

Graph Theory

Benny Sudakov

18 August 2016

Acknowledgement

Much of the material in these notes is from the books *Graph Theory* by Reinhard Diestel and *Introduction to Graph Theory* by Douglas West.

Contents

1	Basic notions	4
1.1	Graphs	4
1.2	Graph isomorphism	4
1.3	The adjacency and incidence matrices	5
1.4	Degree	6
1.5	Subgraphs	7
1.6	Special graphs	7
1.7	Walks, paths and cycles	8
1.8	Connectivity	9
1.9	Graph operations and parameters	10
2	Trees	10
2.1	Trees	10
2.2	Equivalent definitions of trees	11
2.3	Cayley's formula	13
3	Connectivity	17
3.1	Vertex connectivity	17
3.2	Edge connectivity	19
3.3	Blocks	19
3.4	2-connected graphs	21
3.5	Menger's Theorem	21
4	Eulerian and Hamiltonian cycles	24
4.1	Eulerian trails and tours	24
4.2	Hamilton paths and cycles	25
5	Matchings	28
5.1	Real-world applications of matchings	29
5.2	Hall's Theorem	29
5.3	Matchings in general graphs: Tutte's Theorem	31
6	Planar Graphs	34
6.1	Platonic Solids	36
7	Graph colouring	38
7.1	Vertex colouring	38
7.2	Some motivation	39
7.3	Simple bounds on the chromatic number	39
7.4	Greedy colouring	40
7.5	Colouring planar graphs	42
8	More colouring results	43
8.1	Large girth and large chromatic number	44
8.1.1	Digression: the probabilistic method	45
8.1.2	Proof of Theorem 8.5	46
8.2	Chromatic number and clique minors	47
8.3	Edge-colourings	48

8.4	List colouring	50
9	The Matrix Tree Theorem	52
9.1	Lattice paths and determinants	54
10	More Theorems on Hamiltonicity	57
10.1	Pósa's Lemma	59
10.2	Tournaments	60
11	Kuratowski's Theorem	61
11.1	Convex drawings of 3-connected graphs	62
11.2	Reducing the general case to the 3-connected case	65
12	Ramsey Theory	68
12.1	Applications	71
12.2	Bounds on Ramsey numbers	72
12.3	Ramsey theory for integers	72
12.4	Graph Ramsey numbers	73
13	Extremal problems	73
13.1	Turán's theorem	73
13.2	Bipartite Turán Theorems	76

1 Basic notions

1.1 Graphs

Definition 1.1. A *graph* G is a pair $G = (V, E)$ where V is a set of *vertices* and E is a (multi)set of unordered pairs of vertices. The elements of E are called *edges*. We write $V(G)$ for the set of vertices and $E(G)$ for the set of edges of a graph G . Also, $|G| = |V(G)|$ denotes the number of vertices and $e(G) = |E(G)|$ denotes the number of edges.

Definition 1.2. A *loop* is an edge (v, v) for some $v \in V$. An edge $e = (u, v)$ is a *multiple edge* if it appears multiple times in E . A graph is *simple* if it has no loops or multiple edges.

Unless explicitly stated otherwise, we will only consider simple graphs. General (potentially non-simple) graphs are also called *multigraphs*.

Definition 1.3.

- Vertices u, v are *adjacent* in G if $(u, v) \in E(G)$.
- An edge $e \in E(G)$ is *incident* to a vertex $v \in V(G)$ if $v \in e$.
- Edges e, e' are incident if $e \cap e' \neq \emptyset$.
- If $(u, v) \in E$ then v is a *neighbour* of u .


Example 1.4. Any symmetric relation between objects gives a graph. For example:

- let V be the set of people in a room, and let E be the set of pairs of people who met for the first time today;
- let V be the set of cities in a country, and let the edges in E correspond to roads connecting them;
- the internet: let V be the set of computers, and let the edges in E correspond to the links connecting them.

The usual way to picture a graph is to put a dot for each vertex and to join adjacent vertices with lines. The specific drawing is irrelevant, all that matters is which pairs are adjacent.

1.2 Graph isomorphism


Question 1.5.



are these graphs in some sense the same?

Definition 1.6. Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be graphs. An *isomorphism* $\phi : G_1 \rightarrow G_2$ is a bijection (a one-to-one correspondence) from V_1 to V_2 such that $(u, v) \in E_1$ if and only if $(\phi(u), \phi(v)) \in E_2$. We say G_1 is *isomorphic* to G_2 if there is an isomorphism between them.

Example 1.7. Recall the graphs in Question 1.5:



The function $\phi : G_1 \rightarrow G_2$ given by $\phi(1) = a$, $\phi(2) = c$, $\phi(3) = b$, $\phi(4) = d$ is an isomorphism.

Remark 1.8. Isomorphism is an equivalence relation of graphs. This means that

- Any graph is isomorphic to itself
- if G_1 is isomorphic to G_2 then G_2 is isomorphic to G_1
- If G_1 is isomorphic to G_2 and G_2 is isomorphic to G_3 , then G_1 is isomorphic to G_3 .

Definition 1.9. An *unlabelled graph* is an isomorphism class of graphs. In the previous example G_1 and G_2 are different labelled graphs but since they are isomorphic they are the same unlabelled graph.


1.3 The adjacency and incidence matrices

Let $[n] = \{1, \dots, n\}$.

Definition 1.10. Let $G = (V, E)$ be a graph with $V = [n]$. The *adjacency matrix* $A = A(G)$ is the $n \times n$ symmetric matrix defined by

$$a_{ij} = \begin{cases} 1 & \text{if } (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases}$$

Example 1.11.




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

Remark 1.12. Any adjacency matrix A is real and symmetric, hence the spectral theorem proves that A has an orthogonal basis of eigenvalues with real eigenvectors. This important fact allows us to use *spectral methods* in graph theory. Indeed, there is a large subfield of graph theory called *spectral graph theory*.

Definition 1.13. Let $G = (V, E)$ be a graph with $V = \{v_1, \dots, v_n\}$ and $E = \{e_1, \dots, e_m\}$. Then the *incidence matrix* $B = B(G)$ of G is the $n \times m$ matrix defined by

$$b_{ij} = \begin{cases} 1 & \text{if } v_i \in e_j, \\ 0 & \text{otherwise.} \end{cases}$$

Example 1.14.




Remark 1.15. Every column of B has $|e| = 2$ entries 1.

1.4 Degree

Definition 1.16. Given $G = (V, E)$ and a vertex $v \in V$, we define the *neighbourhood* $N(v)$ of v to be the set of neighbours of v . Let the *degree* $d(v)$ of v be $|N(v)|$, the number of neighbours of v . A vertex v is *isolated* if $d(v) = 0$.

Remark 1.17. $d(v)$ is the number of 1s in the row corresponding to v in the adjacency matrix $A(G)$ or the incidence matrix $B(G)$.

Example 1.18.



Fact 1. For any graph G on the vertex set $[n]$ with adjacency and incidence matrices A and B , we have $BB^T = D + A$, where

$$D = \begin{pmatrix} d(1) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & d(n) \end{pmatrix}.$$

Notation 1.19. The *minimum degree* of a graph G is denoted $\delta(G)$; the *maximum degree* is denoted $\Delta(G)$. The *average degree* is

$$\bar{d}(G) = \frac{\sum_{v \in G} d(v)}{|V(G)|}.$$

Note that $\delta \leq \bar{d} \leq \Delta$.

Definition 1.20. A graph G is d -regular if and only if all vertices have degree d .

Question 1.21. Is there a 3-regular graph on 9 vertices?

Proposition 1.22. For every $G = (V, E)$, $\sum_{v \in G} d(v) = 2|E|$.

Proof. In the sum $\sum_{v \in G} d(v)$ every edge $e = (u, v)$ is counted twice: once from u and once from v . □

Corollary 1.23. Every graph has an even number of vertices of odd degree.


This shows that the answer to Question 1.21 is “no”.

1.5 Subgraphs

Definition 1.24. A graph $H = (U, F)$ is a *subgraph* of a graph $G = (V, E)$ if $U \subseteq V$ and $F \subseteq E$. If $U = V$ then H is called *spanning*.

Definition 1.25. Given $G = (V, E)$ and $U \subseteq V$ ($U \neq \emptyset$), let $G[U]$ denote the graph with vertex set U and edge set $E(G[U]) = \{e \in E(G) : e \subseteq U\}$. (We include all the edges of G which have both endpoints in U). Then $G[U]$ is called the subgraph of G *induced* by U .


Example 1.26.



