mathcal{M}$. First suppose that there is an augmenting path $\mathcal{P} = v_1v_2 \dots v_{2k}$ for the matching $\mathcal{M}$. So $v_2v_3, \dots, v_{2i}v_{2i+1}, \dots, v_{2k-2}v_{2k-1} \in \mathcal{M}$. Then we can get a larger matching by removing these edges from $\mathcal{M}$ and replacing them by $v_1v_2, \dots, v_{2i-1}v_{2i}, \dots, v_{2k-1}v_{2k}$. In other words, we replace $\mathcal{M}$ by $\mathcal{M}' = \mathcal{M} \triangle E(\mathcal{P})$.[1] Then $\mathcal{M}'$ is a matching since v_1 and v_2k were unmatched by $\mathcal{M}$, and we have $|\mathcal{M}'| = |\mathcal{M}| + 1$, hence $\mathcal{M}$ is not maximum.

Now suppose $\mathcal{M}$ is not maximum, i.e. there is a matching $\mathcal{M}'$ with $|\mathcal{M}'| > |\mathcal{M}|$. Let H be the graph with $V(H) = V(G)$ and $E(H) = \mathcal{M} \triangle \mathcal{M}'$. Every vertex of H has degree 0, 1 or 2, so each component of H is either a cycle, a path, or an isolated vertices. A cycle in H has the same number of edges from $\mathcal{M}$ and $\mathcal{M}'$, so $|\mathcal{M}'| > |\mathcal{M}|$ implies that there is a path $\mathcal{P}$ in H with more edges from $\mathcal{M}'$ than from $\mathcal{M}$. Then $\mathcal{P}$ is an augmenting path for $\mathcal{M}$. □

[1] We write $S \triangle T$ for the *symmetric difference* of two sets, i.e. $S \triangle T = (S \setminus T) \cup (T \setminus S)$.

This result shows how to extend a matching: we just need to find an augmenting path. If no such path exists, then the matching is maximum. But it is not so easy to check if an augmenting path exists; a naive algorithm would take exponentially many steps. An efficient algorithm for a maximum matching was found by Edmonds in 1965 which is not covered in this course. However, the problem is much simpler in bipartite graphs.

7.3 Matchings in bipartite graphs

Let $G = (A \cup B, E)$ be a bipartite graph, and $\mathcal{M}$ be a matching in it. Note that the end vertices of every augmenting path $\mathcal{P}$ are in different parts because $\mathcal{P}$ has odd length. So if there is an augmenting path in G, then it starts in $S = A \backslash V(\mathcal{M})$ and ends in $T = B \backslash V(\mathcal{M})$. Crucially, such a path starting in S will always use an edge *not in* $\mathcal{M}$ to jump to B, and an edge *in* M to jump back to A.

So let us define $D_{\mathcal{M}}$ to be the directed graph obtained from G by orienting edges not in $\mathcal{M}$ from A to B, and orienting edges in $\mathcal{M}$ from B to A. The observations above show that augmenting paths in G correspond to directed S-T paths in $D_{\mathcal{M}}$; an S-T path being a path that starts in S and ends in T.

The point is, finding S-T paths in a directed graph is easy, using breadth-first-search for example. Together with Theorem 7.2, we arrive at the following algorithm for finding a maximum matching in a bipartite graph.

Input: A bipartite graph $G = (A \cup B, E)$.
Output: A maximum matching $\mathcal{M}$ of G.

```
Initialize the matching: M = ∅.
Initialize the subsets of vertices S and T: S = A , T = B.
Initialize the directed graph D_M: D_M = (A ∪ B, {ab| a ∈ A, b ∈ B, ab ∈ E(G)})
while there exists a directed S-T path in D_M do
    Find such a directed S-T path P in D_M using BFS.
    Extend the matching using the path: M ← M△P.
    Update the subsets of vertices S and T: S = A \ V(M) , T = B \ V(M).
    Update the set of oriented edges in D_M:
      V(D_M) = {ab| a ∈ A, b ∈ B, ab ∈ E(G) \ M} ∪ {ba| a ∈ A, b ∈ B, ab ∈ M}
```

Algorithm 3: Augmenting path

7.4 Perfect matchings

The next question we study is: when does a bipartite graph have a perfect matching? Apart from the trivial condition $|A| = |B|$, it is easy to see that we need every set $X \subseteq A$ to have at least $|X|$ neighbors. This condition turns out to be sufficient, as shown by the next theorem.

Definition 31. *A matching $\mathcal{M}$* matches *or* covers *a vertex set X if every vertex of X is contained in an edge of $\mathcal{M}$. For a set $X \subset V(G)$, we define the* neighborhood *of X in G as $N(X) = \{v \in V(G) \setminus X \mid \exists u \in X, uv \in E(G)\}$.*

Theorem 7.3 ("Mariage theorem", Hall - 1935)**.** *Let $G = (A \cup B, E)$ be a bipartite graph. Then G has a matching covering A if and only if for all $X \subseteq A$ we have $|N(X)| \geq |X|$.*

Proof. For the implication, note that if G has a matching $\mathcal{M}$ that matches A, then the vertices of any $X \subset A$ are matched by $\mathcal{M}$ to $|X|$ distinct neighbors in B, which implies $|N(X)| \geq |X|$.

For the reciprocal, suppose $|N(X)| \geq |X|$ for every $X \subseteq A$. Let $\mathcal{M}$ be a maximum matching, and let $S = A \setminus V(\mathcal{M})$ and $T = B \setminus V(\mathcal{M})$ be the unmatched vertices in A and B. By Theorem 7.2 and the discussion above, there is no S-T path in $D_{\mathcal{M}}$.

Let X_B denote the set of vertices in B that can be reached from S via a directed path in $D_{\mathcal{M}}$, then X_B is disjoint from T, therefore all vertices in X_B are matched by $\mathcal{M}$. Let $X_A \subseteq A$ be the set of vertices matched to X_B by $\mathcal{M}$, then X_A is clearly disjoint from S. As the matching edges are directed from B to A, the vertices in X_A are also reachable from S in $D_{\mathcal{M}}$ via X_B.

To finish the proof, we then only need to check that $N(X_A \cup S) \subseteq X_B$, because this together with the initial assumption applied to $X = X_A \cup S$ would imply:

$$|X_B| \geq |N(X_A \cup S)| \geq |X_A \cup S| = |X_A| + |S| = |X_B| + |S|$$

and hence $S = \emptyset$, which is what we want to prove.

To prove that $N(X_A \cup S) \subseteq X_B$, first note that all vertices in $X_A \cup S$ are reachable from S in $D_{\mathcal{M}}$. Now if there is a vertex $b \in B \setminus X_B$ adjacent to some $a \in X_A \cup S$, then ab is cannot be a matching edge. Thus ab is directed from a to b in $D_{\mathcal{M}}$. Consequently, $b \notin X_B$ is reachable from S, which is a contradiction. Therefore $N(X_A \cup S) \subseteq X_B$ and hence $S = \emptyset$ which implies that $\mathcal{M}$ covers A. □

If $|A| = |B|$, then of course every matching covering A covers B, as well. So Hall's theorem implies that a bipartite graph $G = (A \cup B, E)$ has a perfect matching if and only if $|A| = |B|$ and $|N(X)| \geq |X|$ for every $X \subseteq A$.

The property "$|N(X)| \geq |X|$ for every $X \subseteq A$" is often referred to as *Hall's condition*. Testing this property algorithmically is slow, but it turns out to be a convenient thing to check in many theoretical applications.

Corollary 7.4. *Let $k \geq 1$ be a integer. Every k-regular bipartite graph has a perfect matching.*

Proof. Let $G = (A \cup B, E)$ be a k-regular bipartite graph. Since every edge touches exactly one vertex of A, and every vertex of A touches exactly k edges, the number of edges in G is exactly $k|A|$. By a similar argument, the number of edges is exactly $k|B|$, so $k \geq 1$ implies $|A| = |B|$.

As discussed above, Hall's condition would now guarantee a perfect matching. To check it, let $X \subseteq A$ be a vertex set. We can double-count the edges touching X. On the one hand, there are exactly $k|X|$ such edges. On the other hand, every such edge has its other endpoint in $N(X)$, so the number of such edges is at most $k|N(X)|$. Hence $k|X| \leq k|N(X)|$, which again implies $|X| \leq |N(X)|$. □

Finding perfect matchings in regular bipartite graphs can be iterated to get a proper edge coloring.

Corollary 7.5. *Let G be a k-regular bipartite graph. Then $\chi'(G) = \Delta(G)$.*

Proof. Recall that $\chi'(G)$ is the edge-chromatic number of G. The inequality $\chi'(G) \geq \Delta(G)$ is thus obvious.

Now, we can prove that $\chi'(G) \leq \Delta(G)$ by induction on k. The inequality is true for $k = 0$. For $k \geq 1$, apply Corollary 7.4 to find a perfect matching $\mathcal{M}$ in G. Color this matching with some color, and apply induction on the $(k-1)$-regular graph $G - \mathcal{M}$ to find a $(k-1)$-edge-coloring of the remaining edges. □

Note that this improves Vizing's theorem for regular bipartite graphs. Actually, we can drop the regularity condition which yields the following theorem.

Theorem 7.6 (König - 1916)**.** *For every bipartite graph G, we have $\chi'(G) = \Delta(G)$.*

Lecture 8

König's theorem. Flows.

8.1 König's Theorem

Definition 32. *Given a graph G, a* vertex cover *for G is a set of vertices $\mathcal{C} \subset V(G)$ such that every edge of G is incident with a vertex in C. The minimum size of a vertex cover in G is denoted by $\tau(G)$.*

We can observe that $\mathcal{C} \subseteq V(G)$ is a vertex cover for G if and only if $V(G) \setminus \mathcal{C}$ is an independent set in G, which readily implies that $|V(G)| = \tau(G) + \alpha(G)$. Just like finding a maximum independent set, finding a minimum vertex cover in a graph is a hard problem in general.

It can also be noticed that vertex covers are related to matchings: in fact the set of all saturated vertices from any maximal matching forms a vertex cover. This is true in particular for a maximum matching. We denote by $\nu(G)$ the size of a maximum matching in G. Note that the "size" of a matching is the number of edges in it, while the "size" of a vertex cover is the number of vertices in it. The following theorem shows that those two characteristics are equal for bipartite graphs.

Theorem 8.1 (König - 1931)**.** *Let $G = (A \cup B, E)$ be a bipartite graph. Then $\nu(G) = \tau(G)$.*

Proof. It is clear that $\nu(G) \leq \tau(G)$, because a vertex cover has to cover every edge of the matching with one of the end vertices of the edge. So to cover a matching of size $\nu(G)$, we need at least this many vertices.