1.6 Special graphs

- K_n is the *complete graph*, or a *clique*. Take n vertices and all possible edges connecting them.
- An *empty graph* has no edges.
- $G = (V, E)$ is *bipartite* if there is a partition $V = V_1 \cup V_2$ into two disjoint sets such that each $e \in E(G)$ intersects both V_1 and V_2 .
- $K_{n,m}$ is the *complete bipartite graph*. Take $n + m$ vertices partitioned into a set A of size n and a set B of size m , and include every possible edge between A and B .

Example 1.27.




1.7 Walks, paths and cycles

Definition 1.28. A *walk* in G is a sequence of vertices $v_0, v_1, v_2, \dots, v_k$, and a sequence of edges $(v_i, v_{i+1}) \in E(G)$. A walk is a *path* if all v_i are distinct. If for such a path with $k \geq 2$, (v_0, v_k) is also an edge in G , then $v_0, v_1, \dots, v_k, v_0$ is a *cycle*. For multigraphs, we also consider loops and pairs of multiple edges to be cycles.

Definition 1.29. The *length* of a path, cycle or walk is the number of edges in it.


Example 1.30.



Proposition 1.31. Every walk from u to v in G contains a path between u and v .

Proof. By induction on the length ℓ of the walk $u = u_0, u_1, \dots, v_\ell = v$.


If $\ell = 1$ then our walk is also a path. Otherwise, if our walk is not a path there is $u_i = u_j$ with $i < j$, then $u = u_0, \dots, u_i, u_{j+1}, \dots, v$ is also a walk from u to v which is shorter. We can use induction to conclude the proof.



□

Proposition 1.32. Every G with minimum degree $\delta \geq 2$ contains a path of length δ and a cycle of length at least $\delta + 1$.

Proof. Let v_1, \dots, v_k be a longest path in G . Then all neighbours of v_k belong to v_1, \dots, v_{k-1} so $k-1 \geq \delta$ and $k \geq \delta+1$, and our path has at least δ edges. Let i ($1 \leq i \leq k-1$) be the minimum index such that $(v_i, v_k) \in E(G)$. Then the neighbours of v_k are among v_i, \dots, v_{k-1} , so $k-i \geq \delta$. Then v_i, v_{i+1}, \dots, v_k is a cycle of length at least $\delta+1$.



□

Remark 1.33. Note that we have also proved that a graph with minimum degree $\delta \geq 2$ contains cycles of at least $\delta-1$ different lengths. This fact, and the statement of Proposition 1.32, are both tight; to see this, consider the complete graph $G = K_{\delta+1}$.

1.8 Connectivity

Definition 1.34. A graph G is *connected* if for all pairs $u, v \in G$, there is a path in G from u to v .


Note that it suffices for there to be a walk from u to v , by Proposition 1.31.

Example 1.35.



Definition 1.36. A (*connected*) *component* of G is a connected subgraph that is maximal by inclusion. We say G is *connected* if and only if it has one connected component.

Example 1.37.



has 4 connected components.

Proposition 1.38. A graph with n vertices and m edges has at least $n-m$ connected components.

Proof. Start with the empty graph (which has n components), and add edges one-by-one. Note that adding an edge can decrease the number of components by at most 1. □

1.9 Graph operations and parameters


Definition 1.39. Given $G = (V, E)$, the *complement* \overline{G} of G has the same vertex set V and $(u, v) \in E(\overline{G})$ if and only if $(u, v) \notin E(G)$.

Example 1.40.



Definition 1.41. A *clique* in G is a complete subgraph in G . An *independent set* is an empty induced subgraph in G .

Example 1.42.



Notation 1.43. Let $\omega(G)$ denote the number of vertices in a maximum-size clique in G ; let $\alpha(G)$ denote the number of vertices in a maximum-size independent set in G .

Exercise 2. In Example 1.42, $\omega(G) = 3$ and $\alpha(G) = 2$.

Claim 1.44. A vertex set $U \subseteq V(G)$ is a clique if and only if $U \subseteq V(\overline{G})$ is an independent set.


Corollary 1.45. We have $\omega(G) = \alpha(\overline{G})$ and $\alpha(G) = \omega(\overline{G})$.

2 Trees

2.1 Trees


Definition 2.1. A graph having no cycle is *acyclic*. A *forest* is an acyclic graph; a *tree* is a connected acyclic graph. A *leaf* (or *pendant vertex*) is a vertex of degree 1.

Example 2.2.



Lemma 2.3. *Every finite tree with at least two vertices has at least two leaves. Deleting a leaf from an n -vertex tree produces a tree with $n - 1$ vertices.*

Proof. Every connected graph with at least two vertices has an edge. In an acyclic graph, the endpoints of a maximum path have only one neighbour on the path and therefore have degree 1. Hence the endpoints of a maximum path provide the two desired leaves.



Suppose v is a leaf of a tree G , and let $G' = G \setminus v$. If $u, w \in V(G')$, then no u, w -path P in G can pass through the vertex v of degree 1, so P is also present in G' . Hence G' is connected. Since deleting a vertex cannot create a cycle, G' is also acyclic. We conclude that G' is a tree with $n - 1$ vertices. \square

2.2 Equivalent definitions of trees

Theorem 2.4. *For an n -vertex simple graph G (with $n \geq 1$), the following are equivalent (and characterize the trees with n vertices).*


- (a) G is connected and has no cycles.
- (b) G is connected and has $n - 1$ edges.
- (c) G has $n - 1$ edges and no cycles.
- (d) For every pair $u, v \in V(G)$, there is exactly one u, v -path in G .

To prove this theorem we will need a small lemma.


Definition 2.5. An edge of a graph is a *cut-edge* if its deletion disconnects the graph.

Lemma 2.6. *An edge contained in a cycle is not a cut-edge.*

Proof. Let (u, v) belong to a cycle.



Then any path $x \dots y$ in G which uses the edge (u, v) can be extended to a walk in $G \setminus (u, v)$ as follows:




Proof of Theorem 2.4. We first demonstrate the equivalence of (a), (b), (c) by proving that any two of {connected, acyclic, $n - 1$ edges} implies the third.

(a) \implies (b), (c): We use induction on n . For $n = 1$, an acyclic 1-vertex graph has no edge. For the induction step, suppose $n > 1$, and suppose the implication holds for graphs with fewer than n vertices. Given G , Lemma 2.3 provides a leaf v and states that $G' = G \setminus v$ is acyclic and connected. Applying the induction hypothesis to G' yields $e(G') = n - 2$, and hence $e(G) = n - 1$.

(b) \implies (a), (c): Delete edges from cycles of G one by one until the resulting graph G' is acyclic. By Lemma 2.6, G is connected. By the paragraph above, G' has $n - 1$ edges. Since this equals $|E(G)|$, no edges were deleted, and G itself is acyclic.

(c) \implies (a), (b): Suppose G has k components with orders n_1, \dots, n_k . Since G has no cycles, each component satisfies property (a), and by the first paragraph the i th component has $n_i - 1$ edges. Summing this over all components yields $e(G) = \sum(n_i - 1) = n - k$. We are given $e(G) = n - 1$, so $k = 1$, and G is connected.

(a) \implies (d): Since G is connected, G has at least one u, v -path for each pair $u, v \in V(G)$. Suppose G has distinct u, v -paths P and Q . Let $e = (x, y)$ be an edge in P but not in Q . The concatenation of P with the reverse of Q is a closed walk in which e appears exactly once. Hence, $(P \cup Q) \setminus e$ is an x, y -walk not containing e . By Proposition 1.31, this contains an x, y -path, which completes a cycle with e and contradicts the hypothesis that G is acyclic. Hence G has exactly one u, v -path.



(d) \implies (a): If there is a u, v -path for every $u, v \in V(G)$, then G is connected. If G has a cycle C , then G has two paths between any pair of vertices on C . \square


Definition 2.7. Given a connected graph G , a *spanning tree* T is a subgraph of G which is a tree and contains every vertex of G .

Corollary 2.8.

- (a) Every connected graph on n vertices has at least $n - 1$ edges and contains a spanning tree;
- (b) Every edge of a tree is a cut-edge;
- (c) Adding an edge to a tree creates exactly one cycle.

Proof.

- (a) delete edges from cycles of G one by one until the resulting graph G' is acyclic. By Lemma 2.6, G is connected. The resulting graph is acyclic so it is a tree. Therefore G had at least $n - 1$ edges and contained a spanning tree.
- (b) note that deleting an edge from a tree T on n vertices leaves $n - 2$ edges, so the graph is disconnected by (a).
- (c) Let $u, v \in T$. There is a unique path in T between u and v , so adding an edge (u, v) closes this path to a unique cycle.