Now we prove that $\nu(G) \geq \tau(G)$. Let $\mathcal{M}$ be a maximum matching of G, and as before, let $S = A \setminus V(\mathcal{M})$ and $T = B \setminus V(\mathcal{M})$ be the set of unmatched vertices. As observed previously, since $\mathcal{M}$ maximum there is no S-T path in $D_{\mathcal{M}}$.

Let X_B be the set of vertices in B that can be reached from S via a directed path in $D_{\mathcal{M}}$, and let Y_A be the set of vertices in A that cannot be reached from S in $D_{\mathcal{M}}$. Clearly, X_B is disjoint from T and Y_A is disjoint from S. Note also, that either both end vertices of an edge in $\mathcal{M}$ can be reached from S, or neither of them can be. In the former case, the edge has an end vertex in X_B, in the latter case it has an end vertex in Y_A. Consequently, $|X_B \cup Y_A| = |\mathcal{M}|$

Now we can prove that $X_B \cup Y_A$ is a vertex cover. Suppose it is not, then there is an edge uv with $u \in A$ and $v \in B$ not covered by this set, i.e. $u \notin Y_A$ and $v \notin X_B$. But uv is not a matching edge, so in $D_{\mathcal{M}}$, it is directed from u to v. But then $u \in A \setminus Y_A$ can be reached from S in $D_{\mathcal{M}}$, thus v should also be reachable via the edge uv. This contradicts $v \notin X_B$.

Therefore $X_B \cup Y_A$ is a vertex cover of size $\nu(G)$, implying $\nu(G) \geq \tau(G)$. □

8.2 Flows

Definition 33. *A* network *is a directed graph $G = (V, E)$ with two special vertices, the* source *$s \in V$ and the* sink *$t \in V$, together with a non-negative capacity function $c : E \to \mathbb{R}_+$.*

Definition 34. *A* flow *in a network* $G = (V, E)$ *is a function* $f : V^2 \to \mathbb{R}$ *such that:*

1. $f(u, v) = 0$ *if* $\overrightarrow{uv} \notin E$
2. *capacity constraint:* $0 \leq f(u, v) \leq c(u, v)$ *for every* $\overrightarrow{uv} \in E$
3. *flow conservation:*
   - *for every* $v \in V \setminus \{s, t\}$, $\displaystyle\sum_{w \in V} f(v, w) = \sum_{w \in V} f(w, v)$
   - $\displaystyle\sum_{w \in V} f(s, w) - f(w, s) = \sum_{w \in V} f(w, t) - f(t, w)$

The *value* $|f|$ of a flow is defined as $\sum_{w \in V} f(s, w) - f(w, s)$. Our main question of study is: what is the maximum value of a flow in a network G?

A *cut* in a network G is a partition (S, T) of the vertices (i.e. $S \cap T = \emptyset$ and $S \cup T = V$) such that $s \in S$ and $t \in T$. For two vertex sets $X, Y \subseteq V$ we define:

$$f(X, Y) = \sum_{x \in X,\, y \in Y} f(x, y) - f(y, x) \quad \text{and} \quad c(X, Y) = \sum_{x \in X,\, y \in Y} c(x, y)$$

We call $c(S, T)$ the *capacity of the cut* (S, T).

Proposition 8.2. *For every cut* (S, T) *we have* $f(S, T) = |f|$.

Proof. By definition, $|f| = \sum_{w \in V} f(s, w) - f(w, s)$. Besides, by conservation of the flow, for all $v \in S \setminus \{s\}$, $\sum_{w \in V} f(v, w) = \sum_{w \in V} f(w, v)$. Hence:

$$\begin{aligned} |f| &= \sum_{w \in V} f(s, w) - f(w, s) + \sum_{v \in S \setminus \{s\}} \sum_{w \in V} f(v, w) - f(w, v) \\ &= \underbrace{\sum_{u, v \in S} f(u, v) - f(v, u)}_{=0} + \sum_{u \in S,\, v \in T} f(u, v) - f(v, u) = f(S, T) \end{aligned}$$

where the second equality holds by simply rearranging the terms. □

8.3 Ford-Fulkerson algorithm

So how large can a flow be in a network? Proposition 8.2 implies that for any cut (S, T),

$$|f| = f(S, T) = \sum_{u \in S, v \in T} f(u, v) - f(v, u) \leq \sum_{u \in S, v \in T} f(u, v) \leq \sum_{u \in S, v \in T} c(u, v) = c(S, T)$$

In particular, the maximum value of a flow is bounded by the minimum capacity of a cut.

Theorem 8.3 ("Max-flow, min-cut", Ford–Fulkerson - 1956)**.** *In every network, the maximum value of a flow equals the minimum capacity of a cut.*

The proof of the lower bound is algorithmic, and uses the concept of *residual graph* and *augmenting path*.

Definition 35. *Given a flow f, the* residual graph G_f *is a directed graph defined by:* $V(G_f) = V(G)$ *and* $E(G_f) = \{\overrightarrow{uv} \mid \overrightarrow{uv} \in E(G) \text{ or } \overrightarrow{vu} \in E(G)\}$. *It has a modified capacity function c_f defined by:*

$$\text{for all } \overrightarrow{uv} \in E(G_f), \quad c_f(u,v) = \begin{cases} c(u,v) - f(u,v) + f(v,u) & \text{if } \overrightarrow{uv} \in E(G) \\ -f(u,v) + f(v,u) & \text{otherwise} \end{cases}$$

An augmenting path *is a directed s-t path $v_0, \ldots, v_\ell$ in G_f ($v_0 = s$ and $v_\ell = t$) such that $c_f(v_{i-1}v_i) > 0$ for all $i \in \{1, \ldots, \ell\}$.*

Intuitively, the residual capacity of each edge of G_f corresponds to the additional flow that can go through it in G: $c(u,v) - f(u,v)$ is the flow that can be added on $\overrightarrow{uv}$, while $f(v,u)$ is the flow that can be removed from the edge in the opposite direction, $\overrightarrow{vu}$. The capacity of edges $\overrightarrow{uv}$ that are not originally present in G is set to 0.

Ford-Fulkerson algorithm repeatedly finds an augmenting path and pushes some extra flow through it to increase the value of f.

Input: A network $G = (V, E)$ with source $s \in V$, sink $t \in V$ and capacity function c.
Output: A flow f and a cut (S, T) such that $|f| = c(S, T)$.

1 Initialize the flow and the residual graph: for all u, v connected in G, $f(u,v) = 0$.
2 **while** *there exists an augmenting path* $\mathcal{P}$ **do**
3 Define δ the largest possible increment of flow along the path: $\delta = \min_{\overrightarrow{uv} \in \mathcal{P}} \{c_f(u,v)\}$
4 Increase the flow by δ along the path: for all $\overrightarrow{uv} \in \mathcal{P}$,
5 if $f(u,v) - f(v,u) + \delta \geq 0$ $\begin{cases} f(u,v) \leftarrow f(u,v) - f(v,u) + \delta \\ f(v,u) \leftarrow 0 \end{cases}$
otherwise $\begin{cases} f(u,v) \leftarrow 0 \\ f(v,u) \leftarrow -(f(u,v) - f(v,u) + \delta) \end{cases}$
6 Update G_f.
7 Define S as the set of vertices reachable by a directed path from s in G_f such that all intermediate edges have non-zero residual capacity. Define T as: $T = V \setminus S$.

Algorithm 4: Ford-Fulkerson

We now verify that the algorithm terminates and is correct in the specific case where c is integer valued.

Proof of Theorem 8.3.
Termination: Assume that all capacities in the network are integral, we then have that $\delta > 0$ is an integer at each step, and so $\delta \geq 1$. Since the flow must be finite, bounded by the capacity of any cut, the algorithm therefore terminates in a finite number of steps.

Correctness: First note that, at each step, f is feasible in G, i.e. it satisfies the three conditions in the definition of a flow for the initial graph G. Now, at the last iteration of the `while` loop, line 4, the condition is not met, i.e. there exist no augmenting path. Therefore, by definition of S, for all $\overrightarrow{uv} \in E(G_f)$ such that $u \in S$ and $v \in T$, $c_f(u,v) = 0$ which is equivalent to $f(u,v) - f(v,u) = c(u,v)$. Therefore we have:

$$f(S,T) = \sum_{u \in S,\ v \in T} f(u,v) - f(v,u) = \sum_{u \in S,\ v \in T} c(u,v) = c(S,T)$$

□

Since δ is always an integer, following corollary is also true.

Corollary 8.4. *If the capacity of the network c is integer valued function, then there is an integer maximum flow.*

Note that for the algorithm to stop, it was important that c takes integer values. The same argument works for rational c as well, e.g. by multiplying all values with a number to make c integral, but the algorithm might not stop if c has non-rational values. However, Theorem 8.3 holds for any non-negative real capacity function. This can be proved by approximating it with a rational function and taking the limit. Alternatively, Edmonds and Karp showed that the algorithm stops in a bounded number of steps no matter what the capacity function is, provided that it chooses the *shortest* augmenting path of G_f, line 2 of the algorithm.

As an illustration of application of the Ford-Fulkerson theorem, we can use it to prove the theorem by König detailed at the beginning of the chapter.

Proof of König's theorem (Theorem 8.1) via Ford-Fulkerson.

Given $G = (A \cup B, E)$, let us create a network by adding a source s and a sink t: s is connected to each vertex $a \in A$ via an edge $\overrightarrow{sa}$ of capacity 1; t is connected to each vertex $b \in B$ via an edge $\overrightarrow{bt}$ of capacity 1. We also orient all existing edges of G from A to B and assign them a "large enough" integer capacity K, i.e. $K > |A|$.

Let f be a maximum flow. Since all capacities are integers, f is an integer flow by application of Corollary 8.4. We claim that the G-edges with strictly positive flow form a matching. Indeed, each such edge $uv \in E(G)$ has a flow $f(u, v) \geq 1$. If two such edges shared a vertex $a \in A$, then the inflow at a is at most 1 and the outflow is at least 2, which is not possible according to the flow conservation rule. Similarly, two such edges cannot meet in B, either. Hence each edge $uv \in E(G)$ with non-zero flow is such that $f(u, v) = 1$ and thus $\nu(G) \geq |f|$.