□

2.3 Cayley's formula

Question 2.9. What is the number of spanning trees in a labelled complete graph on n vertices?

Example 2.10.




Theorem 2.11 (Cayley's Formula). *There are n^{n-2} trees with vertex set $[n]$.*

We give two proofs of Cayley's formula. In our first proof, we establish a bijection between trees on $[n]$ and sequences in $[n]^{n-2}$.

Definition 2.12 (Prüfer code). Let T be a tree on an ordered set S of n vertices. To compute the Prüfer sequence $f(T)$, iteratively delete the leaf with the smallest label and append the label of its neighbour to the sequence. After $n - 2$ iterations a single edge remains and we have produced a sequence $f(T)$ of length $n - 2$.

Example 2.13.



We compute the Prüfer code for T as follows:

- 7 (delete 2)
- 4 (delete 3)
- 4 (delete 5)
- 1 (delete 4)
- 7 (delete 6)
- 1 (delete 7)

The edge remaining is $(1, 8)$. We then have $f(T) = (7, 4, 4, 1, 7, 1)$.

Proposition 2.14. *For an ordered n -element set S , the Prüfer code f is a bijection between the trees with vertex set S and the sequences in S^{n-2} .*

Proof. We need to show every sequence $(a_1, \dots, a_{n-2}) \in S^{n-2}$ defines a unique tree T such that $f(T) = (a_1, \dots, a_{n-2})$. If $n = 2$, then there is exactly one tree on 2 vertices and the algorithm defining f always outputs the empty sequence, the only sequence of length zero. So the claim clearly holds for $n = 2$.


Now, assume $n > 2$ and the claim holds for all ordered vertex sets S' of size less than n . Consider a sequence $(a_1, \dots, a_{n-2}) \in S^{n-2}$. We need to show that (a_1, \dots, a_{n-2}) can be uniquely produced by the algorithm.

Suppose that the algorithm produces $f(T) = (a_1, \dots, a_{n-2})$ for some tree T . Then the vertices $\{a_1, \dots, a_{n-2}\}$ are precisely those that are not a leaf in T . Indeed, if a vertex v is a leaf in T then it can only appear in $f(T)$ if its neighbour gets deleted during the algorithm. But this would leave v as an isolated vertex, which is impossible. Conversely, if a vertex v is not a leaf then one of its neighbours must be deleted during the algorithm (it cannot be itself deleted before this happens). When this neighbour of v is deleted, v will be added to the Prüfer code for T , so is in $\{a_1, \dots, a_{n-2}\}$.

This implies that the label of the first leaf removed from T is the minimum element of the set $S \setminus \{a_1, \dots, a_{n-2}\}$. Let v be this element. In other words, in every tree T such that $f(T) = (a_1, \dots, a_{n-2})$ the vertex v is a leaf whose unique neighbour is a_1 .

By induction, there is a unique tree T' with vertex set $S \setminus v$ such that $f(T') = (a_2, \dots, a_{n-2})$. Adding the vertex v and the edge (a_1, v) to T' yields the desired unique tree T with $f(T) = (a_1, \dots, a_{n-2})$. \square

Example 2.15. We use the idea of the above proof to compute the tree with Prüfer code 16631.




6631
4 is the smallest leaf





{1, 3, 5, 6, 7} left

631
5 is the smallest leaf




{1, 3, 6, 7} left

31
6 is the smallest leaf




{1, 3, 7} left

1
3 is the smallest leaf



{1, 7} left

Now add an edge between the remaining vertices {1, 7}.



To prove Cayley's formula, just apply Proposition 2.14 with the vertex set $[n]$ (note that there are n^{n-2} sequences in $[n]^{n-2}$).

For our second proof of Cayley's formula we need the following definition.

Definition 2.16. A *directed graph*, or *digraph* for short, is a vertex set and an edge (multi-)set of *ordered* pairs of vertices. Equivalently, a digraph is a (possibly not-simple) graph where each edge is assigned a direction. The *out-degree* (respectively *in-degree*) of a vertex is the number of edges incident to that vertex which point away from it (respectively, towards it).


Proof of Cayley's formula (due to Joyal 1981). We count trees on n vertices which have two distinguished vertices called the "left end" L and the "right end" R , where L and R can coincide. Let t_n

be the number of labelled trees on n vertices, and let T_n be the family of labelled trees with two distinguished vertices L, R . Clearly, $|T_n| = t_n n^2$, and it is thus enough to prove that $|T_n| = n^n$. We'll describe a bijection between the set of all mappings $f : [n] \rightarrow [n]$, and T_n . As the number of such mappings is clearly n^n , the result will follow.

So, let $f : [n] \rightarrow [n]$ be a mapping. We represent f as a directed graph G_f with vertex set $[n]$ and the set of directed edges $E(G_f) = \{(i, f(i)) : 1 \leq i \leq n\}$.

Example.

$$f = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 \\ 7 & 5 & 5 & 9 & 1 & 2 & 5 & 8 & 4 & 7 \end{pmatrix},$$



Observe that G_f is a digraph in which the outdegree of every vertex is exactly one ($f(i)$ is the only out-neighbour of i).

Let us look at a component of G_f (ignoring edge directions for a moment). Since the out-degree of every vertex is exactly one, each such component contains as many edges as vertices and has therefore exactly one cycle (by Corollary 2.8). This is easily seen to be a directed cycle (just follow an edge leaving a current vertex until you hit a previously visited vertex).

Let M be the union of the vertex sets of these cycles. In order to create a tree, we need to get rid of these cycles. It is easy to see that f restricted to M is a bijection; moreover, M is the unique maximal set on which f acts as a bijection.

Let us write


$$f_M = \begin{pmatrix} v_1 & \dots & v_k \\ f(v_1) & \dots & f(v_k) \end{pmatrix},$$

where $v_1 < v_2 < \dots < v_k$ (and $M = \{v_1, v_2, \dots, v_k\}$). This gives us the ordering $(f(v_1), \dots, f(v_k))$. Now we can choose $L = f(v_1)$, $R = f(v_k)$. The tree T corresponding to f is constructed as follows: Draw a (directed) path $f(v_1), f(v_2), \dots, f(v_k)$, and fill in the remaining vertices as in G_f (removing edge directions).

Example (continued).


$$M = \{1, 4, 5, 7, 8, 9\},$$

$$f|_M = \begin{pmatrix} 1 & 4 & 5 & 7 & 8 & 9 \\ 7 & 9 & 1 & 5 & 8 & 4 \end{pmatrix},$$



Reversing the correspondence is easy: given a tree T with two special vertices L and R , look at the unique path P of T connecting L and R . The vertices of P form the set M . Ordering the vertices of M gives us the first line of $f|_M$, the second line is given by the order of the vertices in P , from L to R .

Example (continued).



$$M = \{1, 4, 5, 7, 8, 9\},$$

$$f|_M = \begin{pmatrix} 1 & 4 & 5 & 7 & 8 & 9 \\ 7 & 9 & 1 & 5 & 8 & 4 \end{pmatrix}.$$

The remaining values of f are then filled in accordance with the unique paths from the remaining vertices to P (directing these paths towards P). □

3 Connectivity

3.1 Vertex connectivity

Definition 3.1. A *vertex cut* in a connected graph $G = (V, E)$ is a set $S \subseteq V$ such that $G \setminus S := G[V \setminus S]$ has more than one connected component. A *cut vertex* is a vertex v such that $\{v\}$ is a cut.

Definition 3.2. G is called *k-connected* if $|V(G)| > k$ and if $G \setminus X$ is connected for every set $X \subseteq V$ with $|X| < k$. In other words, no two vertices of G are separated by fewer than k other vertices. Every (non-empty) graph is 0-connected and the 1-connected graphs are precisely the non-trivial connected graphs. The greatest integer k such that G is k -connected is the *connectivity* $\kappa(G)$ of G .

- $G = K_n$: $\kappa(G) = n - 1$
- $G = K_{m,n}$, $m \leq n$: $\kappa(G) = m$. Indeed, let G have bipartition $A \cup B$, with $|A| = m$ and $|B| = n$. Deleting A disconnects the graph. On the other hand, deleting $S \subset V$ with $|S| < m$ leaves both $A \setminus S$ and $B \setminus S$ non-empty and any $a \in A \setminus S$ is connected to any $b \in B \setminus S$. Hence $G \setminus S$ is connected.

-  : $\kappa(G) = 1$. Deleting v disconnects G , so v is a cut vertex.

Proposition 3.3. For every graph G , $\kappa(G) \leq \delta(G)$.

Proof. If G is a complete graph then trivially $\kappa(G) = \delta(G) = |G| - 1$. Otherwise let $v \in G$ be a vertex of minimum degree $d(v) = \delta(G)$. Deleting $N(v)$ disconnects v from the rest of G . \square

Remark 3.4. High minimum degree does not imply connectivity. Consider two disjoint copies of K_n .

Theorem 3.5 (Mader 1972). Every graph of average degree at least $4k$ has a k -connected subgraph.

Proof. For $k \in \{0, 1\}$ the assertion is trivial; we consider $k \geq 2$ and a graph $G = (V, E)$ with $|V| = n$ and $|E| = m$. For inductive reasons it will be easier to prove the stronger assertion that G has a k -connected subgraph whenever

- (i) $n \geq 2k - 1$ and
- (ii) $m \geq (2k - 3)(n - k + 1) + 1$.

(This assertion is indeed stronger, i.e. (i) and (ii) follow from our assumption of $\bar{d}(G) \geq 4k$: (i) holds since $n > \Delta(G) \geq 4k$, while (ii) follows from $m = \frac{1}{2}\bar{d}(G)n \geq 2kn$.)

We apply induction on n . If $n = 2k - 1$, then $k = \frac{1}{2}(n + 1)$, and hence

$$m \geq (n - 2)\frac{n + 1}{2} + 1 = \frac{1}{2}n(n - 1)$$

by (ii). Thus $G = K_n \supseteq K_{k+1}$, proving our claim. We therefore assume that $n \geq 2k$. If v is a vertex with $d(v) \leq 2k - 3$, we can apply the induction hypothesis to $G \setminus v$ and are done. So we assume that $\delta(G) \geq 2k - 2$. If G is itself not k -connected, then there is a separating set $X \subseteq V$ with less than k vertices, such that $G \setminus X$ has two components on the vertex sets V_1, V_2 . Let $G_i = G[V_i \cup X]$, so that $G = G_1 \cup G_2$, and every edge of G is either in G_1 or G_2 (or both). Each vertex in each V_i has at least $\delta(G) \geq 2k - 2$ neighbours in G and thus also in G_i , so $|G_1|, |G_2| \geq 2k - 1$. Note that each $|G_i| < n$, so by the induction hypothesis, if no G_i has a k -connected subgraph then each

$$e(G_i) \leq (2k - 3)(|G_i| - k + 1).$$

Hence,

$$\begin{aligned} m &\leq e(G_1) + e(G_2) \\ &\leq (2k - 3)(|G_1| + |G_2| - 2k + 2) \\ &\leq (2k - 3)(n - k + 1) \quad (\text{since } |G_1 \cap G_2| \leq k - 1), \end{aligned}$$


contradicting (ii). \square

3.2 Edge connectivity

Definition 3.6. A *disconnecting set* of edges is a set $F \subseteq E(G)$ such that $G \setminus F$ has more than one component. Given $S, T \subset V(G)$, the notation $[S, T]$ specifies the set of edges having one endpoint in S and the other in T . An *edge cut* is an edge set of the form $[S, \bar{S}]$, where S is a non-empty proper subset of $V(G)$. A graph is *k-edge-connected* if every disconnecting set has at least k edges. The *edge-connectivity* of G , written $\kappa'(G)$, is the minimum size of a disconnecting set. One edge disconnecting G is called a *bridge*.

Example 3.7.

- $G = K_n: \kappa'(G) = n - 1$.




Remark 3.8. An edge cut is a disconnecting set but not the other way around. However, every minimal disconnecting set is a cut.

Theorem 3.9. $\kappa(G) \leq \kappa'(G) \leq \delta(G)$.

Proof. The edges incident to a vertex v of minimum degree form a disconnecting set; hence $\kappa'(G) \leq \delta(G)$. It remains to show $\kappa(G) \leq \kappa'(G)$. Suppose $|G| > 1$ and $[S, \bar{S}]$ is a minimum edge cut, having size $\kappa'(G)$.

If every vertex of S is adjacent to every vertex of \bar{S} , then $\kappa'(G) = |S||\bar{S}| = |S|(|G| - |S|)$. This expression is minimized at $|S| = 1$. By definition, $\kappa(G) \leq |G| - 1$, so the inequality holds.

Hence we may assume there exists $x \in S, y \in \bar{S}$ with x not adjacent to y . Let T be the vertex set consisting of all neighbours of x in \bar{S} and all vertices of $S \setminus x$ that have neighbours in \bar{S} (illustrated below). Deleting T destroys all the edges in the cut $[S, \bar{S}]$ (but does not delete x or y), so T is a separating set. Now, by the definition of T we can injectively associate at least one edge of $[S, \bar{S}]$ to each vertex in T , so $\kappa(G) \leq |T| \leq |[S, \bar{S}]| = \kappa'(G)$.




□

3.3 Blocks

Definition 3.10. A *block* of a graph G is a maximal connected subgraph of G that has no cut-vertex. If G itself is connected and has no cut-vertex, then G is a block.

Example 3.11. If B is a block of G , then B as a graph has no cut-vertex, but B may contain vertices that are cut vertices of G . For example, the graph drawn below has five blocks; three copies of K_2 , one of K_3 , and one subgraph that is neither a cycle nor a clique.



Remark 3.12. If a block B has at least three vertices, then B is 2-connected. If an edge is a block of G then it is a cut-edge of G .


Proposition 3.13. Two blocks in a graph share at most one vertex.

Proof. A single vertex deletion cannot disconnect either block. If blocks B_1, B_2 share two vertices, then after deleting any single vertex x there remains a path within B_i from every vertex remaining in B_i to each vertex of $(B_1 \cap B_2) \setminus x$. Hence $B_1 \cup B_2$ is a subgraph with no cut-vertex, which contradicts the maximality of the original blocks. \square

Hence the blocks of a graph partition its edge set. When two blocks of G share a vertex, it must be a cut-vertex of G . The interaction between blocks and cut-vertices is described by a special graph.

Definition 3.14. The *block graph* of a graph G is a bipartite graph H in which one partite set consists of the cut-vertices of G , and the other has a vertex b_i for each block B_i of G . We include (v, b_i) as an edge of H if and only if $v \in B_i$.

Example 3.15.



Proposition 3.16. The block graph of a connected graph is a tree.

Proof. Similar to Proposition 3.13. \square


3.4 2-connected graphs

Definition 3.17. Two paths are *internally disjoint* if neither contains a non-endpoint vertex of the other. We denote the length of the shortest path from u to v (the *distance* from u to v) by $d(u, v)$.

Theorem 3.18 (Whitney 1932). *A graph G having at least three vertices is 2-connected if and only if each pair $u, v \in V(G)$ is connected by a pair of internally disjoint u, v -paths in G .*

Proof. When G has internally disjoint u, v -paths, deletion of one vertex cannot separate u from v . Since this is given for every u, v , the condition is sufficient. For the converse, suppose that G is 2-connected. We prove by induction on $d(u, v)$ that G has two internally disjoint u, v paths. When $d(u, v) = 1$, the graph $G \setminus (u, v)$ is connected, since $\kappa'(G) \geq \kappa(G) = 2$. A u, v -path in $G \setminus (u, v)$ is internally disjoint in G from the u, v -path consisting of the edge (u, v) itself.

For the induction step, we consider $d(u, v) = k > 1$ and assume that G has internally disjoint x, y -paths whenever $1 \leq d(x, y) < k$. Let w be the vertex before v on a shortest u, v -path. We have $d(u, w) = k - 1$, and hence by the induction hypothesis G has internally disjoint u, w -paths P and Q . Since $G \setminus w$ is connected, $G \setminus w$ contains a u, v -path R . If this path avoids P or Q , we are finished, but R may share internal vertices with both P and Q . Let x be the last vertex of R belonging to $P \cup Q$. Without loss of generality, we may assume $x \in P$. We combine the u, x -subpath of P with the x, v -subpath of R to obtain a u, v -path internally disjoint from $Q \cup \{(w, v)\}$.



□

Corollary 3.19. *G is 2-connected and $|G| \geq 3$ if and only if every two vertices in G lie on a common cycle.*

3.5 Menger's Theorem

Definition 3.20. Let $A, B \subseteq V$. An A - B path is a path with one endpoint in A , the other endpoint in B , and all interior vertices outside of $A \cup B$. Any vertex in $A \cap B$ is a trivial A - B path.

If $X \subseteq V$ (or $X \subseteq E$) is such that every A - B path in G contains a vertex (or an edge) from X , we say that X *separates* the sets A and B in G . This implies in particular that $A \cap B \subseteq X$.