Now let (S, T) be a minimum cut. We prove by contradiction that $(S \cap B) \cup (T \cap A)$ is a vertex cover. Assume there exists an uncovered edge $uv \in E(G)$ with $u \in S \cap A$ and $v \in T \cap B$. Then for the capacity of the cut, $c(S, T) \geq c(u, v) = K$. But (S, T) is a minimum cut, thus $c(S, T) \leq c(s, V \setminus \{s\}) \leq |A| < K$ which leads to a contradiction. Furthermore, every s-$(T \cap A)$ edge and every $(S \cap B)$-t edge contributes 1 to the capacity of the S, T cut, so we have: $\tau(G) \leq |S \cap B| + |T \cap A| \leq c(S, T)$. Hence:

$$\tau(G) \leq c(S, T) = |f| \leq \nu(G)$$

Together with the trivial inequality $\nu(G) \leq \tau(G)$, this finishes the proof. □

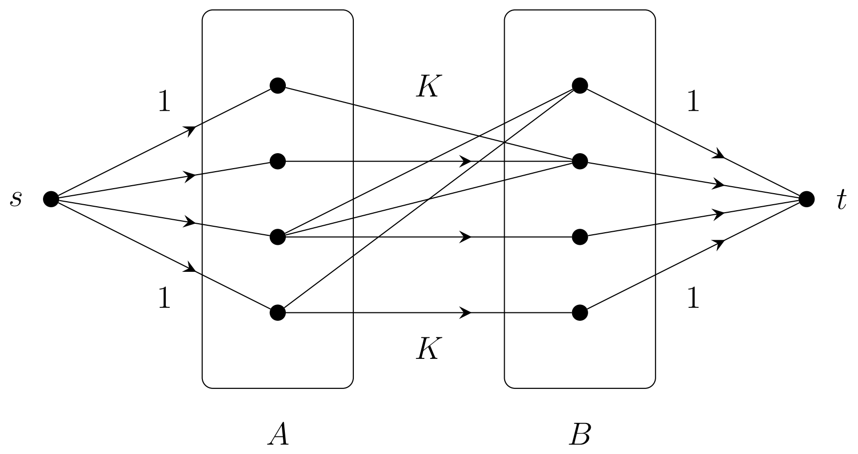

Figure 7: Network created from the bipartite graph G.

Lecture 9

Connectivity.

9.1 Menger's theorem

In the following, we look at what the Ford-Fulkerson theorem says about networks where every edge has capacity 1. The results from this section hold for directed and undirected graphs. We use the following notation: if $G = (V, E)$ is a graph and $W \subseteq V, F \subseteq E$, then $G - W$ is the graph with vertex set $V \setminus W$ and edge set $E \setminus \{e = uv \in E : \{u, v\} \cap W \neq \emptyset\}$, and the graph $G \setminus F$ is the graph with vertex set V and edge set $E \setminus F$.

Definition 36. *Let $G = (V, E)$ be a (directed) graph, and $s, t \in V$.*

We say that two paths s-t paths in G are edge-disjoint*, if they don't share any edges. We say that they are* internally vertex-disjoint*, if they don't share any vertices other than s and t.*

A subset $F \subseteq E$ is an s-t edge separator*, if $G \setminus F$ contains no s-t path. A subset $W \subseteq V$ is an s-t* vertex separator*, if $G - W$ contains no s-t path.*

The following theorem is one of the cornerstones of graph theory.

Theorem 9.1 (Menger - 1927)**.** *In a (directed) graph $G = (V, E)$ with $s, t \in V$:*

1. *Maximum number of edge-disjoint s-t paths = minimum size of an $s - t$ edge separator.*
2. *If $st \notin E$, then:*

   Maximum number of internally vertex-disjoint s-t paths = min size of an $s - t$ vertex separator.

Remember that it was used to prove Chvátal-Erdös theorem (4.10) in lecture 4.

Proof.

"$\leq$" is clear in both statements. Indeed, to suppress all the edge-disjoint (resp. internally vertex-disjoint) *s-t* paths, we need to delete one edge (resp. vertex) from each of them, all of which are distinct. So any *s-t* edge separator (resp. vertex separator) has size at least the number of paths.

Now we show "$\geq$" for *directed G*.

1. Take the network on G where all edges have capacity 1. We will prove the following:

   max # *s-t* edge-disjoint paths $\geq$ max flow = min cut $\geq$ min *s-t* edge separator

   For the first inequality, take a maximum *integer* flow f. If $|f| > 0$, then there is an *s-t* path $\mathcal{P}$ using flow edges. Removing $\mathcal{P}$ from f decreases the value of f by 1. More precisely, we "push back" a capacity-1 flow on $\mathcal{P}$ to decrease $|f|$ by 1. Crucially, as all capacities are 1, the new flow does not use any edge of $\mathcal{P}$. By repeating this step $|f|$ times, we get $|f|$ edge-disjoint *s-t* paths, as desired.

   The equality in the middle is just the max-flow min-cut theorem, Theorem 8.3. For the last inequality, note that the edges appearing in a min cut form an *s-t* edge separator. As the edge capacities are all 1, the capacity of this cut is exactly the number of edges in the edge separator.

2. The key to ensure edge-disjointness in the first statement was to limit the edge capacities to 1. To limit the capacities of vertices, we define the network $\tilde{G}$ based on G as follows (see example Figure 8):

   - for every vertex $v \in V$, we add two vertices v_{in} and v_{out} to $\tilde{G}$, connected by a directed edge $\overrightarrow{v_{in}v_{out}}$ of capacity 1, we denote them "vertex edges";
   - for every directed edge $\overrightarrow{uv} \in E$, we add the edge $\overrightarrow{u_{out}, v_{in}}$ of capacity $+\infty$ to $\tilde{G}$.

   Again, we prove the following inequalities regarding *edges* in $\tilde{G}$ that translates into results regarding *vertices* in G:

   max # edge-disjoint s-t paths $\geq$ max s_{out}-t_{in} flow in $\tilde{G}$ = min cut $\geq$ min s-t edge separator.

   For the first inequality, take a maximum integer s_{out}-t_{in} flow f in $\tilde{G}$. If $|f| > 0$, then there is an s-t path $\mathcal{P}$ using flow edges. As $\overrightarrow{st} \notin E$, $\mathcal{P}$ must use at least one "vertex edge" (v_{in}, v_{out}), and so the capacity of the path $\mathcal{P}$ is 1. This means that pushing back capacity 1 on $\mathcal{P}$ decreases $|f|$ by 1, and ensures that the "vertex edges" appearing in $\mathcal{P}$ are not used by the remaining flow. So repeating this we get $|f|$ paths in $\tilde{G}$, and the corresponding paths in G are internally vertex disjoint.

   The equality in the middle is again the max-flow min-cut theorem. For the last inequality, note that the min cut has finite capacity, because $\overrightarrow{st} \notin E$, so all edges in the cut are "vertex edges" corresponding to a vertex separator. They all have capacity 1, so the capacity of this cut is exactly the number of vertices in the separator.

To prove the undirected variants of Menger's theorem, one can just replace every undirected edge with two directed edges, one in each direction, and apply directed version of the theorem to it. □

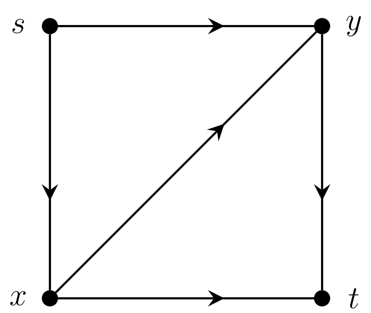

(a) G

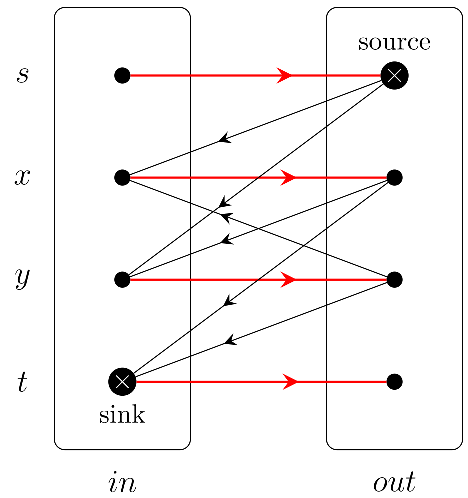

(b) $\tilde{G}$

Figure 8: Example of the network $\tilde{G}$ derived from a directed graph G. Original edges in G are replaced by edges in $\tilde{G}$ (in black) of infinite capacity from *out* to *in* versions of the original end vertices. Capacity 1 edges (in red) connect the *in* to *out* versions of each node.

9.2 Connectivity

For the remaining of this chapter, we consider again undirected graphs.

Definition 37. *$G = (V, E)$ is k-*connected *if $|V| > k$ and $G - X$ is connected for any set X of at most $k - 1$ vertices. The greatest k such that G is k-connected is the* connectivity *$\kappa(G)$. G is k-*edge-connected *if $G \setminus F$ is connected for every set F of at most $k - 1$ edges. The greatest k such that G is k-edge-connected is the* edge-connectivity *$\kappa'(G)$ of G.*

Now we can make some observations:

- G connected $\iff$ G is 1-connected $\iff$ G is 1-edge-connected;
- if G is k-(edge-)connected, then it is $(k - 1)$-(edge-)connected;
- for cliques: $\kappa(K_n) = n - 1$; for cycles: $\kappa(C_n) = \kappa'(C_n) = 2$.

Definition 38. *Let G be a connected graph. A* bridge *is an edge e such that $G \setminus e$ is disconnected. A* cut vertex *is a vertex v such that $G - v$ is disconnected.*

We can observe that:

- G is 2-connected $\iff$ G has no cut vertex;
- G is 2-edge-connected $\iff$ G has no bridge.

Theorem 9.2 (Global version of Menger's theorem)**.** *Let $G = (V, E)$ be a graph.*

1. *G is k-edge-connected $\iff$ G contains k edge-disjoint paths between any two vertices.*
2. *G is k-connected $\iff$ G contains k internally vertex-disjoint paths between any two vertices.*

This is a strong characterization of k-connected graphs that easily implies the following non-trivial facts.

Proposition 9.3. $\kappa(G) \leq \kappa'(G) \leq \delta(G)$.

Proof. $\kappa'(G) \leq \delta(G)$ is trivial, because deleting all edges touching a vertex disconnects G. To show $\kappa(G) \leq \kappa'(G)$, note that by Theorem 9.2, any two vertices of G are connected by $\kappa(G)$ internally vertex-disjoint paths. But then these paths are edge-disjoint, so any two vertices are connected by $\kappa(G)$ edge-disjoint paths, hence G is $\kappa(G)$-edge-connected by the other direction of Theorem 9.2. □

Proposition 9.4. *G is 2-connected $\iff$ any 2 vertices of G are on a cycle.*

Proof. By Theorem 9.2, G is 2-connected if and only if any two vertices are connected by 2 internally vertex-disjoint paths. But this happens if and only if the two vertices are on a cycle, which is the union of the two paths. □

Going back to the global version of Menger's theorem, we need the following lemma to prove the second equivalence regarding vertex connectivity

Lemma 9.5. *Let $G = (V, E)$ be a k-connected graph and s, t two vertices such that $st \in E$. Then $G \setminus st$ is $(k - 1)$-connected.*

Proof. Let us define $G' = G \setminus st$. We reason by contradiction and assume that there is a set X of at most $k-2$ vertices such that $G' - X$ is disconnected. Note that $G - X$ is not disconnected because G is k-connected, and $G' - X = (G - X) \setminus st$, so s and t must be in different components of $G' - X$. Now as G is k-connected, it has at least $k+1$ vertices, so there is a vertex $w \notin X$ and different from s and t. We may assume that w and t are in different components of $G' - X$, if not then w and s are and a similar reasoning can be applied. But now $X' = X \cup \{s\}$ is a set of $k-1$ vertices such that removing X' from G disconnects w and t, because deleting s deletes the edge st, and so $G - X'$ is disconnected. This contradicts G being k-connected. □

We can now we prove the global version of Menger's theorem.