Theorem 3.21 (Menger 1927). *Let $G = (V, E)$ be a graph and let $S, T \subseteq V$. Then the maximum number of vertex-disjoint S - T paths is equal to the minimum size of an S - T separating vertex set.*

Proof. Obviously, the maximum number of disjoint paths does not exceed the minimum size of a separating set, because for any collection of disjoint paths, any separating set must contain a vertex from each path. So we just need to prove there is an S - T separating set and a collection of disjoint S - T paths with the same size.

We use induction on $|E|$, the case $E = \emptyset$ being trivial. We first consider the case where S and T are disjoint.

Let k be the minimum size of an S - T separating vertex set. Choose $e = (u, v) \in E$. Let $G' = (V, E \setminus e)$. If each S - T separating vertex set in G' has size at least k , then inductively there exist k vertex-disjoint S - T paths in G' , hence in G .

So we can assume that G' has an S - T separating vertex set C of size at most $k - 1$. Then $C \cup \{u\}$ and $C \cup \{v\}$ are S - T separating vertex sets of G of size k .

Since C is a separating set for G' , no component of $G' \setminus C$ has elements from both S and T . Let V_S be the union of components with elements from S , and let V_T be the union of components with elements in T . If we were to add the edge (u, v) to $G' \setminus C$ then there would be a path from S to T (because C does not separate S and T in G). So, without loss of generality $u \in V_S$ and $v \in V_T$.

Now, each S - $(C \cup \{u\})$ separating vertex set B of G' has size at least k , as it is S - T separating in G . Indeed, each S - T path P in G intersects $C \cup \{u\}$. Let P' be the subpath of P that goes from S to the first time it touches $C \cup \{u\}$. If P' ends with a vertex in C , then $u \notin P$ so P' is an S - $(C \cup \{u\})$ path in G' . If P' ends in u , then it is disjoint from C and so by the above it contains only vertices in V_S . So $v \notin P'$ and again P' is an S - $(C \cup \{u\})$ path in G' . In both cases we showed that P' is an S - $(C \cup \{u\})$ path in G' so P intersects B .

So by induction, G' contains k disjoint S - $(C \cup \{u\})$ paths. Similarly, G' contains k disjoint $(C \cup \{v\})$ - T paths. Any path in the first collection intersects any path in the second collection only in C , since otherwise G' contains an S - T path avoiding C .

Hence, as $|C| = k - 1$, we can pairwise concatenate these paths to obtain $k - 1$ disjoint S - T paths. We can finally obtain a k th path by inserting the e between the path ending at u and the path starting at v .


It remains to consider the general situation where S and T might not be disjoint. Let $X = S \cap T$ and apply the theorem with the disjoint sets $S' = S \setminus X$ and $T' = T \setminus X$, in the graph $G' = G \setminus X$. Let k' be the size of a maximal separating set in G' . We can obtain a $(k' + |X|)$ -vertex S - T separating set in G by adding every vertex in X to an S' - T' separating set in G' . Similarly we can obtain a collection of $k' + |X|$ vertex-disjoint S - T paths by adding each vertex in X as a trivial path to a collection of vertex-disjoint S' - T' paths in G' . \square

Corollary 3.22. *For $S \subseteq V$ and $v \in V \setminus S$, the minimum number of vertices distinct from v separating v from S in G is equal to the maximum number of paths forming an v - S fan in G . (that is, the maximum number of $\{v\}$ - S paths which are disjoint except at v).*

Proof. Apply Menger's Theorem with $T = N(v)$. If one of the resulting paths passes through v , it contains a subpath that is also an S - T path but does not pass through v (note that in such a path, v must be preceded and succeeded by a vertex of T). So we have a suitable number of vertex-disjoint S - T paths not including v , and we can append v to each path to give a v - S fan. \square

Definition 3.23. The *line graph* of G , written $L(G)$, is the graph whose vertices are the edges of G , with $(e, f) \in E(L(G))$ when $e = (u, v)$ and $f = (v, w)$ in G (i.e. when e and f share a vertex).

Example 3.24.



Corollary 3.25. Let u and v be two distinct vertices of G .

1. If $(u, v) \notin E$, then the minimum number of vertices different from u, v separating u from v in G is equal to the maximum number of internally vertex-disjoint u - v paths in G .
2. The minimum number of edges separating u from v in G is equal to the maximum number of edge-disjoint u - v paths in G .

Proof. For (i), Apply Menger's Theorem with $S = N(u)$ and $T = N(v)$.

For (ii), Apply Menger's Theorem to the line graph of G , with S as the set of edges adjacent to u and T as the set of edges adjacent to v . □

Theorem 3.26 (Global Version of Menger's Theorem).

1. A graph is k -connected if and only if it contains k internally vertex-disjoint paths between any two vertices.
2. A graph is k -edge-connected if and only if it contains k edge-disjoint paths between any two vertices.

Proof. First we prove (i). if a graph G contains k internally disjoint paths between any two vertices, then $|G| > k$ and G cannot be separated by fewer than k vertices; thus, G is k -connected.


Conversely, suppose that G is k -connected (and, in particular, has more than k vertices) but contains vertices u, v not linked by k internally disjoint paths. By Corollary 3.25, u and v are adjacent; let $G' = G \setminus (u, v)$. Then G' contains at most $k - 2$ internally disjoint u, v -paths. By Corollary 3.25, we can separate u and v in G' by a set X of at most $k - 2$ vertices. As $|G| > k$, there is at least one further vertex $w \notin X \cup \{u, v\}$ in G . Now X separates w in G' from either u or v (say, from u). But then $X \cup \{v\}$ is a set of at most $k - 1$ vertices separating w from u in G , contradicting the k -connectedness of G .

Then, (ii) follows straight from Corollary 3.25. □

4 Eulerian and Hamiltonian cycles

4.1 Eulerian trails and tours

Question 4.1. Which of the two pictures below can be drawn in one go without lifting your pen from the paper?

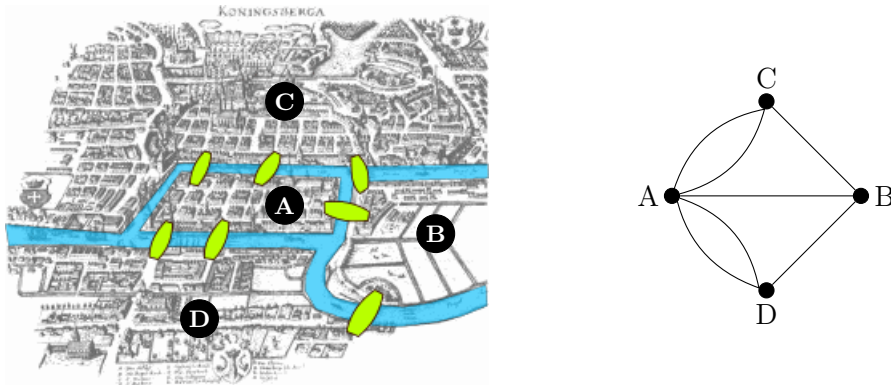


Definition 4.2. A *trail* is a walk with no repeated edges.

Definition 4.3. An *Eulerian trail* in a (multi)graph $G = (V, E)$ is a walk in G passing through every edge exactly once. If this walk is closed (starts and ends at the same vertex) it is called an *Eulerian tour*.

One motivation for this concept is the “7 bridges of Königsberg” problem:

Question 4.4. Is it possible to design a closed walk passing through all the 7 bridges exactly once? Equivalently, does the graph on the right have an Eulerian walk?



Theorem 4.5. A connected (multi)graph has an Eulerian tour if and only if each vertex has even degree.

In order to prove this theorem we use the following lemma.

Lemma 4.6. Every maximal trail in an even graph (i.e., a graph where all the vertices have even degree) is a closed trail.

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 incident to v that is not in T . This edge can be used to extend T to a longer trail, contradicting the maximality of T . \square

Proof of Theorem 4.5. 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, by Lemma 4.6, T is closed, i.e., $v_0 = v_\ell$. If T does not include all the edges of G then, since G is connected, there is an edge e outside of T such that $e = (u, v_i)$ for some vertex v_i in T . But then $T' = ee_{i+1} \dots e_\ell e_1 e_2 \dots e_i$ is a trail in G which is longer than T , contradicting the fact that T is a longest trail in G . Thus, we conclude that T includes all the edges of G and so it is an Eulerian tour. \square

Corollary 4.7. *A connected multigraph G has an Eulerian trail if and only if it has either 0 or 2 vertices of odd degree.*

Proof. Suppose T is an Eulerian trail from vertex u to vertex v . If $u = v$ then T is an Eulerian tour and so by Theorem 4.5 it follows that all the vertices in G have even degree. If $u \neq v$, note that the multigraph $G \cup \{e\}$, where $e = (u, v)$ is a new edge, has an Eulerian tour, namely $T \cup \{e\}$. By Theorem 4.5 it follows that all the degrees in $G \cup \{e\}$ are even. Thus, we conclude that, in the original multigraph G , the vertices u, v are the only ones which have odd degree.