Proof of Theorem 9.2.
We make use of the result from the local version of the theorem.

1. $\Leftarrow$: Let $s, t \in V$ connected by k edge-disjoint paths. Then any s-t edge-separator needs to contain at least one edge from each path. As this is true for any pair of vertices, there is no set F of at most $k-1$ edges such that deleting F disconnects G, thereby separating some two vertices.

   $\Rightarrow$: If G is k-edge-connected, then there is smallest s-t edge separator of size at least k for every $s, t \in V$. Then by Theorem 9.1, s and t are connected by k edge-disjoint paths for any $s, t \in V$.

2. $\Leftarrow$: The above argument works the same with vertices instead of edges.

   $\Rightarrow$: Again the same argument works, except for adjacent pairs $s, t \in V$ because Menger's theorem cannot be applied if $st \in E$. In this case, by applying Lemma 9.5 we know that $G' = G \setminus st$ is $(k-1)$-connected. Then we can apply Theorem 9.1 to G' and get $k-1$ internally vertex-disjoint s-t paths in G'. Together with the edge st, we get k such paths in G, as needed.

□

The following corollary of the global Menger's theorem, called the *fan lemma*, is a nice and highly applicable tool for k-connected graphs.

Lemma 9.6 (fan lemma)**.** *Let $G = (V, E)$ be a k-connected graph. For every $x \in V$ and $U \subseteq V$ with $|U| \geq k$, there are k paths from x to U that are disjoint aside from x, and each has exactly one vertex in U.*

Proof. Add a vertex u that is adjacent to all the vertices in U. The resulting graph is still k-connected. By Theorem 9.2, there are k internally vertex-disjoint paths from x to u. Removing u from these paths gives paths from x to U that are disjoint aside from x. If such a path touches more than one vertex in U, then it contains a subpath with exactly one vertex from U, which we can take instead. □

Lecture 10

Extremal graph theory.

10.1 Triangle-free graphs

Let G be a graph on n vertices that does not contain any triangle as a subgraph, in other words, G is K_3-free. What is the maximum number of edges that G can have?

Bipartite graphs are natural good candidates. With a bit of thinking one can arrive at the conjecture that $K_{\lfloor \frac{n}{2} \rfloor, \lceil \frac{n}{2} \rceil}$ is optimal, i.e. the answer is $\lfloor \frac{n}{2} \rfloor \cdot \lceil \frac{n}{2} \rceil = \lfloor \frac{n^2}{4} \rfloor$. This is indeed the case, as stated by the theorem below.

Theorem 10.1 (Mantel - 1907)**.** *A K_3-free graph on n vertices contains at most $\lfloor \frac{n^2}{4} \rfloor$ edges.*

Proof. Let G be a K_3-free graph with maximum degree Δ, let v be a vertex of such maximum degree and let $S = N(v)$ be its neighborhood, so $|S| = \Delta$. Note that there is no edge with both endpoints in S, otherwise they would form a triangle with v. So every edge of G touches a vertex in $V(G) \setminus S$. Besides, every vertex touches at most Δ edges. Thus the total number of edges is at most:

$$\Delta |V(G) \setminus S| = \Delta(n - \Delta) \leq \lfloor \frac{n^2}{4} \rfloor$$

where we used the fact that $ab \leq \left(\frac{a+b}{2}\right)^2$ for any $a, b \in \mathbb{R}$ to derive the final inequality. □

Example: as an application of Mantel's theorem, we answer the following question by mean of modeling the original problem as a graph problem. Let $a_1, \ldots, a_n \in \mathbb{R}^d$ be vectors such that $\|a_i\| \geq 1$ for each $i \in \{1, \ldots, n\}$. What is the maximum number of pairs satisfying $\|a_i + a_j\| < 1$? There are indeed at most $\lfloor \frac{n^2}{4} \rfloor$ such pairs.

Proof. Define the graph G on $\{1, \ldots, n\}$ where ij is an edge iff $\|a_i + a_j\| < 1$. It is enough to show that G is triangle-free. But this is indeed the case since, for any $i, j, k \in \{1, \ldots, n\}$ with $i \neq j \neq k$:

$$\|a_i + a_j\|^2 + \|a_j + a_k\|^2 + \|a_k + a_i\|^2 = \|a_i + a_j + a_k\|^2 + \|a_i\|^2 + \|a_j\|^2 + \|a_k\|^2 \geq 3$$

So at least one of $\|a_i + a_j\|^2$, $\|a_j + a_k\|^2$, $\|a_k + a_i\|^2$ is larger than 1. □

10.2 Cliques

Instead of triangles, we can ask the same question for arbitrary graphs: for a given graph H, what is the maximum number of edges that an H-free graph on n vertices can have?

Definition 39. *The* extremal number *or* Turán number *of H, denoted* $\mathrm{ex}(n, H)$*, is the maximum value of $|E(G)|$ among all graphs G on n vertices containing no H as a subgraph.*

As a generalization of the triangle-free case, notice that graphs not having K_{r+1} as a subgraph can be obtained by dividing the vertex set V into r pairwise disjoint subsets $V = V_1 \cup \cdots \cup V_r$ with $|V_i| = n_i$ and $n = n_1 + \cdots + n_r$, joining two vertices if and only if they lie in distinct sets V_i and V_j. We denote the resulting graph, called a *complete r-partite graph*, by $K_{n_1, \ldots, n_r}$.

It has $\sum_{i<j} n_i n_j$ edges. For a given $n \in \mathbb{N}^*$, we get the maximum number of edges among such graphs when we distribute the numbers of vertices n_i as evenly as possible among parts, that is $|n_i - n_j| \leq 1$ for all $i \neq j$. Indeed, suppose $n_1 \geq n_2 + 2$. By shifting one vertex from V_1 to V_2, we obtain $K_{n_1-1,n_2+1,\ldots,n_r}$ which contains $(n_1-1)(n_2+1) - n_1 n_2 = n_1 - n_2 - 1 \geq 1$ more edges than $K_{n_1,\ldots,n_r}$. In particular, if r divides n, then choosing $n_i = \frac{n}{r}$ for all i yields a total number of edges of:

$$\binom{r}{2}\left(\frac{n}{r}\right)^2 = \left(1 - \frac{1}{r}\right)\frac{n^2}{2}$$

Turán's theorem states that this is an upper bound for the number of edges in *any* graph on n vertices without an $(r+1)$-clique. This result is considered by many to be the starting point of extremal graph theory. There are many different proofs known. The one we give here is a generalization of our proof for Theorem 10.1.

Definition 40. *We call the graph $K_{n_1,\ldots,n_r}$ with $|n_i - n_j| \leq 1$ for all $i, j \in \{1, \ldots, r\}$, $i \neq j$ the* Turán graph*, denoted by $T(n,r)$.*

Theorem 10.2 (Turán - 1941)**.** *Among all the n-vertex simple graphs with no $(r+1)$-clique, $T(n,r)$ is the unique graph having the maximum number of edges.*

Proof. We apply induction on r. The case $r = 1$ is trivial; the case $r = 2$ is Mantel's theorem. Now assume $r \geq 2$ and let G be a K_{r+1}-free graph on n vertices. Let v be a vertex of maximum degree $\Delta = \Delta(G)$ and let $A = N(v)$ be the neighborhood of v and $B = V(G) \setminus A$ be its complement. By definition $|A| = \Delta$. As G contains no $(r+1)$-clique and v is adjacent to all vertices of A, we note that G contains no K_r with all vertices in A.

In the following, we define A_G and B_G (respectively A_H and B_H) to be the subgraphs induced by G (respectively by H) on the sets of vertices A and B.

Starting from G, we now construct another K_{r+1}-free graph H on $V(G)$ that has at least as many edges as G. We define B_H to be an independent set, i.e. B_H is B_G where all edges $\{xy \in E(G),\ x, y \in B\}$ have been removed. We consider that all possible edges between A_H and B_H exist in H, i.e. all edges in $\{xy,\ x \in A, y \in B\}$. Finally, we define A_H to be isomorphic to $T(\Delta, r-1)$. Then H is a complete r-partite graph. Hence it is K_{r+1}-free.

To see that H has no fewer edges than G, let us define e_B and e_A the number of edges in G respectively touching and not touching B: $e_B = |\{xy \in E(G),\ \{x,y\} \cap B \neq \emptyset\}|$ and $e_A = |\{xy \in E(G),\ x, y \in A\}|$. Then $|E(G)| = e_B + e_A$. Since each vertex of B has degree at most Δ in G, we have: $e_B \leq \Delta |B|$. Besides, $e_A \leq |E(T(\Delta, r-1))|$ by the induction hypothesis and using the fact that G is K_r-free on B. Now observe that H contains exactly $\Delta|B|$ edges touching B and $|E(T(\Delta, r-1))|$ edges not touching B. Therefore $|E(G)| = e_B + e_A \leq |E(H)|$.

This argument shows that $|E(G)| \leq |E(H)|$ for a complete r-partite H. Furthermore, we have seen that $T(n,r)$ maximizes the number of edges among all complete r-partite graphs on n vertices. Thus we have in fact $|E(G)| \leq |E(H)| \leq |E(T(n,r))|$, as needed.

To prove uniqueness, note that equality can only hold in our previous bound if: (i) A induces the complete $r-1$-partite graph $T(\Delta, r-1)$, using the induction hypothesis, and (ii) B touches exactly Δn_r edges in G. But the latter (ii) can only happen if B_G is an independent set in G. Indeed, the sum of the degrees of the vertices in B counts each edge spanned by B twice, and each edge connecting B and A once. As Δ is the maximum degree in G, the sum of degrees is at most Δn_r, so B can only touch this many edges if it spans none of them. But then G is r-partite and since it has the maximum number of edges, $G = T(n,r)$. □

Turán's theorem shows that $\mathrm{ex}(n, K_{r+1}) \leq \left(1 - \frac{1}{r}\right)\frac{n^2}{2}$; there is equality when n is divisible by r. But what happens for other graphs? What is $\mathrm{ex}(n, H)$ for an arbitrary H? Surprisingly the answer pretty much only depends on the chromatic number of H, at least when $\chi(H) \geq 3$. That is what the Erdös-Stone-Simonovits theorem states.

Theorem 10.3 (Erdös, Stone, Simonovits - 1946)**.** *Let H be a graph of chromatic number $\chi(H) = r + 1$ with $r \geq 2$. Then for every $\varepsilon > 0$ and large enough n:*

$$\left(1 - \frac{1}{r}\right) \cdot \binom{n}{2} \leq \mathrm{ex}(n, H) \leq \left(1 - \frac{1}{r}\right) \cdot \binom{n}{2} + \varepsilon n^2$$

10.3 Bipartite graphs

For bipartite H, the Erdös-Stone-Simonovits theorem only says that $\mathrm{ex}(n, H) = o(n^2)$. It is one of the biggest questions in graph theory to determine the order of magnitude of the extremal number of bipartite graphs. There are very few graphs for which we know the answer. One such example is $H = C_4 = K_{2,2}$.

Theorem 10.4. *If a graph G on n vertices contains no $K_{2,2}$, then:*

$$|E(G)| \leq \left\lfloor \frac{n}{4}(1 + \sqrt{4n - 3}) \right\rfloor$$

Proof. Let G be a graph on n vertices without a 4-cycle. Let S be the set elementary configurations defined by $S = \{(u, \{v, w\}) \quad | \quad u, v, w \in V, \quad uv, uw \in E, \ vw \notin E, \quad v \neq w\}$, as illustrated Figure 9.

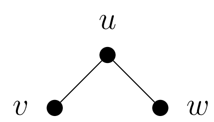

Figure 9: Elementary "cherry-looking" configurations in S

We will count the elements of S of in two different ways. Summing over u, we find $|S| = \sum_{u \in V(G)} \binom{d(u)}{2}$. On the other hand, and this is the crucial observation, every pair $\{v, w\}$ has at most one common neighbor, because G is $K_{2,2}$-free. Hence $|S| \leq \binom{n}{2}$. So far we thus have: $\sum_{u \in V} \binom{d(u)}{2} \leq \binom{n}{2}$, which is equivalent to:

$$\sum_{u \in V} d(u)^2 \leq n(n-1) + \sum_{u \in V} d(u)$$

Now applying the Cauchy-Schwarz inequality to the vectors $(d(u_1), \ldots, d(u_n))$ and $(1, \ldots, 1)$, we get: $\left(\sum_{u \in V} d(u)\right)^2 \leq n \sum_{u \in V} d(u)^2$. This, together with the previous inequality implies:

$$\left(\sum_{u \in V} d(u)\right)^2 \leq n^2(n-1) + n \sum_{u \in V} d(u)$$

Recall that the sum of the degrees is equal to $2|E(G)|$. Therefore we get the following equation in $|E(G)|$:

$$4|E(G)|^2 \leq n^2(n-1) + 2n|E(G)|$$
$$\Leftrightarrow |E(G)|^2 - \frac{n}{2}|E(G)| - \frac{n^2(n-1)}{4} \leq 0$$

Solving this quadratic equation yields the theorem. $\square$

Theorem 10.4 shows that $\text{ex}(n, K_{2,2}) = \mathcal{O}(n^{3/2})$. This upper bound is tight in the sense that there actually exist $K_{2,2}$-free graphs on n vertices with $\Omega(n^{3/2})$ edges.

The proof of this theorem can be generalized to arbitrary $K_{s,t}$, leading to the following theorem.

Theorem 10.5 (Kövári, Sós, Turán - 1954)**.** *For any integers $1 \leq s \leq t$, there is a constant c such that:*

$$\text{ex}(n, K_{s,t}) \leq cn^{2-\frac{1}{s}}$$

However, lower bound constructions where the number of edges has the same order of magnitude is only known for $s = t = 2$ and $s = 2, t = 3$.

For even cycles, the following result is the best known upper bound. Once again, matching constructions are only known for $k = 2$, 3, or 5.

Theorem 10.6 (Bondy, Simonovits - 1974)**.** *For any integer $2 \leq k$, there is a constant c such that:*

$$\text{ex}(n, C_{2k}) \leq cn^{1+\frac{1}{k}}$$

Lecture 11

The probabilistic method.

Probabilistic tools turn out to be extremely useful in discrete mathematics. Indeed certain combinatorial problems that have, a priori, nothing to do with probability sometimes have simple solutions when introducing some randomness. The general idea is of the so-called *probabilistic method* to prove the existence of an object with a given property is to define a probability space on a larger non-empty set of objects and prove that one element of this space has the desired property with non-nil probability. In graph theory more particularly, the application of the method will often consists in defining a probability space on the set of vertices, edges, or specific subsets like partitions, cuts, etc., with carefully selected probabilities. A historical milestone that can be regarded as the inception of the probabilistic method is the proof by Erdös of its theorem stating the existence of graphs with simultaneously large chromatic number and girth. Deeper theoretical results, like the Szemerédi regularity lemma, actually show some inherent connection between graphs and random structures.

11.1 Review of basic notions

We will use standard notions from probability theory, although in a quite restricted setting: we will always work with a probability space that is defined on a *finite* base set Ω with a probability mass function $p : \Omega \to [0, 1]$ satisfying $\sum_{\omega\in\Omega} p(\omega) = 1$. The *events* in this probability space are all subsets $\mathcal{E} \subseteq \Omega$, and the *probability* that $\mathcal{E}$ holds is $\mathbb{P}[\mathcal{E}] = \sum_{\omega\in\mathcal{E}} p(\omega)$.

A *random variable* is just a function $X : \Omega \to \mathbb{R}$. The *expected value* or *expectation* of a random variable X is:

$$\mathbb{E}[X] = \sum_{x\in\mathbb{R}} x \cdot \mathbb{P}[X = x] = \sum_{\omega\in\Omega} p(\omega) \cdot X(\omega)$$

A very important fact that can be easily seen from the right-hand side of this equality, is the *linearity of expectation*, i.e. that for any two random variables X and Y we have $\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$ and for any real c, $\mathbb{E}[cX] = c\mathbb{E}[X]$.

Two events $\mathcal{E}_1$ and $\mathcal{E}_2$ are said to be *independent* if $\mathbb{P}[\mathcal{E}_1 \cap \mathcal{E}_2] = \mathbb{P}[\mathcal{E}_1] \cdot \mathbb{P}[\mathcal{E}_2]$. For more than two events, $\mathcal{E}_1, \dots, \mathcal{E}_k$ are independent if for any $I \subset \{1, \dots, k\}$, we have:

$$\mathbb{P}\left[\bigcap_{i\in I} \mathcal{E}_i\right] = \prod_{i\in I} \mathbb{P}[\mathcal{E}_i]$$

Similarly, the random variables $X_1, \dots, X_k$ are independent if for any $x_1, \dots, x_k$ and any $I \subset \{1, \dots, k\}$, we have:

$$\mathbb{P}\left[\bigcap_{i\in I} X_i = x_i\right] = \prod_{i\in I} \mathbb{P}[X_i = x_i]$$

11.2 Direct applications

As a first example, we provide a new proof of property 2.4 previously covered about the existence of a subgraph containing at least half as many edges as the original graph. We provide here a probabilistic argument.

Proposition 11.1. *Any graph G contains a bipartite subgraph H with $|E(H)| \geq |E(G)|/2$.*

Proof. Let us take a random partition of the vertices into parts A and B, where each vertex is independently assigned to A or B with probability $1/2$. So Ω is the set of all partitions. For every edge $e \in E(G)$, let X_e be the indicator random variable for the event that e "crosses" between A and B. Then $X = \sum_{e \in E(G)} X_e$ is a random variable that counts the number of edges crossing between A and B. Thus the expected number of edges crossing is:

$$\mathbb{E}[X] = \mathbb{E}\left[\sum_{e \in E} X_e\right] = \sum_{e \in E} \mathbb{E}[X_e]$$

Now for an edge $e = uv$, we have:

$$\mathbb{E}[X_e] = \mathbb{P}[e \text{ "crossing"}] = \mathbb{P}[u \in A, v \in B] + \mathbb{P}[u \in B, v \in A] = \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{2}$$

Hence $\mathbb{E}[X] = \sum_{e \in E} 1/2 = |E(G)|/2$. But then there must be a partition $\omega \in \Omega$ such that $X(\omega) \geq |E(G)|/2$, otherwise the expectation would be strictly smaller. □

Note that this is only an existence proof. We do not get any information on how to find such a subgraph.

As another application of the probabilistic method, we now prove the following theorem that connects the independence number of a graph to its degree sequence.

Theorem 11.2 (Caro - 1979)**.** *Let G be a graph. Then G contains an independent set of size at least:*

$$\sum_{v \in V(G)} \frac{1}{1 + d(v)}$$

Proof. Let $n = |V(G)|$ and let $v_1, \dots, v_n$ be a random ordering of the vertex set, that is an ordering chosen from the set of all possible ordering with probability $1/n!$. Let S be the set of vertices v_i such that v_i has no neighbor in $\{v_1, \dots, v_{i-1}\}$. Then S is an independent set. Let us calculate the expected value of the size of S.

For a vertex v, let X_v be the indicator random variable that v is in S. That is: $X_v = 1$ if $v \in S$, $X_v = 0$, otherwise. Then $|S| = \sum_{v \in V(G)} X_v$. Now, the probability of the event $v \in S$ is exactly $\frac{1}{1+d(v)}$, because $v \in S$ if v is the first vertex in the given ordering among all vertices in $\{v\} \cup N(v)$. Hence, $\mathbb{E}(X_v) = \frac{1}{1+d(v)}$. By the linearity of expectation, we have:

$$\mathbb{E}(|S|) = \sum_{v \in V(G)} \mathbb{E}(X_v) = \sum_{v \in V(G)} \frac{1}{1 + d(v)}$$

Therefore there exists an ordering of the vertices such that $|S| \geq \sum_{v \in V(G)} \frac{1}{1+d(v)}$, which concludes the proof. □

11.3 Constructions for the Kövári-Sós-Turán theorem

By the Kövári-Sós-Turán theorem 10.5, for every positive integer $s \geq 2$, we have:

$$\mathrm{ex}(n, K_{s,s}) \leq cn^{2-\frac{1}{s}}$$

where $c > 0$ is a constant depending only on s. However, if $s \geq 4$, it is not known whether $\text{ex}(n, K_{s,s}) \geq c' n^{2-\frac{1}{s}}$ also hold with some constant $c' > 0$. This is one of the central open questions in extremal graph theory. In the following, we use the probabilistic method to find graphs with slightly less than $n^{2-\frac{1}{s}}$ edges that do not contain $K_{s,s}$.

Theorem 11.3. *Let $s \geq 2$ be a positive integer. Then:*

$$\text{ex}(n, K_{s,s}) > \frac{1}{16} n^{2-\frac{2}{(s+1)}}$$

Proof. Let $p = \frac{1}{2} n^{\frac{-2}{s+1}}$ and let G be the random graph on an n elements vertex set in which each pair of vertices is connected by an edge independently with probability p. The graph G might contain $K_{s,s}$, but we can make it $K_{s,s}$-free by deleting an edge from each copy of $K_{s,s}$. Let X be the number of copies of $K_{s,s}$ in G. More precisely, X is the number of pairs (A, B) such that $A, B \subset V(G)$, $|A| = |B| = s$, $A \cap B = \emptyset$ and $ab \in E(G)$ for every $a \in A$, $b \in B$.