Now we prove the other direction of the corollary. If G has no vertices of odd degree then by Theorem 4.5 it contains an Eulerian tour which is also an Eulerian trail. Suppose now that G has 2 vertices u, v of odd degree. Then $G \cup \{e\}$, where $e = (u, v)$ is a new edge, only has vertices of even degree and so, by Theorem 4.5, it has an Eulerian tour C . Removing the edge e from C gives an Eulerian trail of G from u to v . \square

4.2 Hamilton paths and cycles

Definition 4.8. A Hamilton path/cycle in a graph G is a path/cycle visiting every vertex of G exactly once. A graph G is called Hamiltonian if it contains a Hamilton cycle.

Hamilton cycles were introduced by Kirkman in 1859, and were named after Sir William Hamilton, who produced a puzzle whose goal was to find a Hamilton cycle in a specific graph.


Example 4.9. Hamilton cycle in the skeleton of the 3-dimensional cube.



We give some necessary conditions for Hamiltonicity.

Proposition 4.10. *If G is Hamiltonian then for any set $S \subseteq V$ the graph $G \setminus S$ has at most $|S|$ connected components.*

Proof. Let C_1, \dots, C_k be the components of $G \setminus S$. Imagine that we are moving along a Hamilton cycle in some order, vertex-by-vertex (in the picture below, we are moving clockwise, starting from some vertex in C_1 , say). We must visit each component of $G \setminus S$ at least once; when we leave C_i for the first time, let v_i be the subsequent vertex visited (which must be in S). Each v_i must be distinct because a cycle cannot intersect itself. Hence, S must have at least as many vertices as the number of connected components of $G \setminus S$.




□

Corollary 4.11. *If a connected bipartite graph $G = (V, E)$ with bipartition $V = A \cup B$ is Hamiltonian then $|A| = |B|$.*

Proof. By deleting the vertices in A from G we get $|B|$ isolated vertices and so $G \setminus A$ has $|B|$ connected components. Thus, by Proposition 4.10 we conclude that $|A| \geq |B|$. By symmetry we can also show that $|B| \geq |A|$. Thus, we conclude that $|A| = |B|$. □

Example 4.12. The condition in Proposition 4.10 is not sufficient to ensure that a graph is Hamiltonian. The graph G on the right satisfies the condition of Proposition 4.10 but is not Hamiltonian. Indeed, one would need to include all the edges incident to the vertices v_1, v_2 and v_3 in a Hamilton cycle of G ; however, in that case the vertex u would have degree at least 3 in that Hamilton cycle, which is impossible.




We also give some sufficient conditions for Hamiltonicity.

Theorem 4.13 (Dirac 1952). *If G is a simple graph with $n \geq 3$ vertices and if $\delta(G) \geq n/2$, then G is Hamiltonian.*

Example 4.14. (best-possible minimum degree bound):

- The graph consisting of two cliques of orders $\lfloor (n+1)/2 \rfloor$ and $\lceil (n+1)/2 \rceil$ sharing a vertex has minimum degree $\lfloor (n-1)/2 \rfloor$ but is not Hamiltonian (it is not even 2-connected).




- If n is odd, then the complete bipartite graph $K_{(n-1)/2, (n+1)/2}$ has minimum degree $\frac{n-1}{2}$ but is not Hamiltonian.

Proof of Theorem 4.13. The condition that $n \geq 3$ must be included since K_2 is not Hamiltonian but satisfies $\delta(K_2) = |K_2|/2$.

If there is a non-Hamiltonian graph satisfying the hypotheses, then adding edges cannot reduce the minimum degree, so we may restrict our attention to *maximal* non-Hamiltonian graphs G with minimum degree at least $n/2$. By “maximal” we mean that for every pair (u, v) of non-adjacent vertices of G , the graph obtained from G by adding the edge $e = (u, v)$ is Hamiltonian.

The maximality of G implies that G has a Hamilton path, say from $u = v_1$ to $v = v_n$, because every Hamilton cycle in $G \cup \{e\}$ must contain the new edge e . We use most of this path v_1, \dots, v_n , with a small switch, to obtain a Hamilton cycle in G . If some neighbour of u immediately follows a neighbour of v on the path, say $(u, v_{i+1}) \in E(G)$ and $(v, v_i) \in E(G)$, then G has the Hamilton cycle $(u, v_{i+1}, v_{i+2}, \dots, v_{n-1}, v, v_i, v_{i-1}, \dots, v_2)$ shown below.



To prove that such a cycle exists, we show that there is a common index in the sets S and T defined by $S = \{i : (u, v_{i+1}) \in E(G)\}$ and $T = \{i : (v, v_i) \in E(G)\}$. Summing the sizes of these sets yields

$$|S \cup T| + |S \cap T| = |S| + |T| = d(u) + d(v) \geq n.$$

Neither S nor T contains the index n . This implies that $|S \cup T| < n$, and hence $|S \cap T| \geq 1$, as required. This is a contradiction. \square

Ore observed that this argument uses only that $d(u) + d(v) \geq n$. Therefore, we can weaken the requirement of minimum degree $n/2$ to require only that $d(u) + d(v) \geq n$ whenever u is not adjacent to v .

Theorem 4.15 (Ore 1960). *If G is a simple graph with $n \geq 3$ vertices such that for every pair of non-adjacent vertices u, v of G we have $d(u) + d(v) \geq |G|$, then G is Hamiltonian.*


5 Matchings

Definition 5.1. A set of edges $M \subseteq E(G)$ in a graph G is called a *matching* if $e \cap e' = \emptyset$ for any pair of edges $e, e' \in M$.

A matching is *perfect* if $|M| = \frac{|V(G)|}{2}$, i.e. it covers all vertices of G . We denote the size of the maximum matching in G , by $\nu(G)$.

Example 5.2.

- $G = K_n$: $\nu(G) = \lfloor \frac{n}{2} \rfloor$
- $G = K_{s,t}$, $s \leq t$: $\nu(G) = s$




Remark 5.3. A matching in a graph G corresponds to an independent set in the line graph $L(G)$

Definition 5.4. A set of vertices $T \subseteq V(G)$ of a graph G is called a *cover* of G if every edge $e \in E(G)$ intersects T ($e \cap T \neq \emptyset$), i.e., $G \setminus T$ is an empty graph. Then, $\tau(G)$ denotes the size of the minimum cover.

Example 5.5.

- $G = K_n$: $\tau(G) = n - 1$
- $G = K_{s,t}$, $s \leq t$: $\nu(G) = s$



To see this, note that the graphs induced by the outer 5 vertices and inner 5 vertices are both 5-cycles C_5 . Since $\tau(C_5) = 3$, at least 3 of the outer vertices and 3 of the inner vertices must be included in a vertex cover.

Proposition 5.6. $\nu(G) \leq \tau(G) \leq 2\nu(G)$.

Proof. Let M be a maximum matching in G . Since every cover has at least one vertex on each edge of M and edges are disjoint, we have $\nu(G) \leq \tau(G)$. Note also that since M is maximum, every edge $e \in E(G)$ intersects some edge $e' \in M$, otherwise we get a larger matching. So the vertices covered by M form a cover for G , hence $\tau(G) \leq 2|M| = 2\nu(G)$. \square

5.1 Real-world applications of matchings

Here is just a short list of situations where it is useful to think about and look for matchings.

- Suppose certain workers can operate certain machines, but only one at a time; this gives a bipartite graph between workers and machines. If we want to have many machines operating at the same time, we need a large matching in our bipartite graph.
- Suppose we have a number of hour-long jobs to perform on two computers. Certain jobs can only be started once other jobs are finished. We can define a graph by putting an edge between every pair of jobs that can be performed simultaneously; to finish all the jobs as quickly as possible we would like to find a large matching in our graph.
- The molecular structure of a compound can be described by a graph. For certain kinds of hydrocarbon molecules (*benzenoids*), a perfect matching of this graph gives information about the location of its “double bonds”.
- When students apply to universities, each student has a list of university preferences, and each university also has a list of preferences for students. In order to decide which students should go to which university, we need to find a bipartite matching that is somehow compatible with these preferences. This kind of situation is called the *stable matching* problem, and is extremely important in economics and operations research. The Gale-Shapley algorithm for efficiently computing a stable matching was worth the Nobel prize in economics in 2012.
- Algorithms to find large matchings are essential subroutines for solving optimization problems. The *Chinese postman problem* involves visiting several designated points while travelling as short a total distance as possible. This problem can be efficiently solved by first solving a set of shortest path problems, then solving a certain matching problem.

5.2 Hall’s Theorem

Theorem 5.7 (Hall 1935). *A bipartite graph $G = (V, E)$ with bipartition $V = A \cup B$ has a matching covering A if and only if*

$$|N(S)| \geq |S| \quad \forall S \subseteq A. \tag{1}$$

Proof. It is easy to see that if G has such a matching then (1) holds.

To show the other direction, we apply induction on $|A|$. For $|A| = 1$ the assertion is true. Now let $|A| \geq 2$, and assume that (1) is sufficient for the existence of a matching covering A when $|A|$ is smaller.

If $|N(S)| \geq |S| + 1$ for every non-empty set $S \subsetneq A$, then we pick an edge $(a, b) \in G$ and consider the graph $G' = G \setminus \{a, b\}$ obtained by deleting the vertices a and b . Then every non-empty set $S \subseteq A \setminus \{a\}$ satisfies

$$|N_{G'}(S)| \geq |N_G(S)| - 1 \geq |S|,$$

so by the induction hypothesis G' contains a matching covering $A \setminus \{a\}$. Together with the edge ab , this yields a matching covering A in G .

Suppose now that A has a non-empty proper subset A' with neighbourhood $B' = N(A')$ such that $|A'| = |B'|$. By the induction hypothesis, $G' = G[A' \cup B']$ contains a matching covering

A' . But $G \setminus G'$ satisfies (1) as well: for any set $S \subseteq A \setminus A'$ with $|N_{G \setminus G'}(S)| < |S|$ we would have $|N_G(S \cup A')| = |N_{G \setminus G'}(S)| + |B'| < |S \cup A'|$, contrary to our assumption. Again, by induction, $G \setminus G'$ contains a matching of $A \setminus A'$. Putting the two matchings together, we obtain a matching in G covering A . \square

The next corollary gives a so-called *defect version* of Hall's theorem.

Corollary 5.8. *If in a bipartite graph $G = (A \cup B, E)$ we have $|N(S)| \geq |S| - d$ for every set $S \subseteq A$ and some fixed $d \in \mathbb{N}$, then G contains a matching of cardinality $|A| - d$.*

Proof. We add d new vertices to B , joining each of them to all the vertices in A . Call the resulting graph G' . Note that the new graph has

$$N_{G'}(S) \geq |N_G(S)| + d \geq |S| - d + d = |S|$$

for any $S \subseteq A$, so by Hall's theorem, G' contains a matching of A . At least $|A| - d$ edges in this matching must be edges of G . \square

Corollary 5.9. *If a bipartite graph $G = (A \cup B, E)$ is k -regular with $k \geq 1$, then G has a perfect matching.*

Proof. If G is k -regular, then clearly $|A| = |B|$, since the total number of edges is $k|A| = \sum_{x \in A} d(x) = \sum_{y \in B} d(y) = k|B|$. It thus suffices to show by Theorem 5.7 that G contains a matching covering A . Now every set $S \subseteq A$ is joined to $N(S)$ by a total of $k|S|$ edges, and these are among the $k|N(S)|$ edges of G incident with $N(S)$. Therefore $k|S| \leq k|N(S)|$, so G does indeed satisfy (1). \square

Corollary 5.10. *Every regular graph of positive even degree has a 2-factor (a spanning 2-regular subgraph).*

Proof. Let G be any connected $2k$ -regular graph. By Theorem 4.5 G contains an Euler tour. Define a new graph G' by splitting every vertex v into two vertices v^- and v^+ . If an edge of the Euler tour goes from v to w , put an edge in G' from v^+ to w^- . So, the edges in G and in G' naturally correspond to each other. It is easy to see that G' is bipartite and k -regular so contains a perfect matching. Collapsing each pair of vertices v^-, v^+ back into a single vertex v , a perfect matching of G' corresponds to a 2-factor of G . (Each vertex v is incident to one edge which was incident to v^+ in G' , and one edge incident to v^- in G'). \square


Remark 5.11. A 2-factor is a disjoint union of cycles covering all the vertices of a graph

Definition 5.12. Let A_1, \dots, A_n be a collection of sets. A family $\{a_1, \dots, a_n\}$ is called a *system of distinct representatives* (SDR) if all the a_i are distinct, and $a_i \in A_i$ for all i .

Corollary 5.13. *A collection A_1, \dots, A_n has an SDR if and only if for all $I \subseteq [n]$ we have $|\bigcup_{i \in I} A_i| \geq |I|$.*

Proof. Define a bipartite graph with parts $A = [n]$ and $X = \bigcup_i A_i$ such that (i, a) is an edge if and only if $a \in A_i$. A matching of $[n]$ in this graph corresponds exactly to an SDR, where an edge (i, a) in the matching means that $a_i = a$. But the condition $|\bigcup_{i \in I} A_i| \geq |I|$ for all $I \subseteq [n]$ is precisely Hall's condition for the existence of a matching covering A , so Hall's theorem provides the desired equivalence. \square

Example 5.14. $A_1 = \{1, 2\}$, $A_2 = \{1, 2, 3\}$, $A_3 = \{4\}$, $A_4 = \{1, 3, 4\}$.




Theorem 5.15 (König 1931). *If $G = (A \cup B, E)$ is a bipartite graph, then the maximum size of a matching in G equals the minimum size of a vertex cover of G .*

Proof. We have already seen that a minimum cover has at least the size of a maximum matching. Now take a minimum vertex cover U of G . We construct a matching of size $|U|$ to prove that equality can always be achieved.

Let $R = U \cap A$ and $T = U \cap B$. Let H, H' be the subgraphs of G induced by $R \cup (B \setminus T)$ and $T \cup (A \setminus R)$. We use Hall's theorem to show that H has a complete matching of R into $B \setminus T$ and H' has a complete matching of T into $A \setminus R$. Since these subgraphs are disjoint, the two matchings together form a matching of size $|U|$ in G .

Since $R \cup T$ is a vertex cover, G has no edge from $B \setminus T$ to $A \setminus R$. Suppose $S \subseteq R$, and consider $N_H(S) \subseteq B \setminus T$. If $|N_H(S)| < |S|$, then we can substitute $N_H(S)$ for S in U and obtain a smaller vertex cover, since $N_H(S)$ covers all edges incident to S that are not covered by T . The minimality of U thus implies that Hall's condition holds in H , and hence H has a complete matching of R into $B \setminus T$. Applying the same argument to H' yields the rest of the matching.



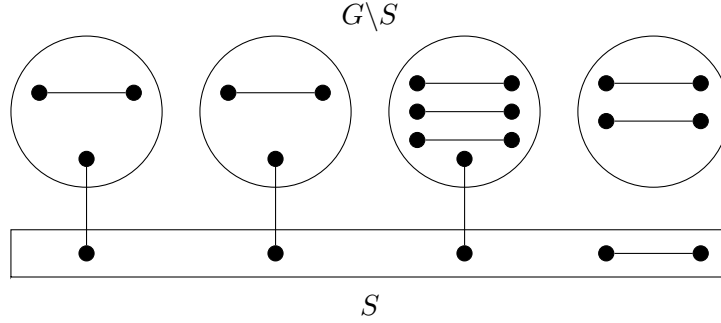
□

5.3 Matchings in general graphs: Tutte's Theorem

Given a graph G , let $q(G)$ denote the number of its *odd components*, i.e. the ones of odd order. If G has a perfect matching then clearly

$$q(G \setminus S) \leq |S| \quad \text{for all } S \subseteq V(G), \quad (2)$$

since every odd component of $G \setminus S$ will send an edge of the matching to S , and each such edge covers a different vertex in S .




Theorem 5.16 (Tutte, 1947). *A graph G has a perfect matching if and only if $q(G \setminus S) \leq |S|$ for all $S \subseteq V(G)$.*

Proof. As noted above, Tutte's condition is necessary; we prove sufficiency. Tutte's condition is preserved by addition of edges: if $G' = G \cup \{e\}$ and $S \subseteq V(G)$, then $q(G' \setminus S) \leq q(G \setminus S)$, because when the addition of e combines two components of $G \setminus S$ into one, the number of components that have odd order does not increase. Therefore, it suffices to consider a simple graph G such that G satisfies (2), has no perfect matching, but adding any edge to G creates a perfect matching. We will obtain a contradiction in every case by constructing a perfect matching in G .