Let us calculate the expectation of X. For any pair (A, B) such that $A, B \subset V(G)$, $|A| = |B| = s$, $A \cap B = \emptyset$, let $X_{A,B}$ be the indicator random variable that $ab \in E(G)$ for every $a \in A, b \in B$. Then $X = \sum X_{A,B}$. The probability that $ab \in E(G)$ for every $a \in A, b \in B$ is p^{s^2}, as this is the probability that s^2 given edges are present in G. Hence, $\mathbb{E}(X_{A,B}) = p^{s^2}$. The number of such (A, B) pairs is $\binom{n}{s}\binom{n-s}{s} < n^{2s}$. By the linearity of expectation, we get:

$$\mathbb{E}(X) = \sum \mathbb{E}(X_{A,B}) \leq n^{2s} p^{s^2} = \frac{1}{2^{s^2}} n^{2-\frac{2}{s+1}}$$

Let Y be the number of edges of G. Then $\mathbb{E}(Y) = \binom{n}{2} p > \frac{1}{4} n^2 p = \frac{1}{8} n^{2-\frac{2}{s+1}}$.

Consider the random variable $Z = Y - X$. Then we have:

$$\mathbb{E}(Z) = \mathbb{E}(Y) - \mathbb{E}(X) = \frac{1}{8} n^{2-\frac{2}{s+1}} - \frac{1}{2^{s^2}} n^{2-\frac{2}{s+1}} \geq \frac{1}{16} n^{2-\frac{2}{s+1}}$$

Hence there exists a graph G satisfying $Z(G) \geq \frac{1}{16} n^{2-\frac{2}{s+1}}$. Remove an arbitrary edge from each copy of $K_{s,s}$ in G and let H be the resulting graph. Then H is $K_{s,s}$-free, and we removed at most $X(G)$ edges. Therefore we have $E(H) \geq Y(G) - X(G) = Z(G) \geq \frac{1}{16} n^{2-\frac{2}{s+1}}$. □

11.4 Erdös theorem about chromatic number and girth

Recall Erdös Theorem 6.3 covered in a previous lecture about coloring. It states that: *for any integers k and ℓ, there exists a graph G such that $\chi(G) \geq k$ and $g(G) \geq \ell$*, where $g(G)$ denotes the girth of G.

The meaning of this theorem is counter-intuitive: even graphs with large girth, which locally "look" like trees and are therefore 2-colorable, can still be globally "hard" to color, i.e. be properly colored only with a large number of colors. The theorem does not say how to build such graph, constructing it is indeed extremely difficult, it only states that it exists for any arbitrary pair (k, ℓ). To prove the theorem, we only have to find a graph G without "short" cycles, i.e. of length smaller than ℓ, and without "big" independent sets, because they would allow coloring G with fewer than k colors.

We make use of the same type of random graph construction as in the proof of Theorem 11.3. For $n \in \mathbb{N}^*$ and $p \in [0, 1]$, we define $\mathcal{G}(n, p)$ as the set graphs on n vertices that are the possible realisations of the following random process: each pair of vertices u, v is connected by an edge with probability p independently from other pairs.

Now, the idea is to carefully select p such that, for n large enough, there exists with strictly positive probability a graph $G \in \mathcal{G}(n,p)$ featuring the desired properties. If p is small enough, G is unlikely to contain short cycles; if p is large enough, G is unlikely to contain big independent sets.

In the proof, we make use of the following two lemmas and Markov inequality, restated as follows: *let X a random variable on $\mathbb{R}_+$ and $a > 0$, then* $\mathbb{P}[X \geq a] \leq \frac{\mathbb{E}(X)}{a}$.

Lemma 11.4. *The expectation of the number of ℓ-cycles, $3 \leq \ell \leq n$, in $G \in \mathcal{G}(n,p)$ is:*

$$\frac{n(n-1)\dots(n-\ell+1)}{2\ell}p^\ell = \frac{n!}{(n-\ell)!\,2\ell}p^\ell$$

Lemma 11.5. *For any integers n and k such that $n \geq k \geq 2$, the probability that a graph $G \in \mathcal{G}(n,p)$ has an independent set larger than k is at most:*

$$\mathbb{P}[\alpha(G) \geq k] \leq \binom{n}{k}(1-p)^{\binom{k}{2}}$$

Proof of Erdös Theorem 6.3.

<u>Step 1: setup</u>

Assume that $\ell \geq 3$. Let $0 < \varepsilon < \frac{1}{\ell}$ and define $p = n^{\varepsilon - 1}$. Let $X(G)$ be the random variable that corresponds to the number of cycles of length at most ℓ in a graph $G \in \mathcal{G}(n,p)$ for a given $n \in \mathbb{N}^*$.

<u>Step 2: there are few short cycles in G</u>

By Lemma 11.4 we can evaluate the expectation of X and give an upper-bound using the fact that $np = n^\varepsilon \geq 1$:

$$\mathbb{E}(X) = \sum_{i=3}^{\ell} \frac{n!}{(n-i)!\,2i}p^i \leq \frac{1}{2}\sum_{i=3}^{\ell} n^i p^i \leq \frac{1}{2}(\ell-2)n^\ell p^\ell$$

Now, applying Markov inequality we have:

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

Since $\ell\varepsilon < 1$, it implies: $\lim_{n\to+\infty} \mathbb{P}\left[X \geq \frac{n}{2}\right] = 0$.

<u>Step 3: large independent sets in G are unlikely</u>

For $n \in \mathbb{N}^*$ and r such that $n \geq r \geq 2$, using Lemma 11.5 we have:

$$\mathbb{P}[\alpha(G) \geq r] \leq \binom{n}{r}(1-p)^{\binom{r}{2}} \leq 2^n(1-p)^{\binom{r}{2}} \leq 2^n \exp^{-p\binom{r}{2}}$$

because $1 - p \leq \exp^{-p}$ for any $p \in \mathbb{R}$.

Now consider $r \geq \frac{1}{2}\frac{n}{k} \geq 2$. Since $\binom{r}{2} = \frac{r(r-1)}{2} \geq \frac{r^2}{4}$ for $r \geq 2$, then $p\binom{r}{2} \geq p\frac{r^2}{4} = \frac{1}{16}\frac{n^{\varepsilon+1}}{k^2}$ and thus the upper-bound on the probability of independent sets larger than r is :

$$\mathbb{P}[\alpha(G) \geq r] \leq 2^n \exp^{-\frac{1}{16}\frac{n^{\varepsilon+1}}{k^2}}$$

Since $\alpha(G)$ is an integer, $\mathbb{P}[\alpha(G) \geq r] = \mathbb{P}[\alpha(G) \geq \frac{1}{2}\frac{n}{k}]$ for $r = \lceil \frac{1}{2}\frac{n}{k} \rceil$. Thus, using the previous upper-bound: $\lim_{n\to+\infty} \mathbb{P}\left[\alpha(G) \geq \frac{1}{2}\frac{n}{k}\right] = 0$.

Step 4: a graph with the desired properties exists

Using the two previous limits, we know that there exists $n_0 \in \mathbb{N}^*$ such that for $n \geq n_0$, $\mathbb{P}\left[X \geq \frac{n}{2}\right] < \frac{1}{2}$ and $\mathbb{P}\left[\alpha(G) \geq \frac{1}{2}\frac{n}{k}\right] < \frac{1}{2}$. For $n \geq n_0$ we thus have:

$$\begin{aligned}
\mathbb{P}\left[(X < \frac{n}{2}) \text{ and } (\alpha(G) < \frac{1}{2}\frac{n}{k})\right] &= 1 - \mathbb{P}\left[\neg\left((X < \frac{n}{2}) \text{ and } (\alpha(G) < \frac{1}{2}\frac{n}{k})\right)\right] \\
&\geq 1 - \mathbb{P}\left[\neg(X < \frac{n}{2})\right] - \mathbb{P}\left[\neg(\alpha(G) < \frac{1}{2}\frac{n}{k})\right] \\
&\geq 1 - \mathbb{P}\left[(X \geq \frac{n}{2})\right] - \mathbb{P}\left[(\alpha(G) \geq \frac{1}{2}\frac{n}{k})\right] \\
&> 0
\end{aligned}$$

Hence there exists a graph $G \in \mathcal{G}(n,p)$ for a given $n \geq n_0$ which contains fewer than $\frac{n}{2}$ cycles of length at most ℓ and such that $\alpha(G) < \frac{1}{2}\frac{n}{k}$.

Now we define the graph H by deleting one vertex on each cycle of length at most ℓ in G. Then $|V(H)| \geq \frac{n}{2}$ and by construction, H does not contain any cycle of length ℓ or smaller, i.e. $g(H) \geq \ell$. Furthermore, $\alpha(H) \leq \alpha(G)$ because H is a subgraph of G. Therefore we have:

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

where the first inequality comes from Property 6.4 and is valid for any graph.

The graph H has the desired properties, $g(H) \geq \ell$ and $\chi(H) \geq k$, which concludes the proof of existence.

□

Lecture 12

Ramsey's theorem.

12.1 Ramsey's theorem. Upper bounds

We say that a clique in an edge-colored graph is *monochromatic* if all the edges in the clique have the same color.

Proposition 12.1. *If the edges of K_6 are colored red and blue, then it contains a monochromatic red K_3 or a monochromatic blue K_3.*

Proof. Pick a vertex v. It has five incident edges, among which there must be three red edges or three blue edges. Without loss of generality, assume there are three red edges vx, vy, vz. If one of the edges xy, xz, yz is red, then it forms a red K_3 together with the two corresponding edges to v. Otherwise, the edges xy, xz, yz are all blue, so they form a blue K_3. □

Theorem 12.2 (Ramsey - 1930). *For every positive integer s, there is an N such that any 2-edge-colored K_N contains a monochromatic clique of size s.*

The smallest such N is called the Ramsey number of s. Below is the general definition.

Definition 41. *The* Ramsey number $R(s,t)$ *is the smallest integer N such that whenever the edges of K_N are colored red or blue, this coloring contains a* red K_s *or a* blue K_t.

Proposition 12.1 thus states that $R(3,3) \leq 6$. In fact, we even have $R(3,3) = 6$, because the edges of K_5 can be colored with red and blue without creating a monochromatic K_3. Indeed, we can take a C_5 in K_5 and color its edges red; the remaining edges form another C_5, which we color blue.

Ramsey's theorem, as stated above, says that $R(s,s)$ is always finite. The original statement is a lot more general, for example it also implies that $R(s,t)$ is finite. However, the original bounds that Ramsey got were quite weak. The main results of this lecture are the following lower and upper bounds on $R(s,s)$. The upper bound serves as a proof for Ramsey's theorem.

For every $s \geq 2$:

$$2^{s/2} \leq R(s,s) \leq 2^{2s}$$

Although these bounds seem far apart from each other, and indeed they are, they are basically the state of the art today, as well: no improvement in the exponent is known for either bound. Closing the gap is one of the biggest open problems in combinatorics.

Let us start with the proof of the upper bound. It is actually easier to prove a stronger result:

Theorem 12.3 (Erdös, Szekeres - 1935). *For every $s, t \geq 3$:*

$$R(s,t) \leq \binom{s+t-2}{s-1}$$

Observe that this readily implies $R(s,s) \leq \binom{2s-2}{s-1} \leq 2^{2s-2}$.

Proof. We apply induction on $s+t$.

For $s=2$ the statement is true, because clearly $R(2,t)=t$. Similarly, for $t=2$, we have $R(s,2)=s$, so the bound in the theorem holds.

Now observe that $\binom{s+t-2}{s-1}=\binom{s+t-3}{s-2}+\binom{s+t-3}{s-1}$. Then for the induction step, it is enough to show the following:

$$R(s,t) \leq R(s-1,t)+R(s,t-1) \tag{1}$$

Let $N=R(s-1,t)+R(s,t-1)$ and consider a coloring of the edges of K_N with red and blue. Pick a vertex v. Then in K_N, there are at least $R(s-1,t)$ red edges or at least $R(s,t-1)$ blue edges incident at v. The cases are symmetric so we can assume without loss of generality that v is an endpoint of $R(s-1,t)$ red edges. Now consider the 2-colored complete graph on the corresponding $R(s-1,t)$ neighbors of v. By the definition of $R(s-1,t)$, this graph has either a red K_{s-1} or a blue K_t. In the first case the red K_{s-1} together with v will form a red K_s, and in the second case we get a blue K_t, so we get the monochromatic clique we were looking for in both cases. This establishes the inequality in (1) and finishes the proof. □

Before proving the lower bound, let us show an interesting application. We need to define *multicolor Ramsey numbers* first.

Definition 42. *The* k-color Ramsey number $R_k(s_1,\ldots,s_k)$ *is the smallest integer* N *such that whenever the edges of* K_N *are colored with* k *colors, there is an* i *such that the coloring contains a* K_{s_i} *in color* i.

The following property implies that $R_k(s_1,\ldots,s_k)$ is finite for any choice of $k,s_1,\ldots,s_k$.

Proposition 12.4. *Let* $k,s_1,\ldots,s_k\in\mathbb{N}^*$ *such that* $k\geq 2$ *and* $s_i\geq 3$ *for* $i\in\{1,\ldots,k\}$. *Then we have:*

$$R_k(s_1,\ldots,s_k) \leq R(s_1,\ldots,s_{k-2},R(s_{k-1},s_k))$$

Corollary 12.5 (Schur - 1916)**.** *For every* k *there exists an* N *such that whenever the numbers* $\{1,\ldots,N\}$ *are* k*-colored, there are* x,y,z *of the same color that satisfy* $x+y=z$.

Proof. We prove that $N=R_k(3,\ldots,3)-1$ works. Let $c:\{1,\ldots,N\}\to\{1,\ldots,k\}$ be an arbitrary k-coloring of the integers. We consider the k-edge-coloring c' of K_{N+1} on the vertex set indexed by $\{1,\ldots,N+1\}$, defined as follows:

$$c'(ij)=c(|i-j|)$$

As this defines a k-edge-colored complete graph on $R_k(3,\ldots,3)$ vertices, it must contain a monochromatic triangle hij; without loss of generality assume that $h<i<j$. Then $x=j-i$, $y=i-h$ and $z=j-h$ are numbers in $\{1,\ldots,N\}$ of color $c(x)=c'(ij)$, $c(y)=y'(hi)$, $c(z)=c'(hj)$. These colors are all the same and, by definition of c' they also satisfy the relationship: $x+y=(j-i)+(i-h)=j-h=z$, which proves the property. □

12.2 Ramsey numbers. Lower bounds

We now proceed to the lower bound. Its proof is one of the first applications of the probabilistic method in combinatorics, and was a major breakthrough at the time.

Theorem 12.6 (Erdös - 1947)**.** *For every $s \geq 3$:*

$$2^{s/2} \leq R(s,s)$$

Proof. We actually prove the following:

$$\text{if } 2^{1-\binom{s}{2}}\binom{n}{s} < 1, \text{ then } R(s,s) > n \tag{2}$$

To see that this is sufficient, we just need to check that $n = 2^{s/2}$ satisfies the inequality on the left. Indeed, it is easy to see that $\binom{n}{s} < \frac{n^s}{s!} < \frac{n^s}{2^{1+s/2}}$ for $s \geq 3$. Hence, for $n = 2^{s/2}$ we obtain:

$$2^{1-\binom{s}{2}}\binom{n}{s} < 2^{1-s^2/2+s/2} \cdot \frac{n^s}{2^{1+s/2}} = 2^{-s^2/2}n^s = (n/2^{s/2})^s = 1$$

Now let us prove (2). Consider $n \in \mathbb{N}^*$ and color the edges of K_n randomly red or blue with probability $1/2$, independently of the other edges. For each set A of s vertices, let X_A be the indicator random variable that A forms a monochromatic clique. Then $X = \sum_A X_A$ is the random variable that counts the number of monochromatic s-cliques in this randomly colored graph.

Then:

$$\mathbb{E}[X_A] = \mathbb{P}[A \text{ is monochromatic}] = \mathbb{P}[A \text{ is blue}] + \mathbb{P}[A \text{ is red}] = 2 \cdot \left(\frac{1}{2}\right)^{\binom{s}{2}}$$

Therefore:

$$\mathbb{E}[X] = \sum_{\substack{A \subset V(K_n) \\ |A|=s}} \mathbb{E}[X_A] = \binom{n}{s} 2^{1-\binom{s}{2}}$$

So if $2^{1-\binom{s}{2}}\binom{n}{s} < 1$, then the expected number of monochromatic s-cliques in the graph is strictly smaller than 1. In particular, there is a coloring where the number of monochromatic s-cliques is zero, which concludes the proof. □

Note that with a slightly tighter condition, we can improve the lower bound a little bit using the same reasoning. Indeed, the following is true:

$$\text{if } 2^{1-\binom{s}{2}}\binom{n}{s} = o(n), \text{ then } R(s,s) > (1-o(1))n \tag{3}$$

To prove this, we can use the same argument to see that the expected number of monochromatic s-cliques in the random coloring above is $o(n)$. Now this is not a good Ramsey graph, but we can easily turn it into one by deleting a vertex from each of these monochromatic s-cliques. This way we only delete $o(n)$ vertices, so $(1-o(1))n$ vertices remain in the graph. On the other hand, it no longer contains any monochromatic s-cliques.

Doing the calculations more carefully, (2) implies $R(s,s) \geq \frac{s}{\sqrt{2}e}2^{s/2}(1+o(1))$, whereas (3) implies $R(s,s) \geq \frac{s}{e}2^{s/2}(1+o(1))$ which improves the lower bound by a factor of $\sqrt{2}$. However, it remains very far from the upper bound.

The definition of the Ramsey number $R(s,t)$ is often stated in terms of graphs without edge colors: if one deletes all red edges from a 2-edge-colored complete graph and only keeps the blue edges, then a red clique will correspond to an independent set and a blue clique corresponds to an actual clique in the graph. With this in mind, the Ramsey number $R(s,t)$ can be equivalently defined as the smallest integer N such that *every graph* on N vertices contains an *independent set* of size s or a *clique* of size t, i.e every N-vertex graph G satisfies $\alpha(G) \geq s$ or $\omega(G) \geq t$.

Our lower bound construction is actually the *Erdös-Rényi* random graph $G(n, 1/2)$, where, more generally, $G(n,p)$ is defined as a random graph where each pair of vertices is connected by an edge with probability p independently of the others, as already introduced in the proof of Theorem 11.3

12.3 An upper bound on $R_k(3)$

We finish the lecture by giving an upper bound on the k-color Ramsey number $R_k(3, \ldots, 3)$ sometimes imply denoted $R_k(3)$, that we used in Schur's theorem. The proof is quite similar to the Erdös-Szekeres upper bound.

Proposition 12.7. *For every $k \geq 2$, we have:*

$$R_k(3) \leq \lfloor e \cdot k! \rfloor + 1$$

Proof. We apply induction on k. For $k = 2$ the right-hand side is 6 and indeed $R_2(3) = R(3,3) \leq 6$.

Now suppose $k > 2$ and take a k-edge-colored complete graph on $\lfloor e \cdot k! \rfloor + 1$ vertices. Note that $\lfloor e \cdot k! \rfloor = \sum_{i=0}^{k} \frac{k!}{i!}$. Indeed, $e \cdot k! = \sum_{i=0}^{\infty} \frac{k!}{i!}$ but here $\sum_{i=0}^{k} \frac{k!}{i!}$ is an integer and $\sum_{i=k+1}^{\infty} \frac{k!}{i!} < 1$. Now, observe the following:

$$\lfloor e \cdot k! \rfloor = \sum_{i=0}^{k} \frac{k!}{i!} = 1 + \sum_{i=0}^{k-1} \frac{k!}{i!} = 1 + k \cdot \sum_{i=0}^{k-1} \frac{(k-1)!}{i!} = 1 + k \lfloor e \cdot (k-1)! \rfloor$$

Now a vertex v in our graph has $\lfloor e \cdot k! \rfloor$ edges touching it, colored with k colors. The above equation then implies that at least $\lfloor e \cdot (k-1)! \rfloor + 1$ of these edges have the same color i. Let S be the set of neighbors of v in color i. If the graph contains an edge in color i between two vertices of S then it forms a K_3 in color i with v. Otherwise, the edges between any two vertices in S are colored with $k-1$ colors only, and since $|S| \geq \lfloor e \cdot (k-1)! \rfloor + 1$, we get a monochromatic K_3 by induction. $\square$

Lecture 13

Linear algebra techniques in graph theory.

13.1 Adjacency matrices

If $G = (V, E)$ is a graph on n vertices, the *adjacency matrix* of G is the $n \times n$ sized matrix $A_G \in \mathbb{R}^{n\times n}$, where $A_G(x, y) = 1$ if $xy \in E$, and $A_G(x, y) = 0$ otherwise. Also, $A_G(x, x) = 0$. By definition, A_G is a symmetric binary matrix, i.e. with entries in $\{0, 1\}$, with a 0 leading diagonal. Furthermore, the adjacency matrix has the following properties:

- $\mathrm{Tr}(A_G) = 0$;
- every eigenvalue of A_G is real and A_G has an orthonormal system of eigenvectors;
- if G is d-regular, then $(1, \ldots, 1)$ is an eigenvector of G with eigenvalue d;
- for every $u, w \in \mathbb{R}^n$, we have $u^T A_G w = \sum_{xy\in E(G)} u(x)w(y) + u(y)w(x)$.

Let us now observe what are the eigenvectors and eigenvalues of complete graphs and complete bipartite graphs.

Proposition 13.1. *Let G be the complete graph on n vertices. Then $n-1$ is an eigenvalue of G with eigenvector $(1, ..., 1)$, and -1 is an eigenvalue with multiplicity $n-1$ and eigenspace $U = \{u \in \mathbb{R}^n : \sum_{x\in V(G)} u(x) = 0\}$.*

Proposition 13.2. *Let G be a complete bipartite graph with vertex classes A and B, $|A| = s$ and $|B| = t$. Then 0 is an eigenvalue with multiplicity $n-2$, and $\sqrt{st}$ and $-\sqrt{st}$ are the remaining two eigenvalues. In particular, rank$(A_G) = 2$.*

13.2 Partitioning into complete bipartite graphs

It can be proven that the edges of the complete graph on n vertices can be covered using at most $\lceil \log_2(n)\rceil$ complete bipartite graphs. But how many complete bipartite graphs do we need in order to partition this set of edges, i.e. without edges from different bipartite graphs overlapping? Surprisingly, it turns out that we need much more of them.

Theorem 13.3. *Let G be the complete graph on n vertices and let $B_1, \ldots, B_k$ be edge disjoint complete bipartite graphs with $V(B_i) \subseteq V(G)$ for $i \in \{1, \ldots, k\}$ whose union contains every edge of G. Then $k \geq n-1$.*

Proof. By the conditions of the theorem, we have $A_G = A_{B_1} + \cdots + A_{B_k}$. For every $u \in \mathbb{R}^n$, consider $u^T A_G u$. On one hand, we have:

$$u^T A_G u = \sum_{x,y\in V, x\neq y} u(x)u(y) = \left(\sum_{x\in V} u(x)\right)^2 - \sum_{x\in V} u(x)^2 \tag{4}$$

On the other hand, if X_i and Y_i are the vertex classes of B_i for all $i = 1, \ldots, k$, we have:

$$u^T A_G u = \sum_{i=1}^{k} u^T A_{B_i} u = \sum_{i=1}^{k} \sum_{x \in X_i, y \in Y_i} 2u(x)u(y) = \sum_{i=1}^{k} 2 \left(\sum_{x \in X_i} u(x) \right) \left(\sum_{y \in Y_i} u(y) \right) \quad (5)$$

Now, $U = \{u \in \mathbb{R}^n : \sum_{x \in V(G)} u(x) = 0\}$ and $W_i = \{u \in \mathbb{R}^n : \sum_{x \in X_i} u(x) = 0\}$, for $i \in \{1, \ldots, k\}$ are vector subspaces of $\mathbb{R}^n$ of dimension $n-1$. Therefore the dimension of their intersection is larger than $n - (k+1)$. Assuming that $k < n - 1$, the dimension of the intersection is at least 1. Thus there exists a nonzero vector $u \in \mathbb{R}^n$ such that u satisfies the $k+1$ linear equations:

$$\sum_{x \in V(G)} u(x) = 0 \quad \text{and} \quad \forall i \in \{1, \ldots, k\}, \sum_{x \in X_i} u(x) = 0$$

Plugging this u into (4) we get $u^T A_G u = -\sum_{x \in V} u(x)^2 < 0$, and plugging u into (5) we get $u^T A_G u = 0$, which is a contradiction. □

The bound $k \geq n - 1$ in the theorem is also the best possible as K_n is the disjoint union of $n - 1$ stars.

Let us see what happens if instead of partitioning the edge set, we want to cover every edge an odd number of times by complete bipartite graphs.

Theorem 13.4. *Let G be the complete graph on n vertices and let $B_1, ..., B_k$ be complete bipartite graphs on $V(G)$ that cover each edge of G an odd number of times. Then $k \geq \lfloor \frac{n}{2} \rfloor$.*

Proof. In this proof, all the calculations are done over the field $\mathbb{Z}_2$. For an integer valued matrix A, let $\text{rank}_2(A)$ denote the rank of A over $\mathbb{Z}_2$. Note that $\text{rank}_2(A) \leq \text{rank}(A)$.

Consider the rank of A_G over $\mathbb{Z}_2$. Let $U = \{u \in \mathbb{Z}_2^n : \sum_{x \in V(G)} u(x) = 0\}$. Then U is an $(n-1)$-dimensional subspace of $\mathbb{Z}_2^n$ and $A_G u = u$ for every $u \in U$. Hence, $U \subset \text{Im}(A_G)$, which implies $\text{rank}_2(A_G) \geq n-1$. Now consider $u = (1, 0, \ldots, 0) \notin U$, then $A_G u = (0, 1, \ldots, 1) \notin U$ if n is even. Thus the image of A_G is larger than U, but then $\text{Im}(A_G)$ is $\mathbb{Z}_2^n$ and therefore $\text{rank}_2(A_G) = n$.

Furthermore, for $i = 1, \ldots, k$, as B_i is a complete bipartite graph, hence we have $\text{rank}_2(A_{B_i}) \leq 2$ by Property 13.2. Using the inequality $\text{rank}_2(A+B) \leq \text{rank}_2(A) + \text{rank}_2(B)$, we can write:

$$\text{rank}_2(A_G) = \text{rank}_2 \left(\sum_{i=1}^{k} A_{B_i} \right) \leq \sum_{i=1}^{k} \text{rank}_2 (A_{B_i}) \leq 2k$$

Therefore $k \geq \frac{n}{2}$ if n is even and $k \geq \frac{n-1}{2}$ if n is odd. □

This theorem is also sharp for infinitely many values of n.

13.3 Strongly regular graphs

A graph G is *strongly regular*, if G is regular and there exist two positive integers λ and μ such that: (i) every pair of adjacent vertices have λ common neighbors and (ii) every pair of non-adjacent vertices have μ common neighbors

If G has n vertices and is moreover d-regular, then G is said to be $\text{srg}(n, d, \lambda, \mu)$. For example C_5 is $\text{srg}(5, 2, 0, 1)$ and the Petersen graph is $\text{srg}(10, 3, 0, 1)$. We prove that other than these two graphs, there can be only a few graphs that are $\text{srg}(n, d, 0, 1)$.

Theorem 13.5. *If G is strongly regular on at least three vertices with parameters $\lambda = 0$ and $\mu = 1$, then G is either srg$(5, 2, 0, 1)$, srg$(10, 3, 0, 1)$, srg$(50, 7, 0, 1)$ or srg$(3250, 57, 0, 1)$.*

Proof. Let G be a d-regular graph on n vertices. Assume G is srg$(n, d, 0, 1)$.

First, we show that $n = d^2 + 1$. To do so we count the number of "cherry-looking" configurations, as defined by elements of the set S in the proof of Theorem 10.4, that is the number of pairs $(x, \{y, z\})$, $x \neq y \neq z$ such that $xy, xz \in E(G)$ and $yz \notin E(G)$. Each x appears in exactly $\binom{d}{2}$ such configurations, so $|S| = n\binom{d}{2}$. Also, each pair $\{y, z\}$ appears in exactly one configuration if $yz \notin E(G)$, and in none otherwise, thus $|S| = \binom{n}{2} - \frac{dn}{2}$. From this, we get the equality $n\binom{d}{2} = \binom{n}{2} - \frac{dn}{2}$, which gives $n = d^2 + 1$.

Let us denote $A = A_G$. Now consider A^2 and observe that for any $x, y \in V(G)$, $A_G^2(x, y)$ is the number of common neighbors of x and y. In particular for any $x \in V(G)$, $A(x, x)^2 = d$. Let $I \in \mathbb{R}^{n\times n}$ be the identity matrix, and let $J \in \mathbb{R}^{n\times n}$ be the the matrix whose every entry equals 1. Then, by the definition of A and because G is srg$(n, d, 0, 1)$, we have:

$$A^2 + A = J + (d - 1)I \tag{6}$$

We can use this identity to calculate the eigenvalues of A. Since G is d-regular, d is an eigenvalue with eigenvector $(1, \ldots, 1)$. Let $u \in \mathbb{R}^n$ be an eigenvector of A with eigenvalue $\lambda \neq d$. Then u is orthogonal to $(1, \ldots, 1)$. Multiplying both sides of (6) with u, we get:

$$A^2u + Au = Ju + (d - 1)Iu$$

Now, $Au = \lambda u$ and $A^2u = \lambda^2 u$. Also, $Ju = (0, \ldots, 0)$ because u is orthogonal to $(1, \ldots, 1)$. Finally, $Iu = u$. Therefore:

$$\lambda^2 u + \lambda u = (d - 1)u \ ,$$

from which $\lambda^2 + \lambda = d - 1$. This equation has the two solutions: $\lambda_{1,2} = \frac{-1\pm\sqrt{4d-3}}{2}$. Let us denote n_i the multiplicity of λ_i for $i = 1, 2$. Then $1+n_1+n_2 = n$ and $\mathrm{Tr}(A) = d+n_1\lambda_1+n_2\lambda_2$. Also, by definition of A, $\mathrm{Tr}(A) = 0$. Therefore, we have:

$$n_{1,2} = \frac{1}{2}\left(n - 1 \mp \frac{2d - n + 1}{\sqrt{4d - 3}}\right)$$

Since n_1 and n_2 are integers, $\frac{2d-n+1}{\sqrt{4d-3}}$ must be an integer as well.

Consider as a first case that $2d - n + 1 = 0$. Since $n = d^2 + 1$ also holds, we must have $d = 0$ or $d = 2$. If $d = 0$ then $n = 1$, and if $d = 2$, we have $n = 5$, thus G is srg$(5, 2, 0, 1)$.

Consider now as a second case that $2d - n + 1 \neq 0$. Then $\sqrt{4d - 3}$ must be an integer that divides $2d - n + 1 = 2d - d^2$. Let $s^2 = 4d - 3$, where s is a positive integer. Then $d = \frac{s^2+3}{4}$ and s divides $2d - d^2$. But then s divides $16(2d - d^2)$ as well, and $16(2d - d^2) = 8s^2 + 24 - (s^4 + 6s^2 + 9) = -s^4 + 2s^2 + 15$. Hence s divides 15 and thus $s \in \{1, 3, 5, 15\}$ which implies that $d = 1, 3, 7, 57$. Therefore, G is either srg$(2, 1, 0, 1)$, srg$(10, 3, 0, 1)$, srg$(50, 7, 0, 1)$ or srg$(3250, 57, 0, 1)$. □

There is a construction of a graph G that is srg$(50, 7, 0, 1)$, but it is open whether a graph srg$(3250, 57, 0, 1)$ exists.