By considering $S = \emptyset$, we know that G has an even number of vertices, since a graph of odd order must have a component of odd order. Let U be the set of vertices in G that are connected to all other vertices. Suppose $G \setminus U$ consists of disjoint complete graphs; we build a perfect matching for such a G . The vertices in each component of $G \setminus U$ can be paired up arbitrarily, with one left over in the odd components. Since $q(G \setminus U) \leq |U|$ and each vertex of U is adjacent to all of $G \setminus U$, we can match these leftover vertices arbitrarily to vertices in U to complete a matching.

This leaves the case where $G \setminus U$ is not a disjoint union of cliques. We can therefore find two vertices in the same component which are not adjacent, and on a shortest path between them there are two nonadjacent vertices x, z at distance 2 (which have a common neighbour y). Furthermore, $G \setminus U$ has another vertex w not adjacent to y , since $y \notin U$. By the maximality of G , adding any edge to G produces a perfect matching. Let M_1 and M_2 be perfect matchings in $G \cup (x, z)$ and $G \cup (y, w)$, respectively. It suffices to show that in $M_1 \cup M_2$ we can find a perfect matching avoiding (x, z) and (y, w) , because that would be contained in G .

Let F be the graph on $V(G)$ with the edges that belong to exactly one of M_1 and M_2 . Note that F contains (x, z) and (y, w) . Since every vertex of G has degree 1 in each of M_1 and M_2 , every vertex of G has degree 0 or 2 in F . Hence F is a collection of disjoint even cycles (alternating between edges of M_1 and M_2) and isolated vertices. Let C be the cycle of F containing (x, z) . If C does not also contain (y, w) , then the desired matching consists of the edges of M_2 from C and all of M_1 not in C . If C contains both (y, w) and (x, z) , as illustrated below, then we use the edge (y, x) or the edge (y, z) to obtain a matching of $V(C)$ using only edges of G (avoiding both (x, z) and (y, w)). Specifically, we use (y, x) if the distance between y and x in C is odd, and we use (y, z) otherwise (then the distance between y and z in C is odd). In the illustration below, this second case applies. The remaining vertices of C form two paths of even order. We use the edges of M_1 in one of these paths and the edges of M_2 in the other to produce a matching in C that does not use (x, z) or (y, w) . (In the illustration below, we use the edges of M_1 on the right side of (y, z) and the edges of M_2 on the left). Combined with M_1 or M_2 outside C , we have a perfect matching of G .




□

Corollary 5.17 (Petersen, 1891). *Every 3-regular graph with no cut-edge has a perfect matching.*

Proof. Let $S \subseteq V(G)$. Let H be a component of $G \setminus S$, with $|H|$ odd. The number of edges between S and H cannot be 1, since G has no cut-edge. It also cannot be even, because then the sum of the vertex degrees in H would be odd. Hence there are at least three edges from H to S .

Since G is 3-regular, each vertex of S is incident to at most three edges between S and $G \setminus S$. Combining this fact with the previous paragraph, we have $3q(G \setminus S) \leq 3|S|$ and hence $q(G \setminus S) \leq |S|$, which proves the corollary. □

Example 5.18. The condition that the graph has no cut-edge is necessary. The graph below is 3-regular but has no perfect matching. Deleting the central vertex leaves 3 odd components.



Finally, we give a defect version of Tutte's theorem.

Corollary 5.19 (Berge 1958). *The largest matching in an n -vertex graph G covers $n + \min_{S \subseteq V(G)} (|S| - q(G \setminus S))$ vertices.*

Proof. Let $d(S) = q(G \setminus S) - |S|$ and let $d = \max_{S \subseteq V} d(S)$. Given $S \subseteq V(G)$, at most $|S|$ edges can match vertices of S to vertices in odd components of $G \setminus S$, so every matching has at least $q(G \setminus S) - |S|$ unmatched vertices. We have shown that no matching can have more than $n - d$ vertices; we want to achieve this bound. Considering the case $S = \emptyset$ shows $d \geq 0$. Let G' be obtained by adding a set D of d vertices to G , each of which are adjacent to every other vertex. Since $d(S)$ has the same parity as $|G|$ for each S , we know that $|G'|$ is even. If G' satisfies Tutte's condition, then we can obtain a matching of the desired size in G from a perfect matching in G' , because deleting D eliminates edges that match at most d vertices of G .


The condition $q(G' \setminus S') \leq |S'|$ holds for $S' = \emptyset$ because $|G'|$ is even. If S' is nonempty but does not contain all of D , then $G' \setminus S'$ has only one component, and $1 \leq |S'|$. Finally, if $D \subseteq S'$, let $S = S' \setminus D$. We have $G' \setminus S' = G \setminus S$, so $q(G' \setminus S') = q(G \setminus S) \leq |S| + d = |S'|$, and G' indeed satisfies Tutte's condition. □

6 Planar Graphs

Definition 6.1. A *polygonal path* or *polygonal curve* in the plane is the union of many line segments such that each segment starts at the end of the previous one and no point appears in more than one segment except for common endpoints of consecutive segments. In a *polygonal u, v -path*, the beginning of the first segment is u and the end of the last segment is v .

A *drawing* of a graph G is a function that maps each vertex $v \in V(G)$ to a point $f(v)$ in the plane and each edge uv to a polygonal $f(u), f(v)$ -path in the plane. The images of vertices are distinct. A point in $f(e) \cap f(e')$ other than a common end is a *crossing*. A graph is *planar* if it has a drawing without crossings. Such a drawing is a *planar embedding* of G . A *plane graph* is a particular drawing of a planar graph in the plane with no crossings.

Example 6.2.



Remark 6.3. We get the same class of graphs if we only require images of edges to be continuous curves. This is because any continuous line can be arbitrarily accurately approximated by a polygonal curve.

Definition 6.4. An *open set* in the plane is a set $U \subset \mathbb{R}^2$ such that for every $p \in U$, all points within some small distance from p belong to U . A *region* is an open set U that contains a polygonal u, v -path for every pair $u, v \in U$ (that is, it is “path-connected”). The *faces* of a plane graph are the maximal regions of the plane that are disjoint from the drawing.

Theorem 6.5 (Jordan curve theorem). *A simple closed polygonal curve C consisting of finitely many segments partitions the plane into exactly two faces, each having C as boundary.*


Remark 6.6. This is not true in three dimensions. In \mathbb{R}^3 there is a surface called the Möbius band which has only one side.

Remark 6.7. The faces of G are pairwise disjoint (they are separated by the edges of G). Two points are in the same face if and only if there is a polygonal path between them which does not cross an edge of G . Also, note that a finite graph has a single unbounded face (the area “outside” of the graph).

Proposition 6.8. *A plane forest has exactly one face.*

Definition 6.9. The *length* of the face f in a planar embedding of G is the length of the walk in G that bounds it.

Example 6.10. The following graph has 3 faces of lengths 6, 3 and 7.



Proposition 6.11. *If $l(f_i)$ denotes the length of a face f_i in a plane graph G , then $2e(G) = \sum l(f_i)$.*

Proof. In the sum $\sum l(f_i)$, every edge was counted twice. □

Theorem 6.12 (Euler's formula 1758). *If a connected plane graph G has exactly n vertices, e edges and f faces, then $n - e + f = 2$.*

Proof. We use induction on the number of edges in G . If $e(G) = n - 1$ and G is connected, then G is a tree. We have $f = 1$, $e = n - 1$. Thus $n - e + f = 2$ holds.

If $e(G) \geq n$ and G is connected, G contains a cycle C . Choose any edge g on C . Let $G' = G \setminus g$. Then G' is connected and $e(G') \geq n - 1$. By the inductive hypothesis, for G' , we have

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

Here $n' = n$ and $e' = e - 1$. Also, deleting g unites two faces, so $f' = f - 1$. Thus,

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

□

Remark 6.13. The fact that deleting an edge in a cycle decreases the number of faces by one can be proved formally using the Jordan curve theorem.

Theorem 6.14. *If G is a planar graph with at least three vertices, then $e(G) \leq 3|G| - 6$. If G is also triangle-free, then $e(G) \leq 2|G| - 4$.*

Proof. It suffices to consider connected graphs; otherwise we could add edges to connect the graph. Also, we will assume there are no leaf vertices. Indeed, all 3 of the non-isomorphic graphs with at most 3 vertices which have a leaf satisfy the theorem, and each time we delete a leaf vertex we only decrease $3|G| - 6 - e(G)$.

In the case where G has no edge, the theorem trivially holds. Otherwise, there is more than one face, and every face boundary in a simple graph contains at least three edges. Let $\{f_i\}$ be the list of face lengths. Then $2e = \sum_i f_i \geq 3f$. Hence $f \leq \frac{2}{3}e$. Substitute this into Euler's formula. We have


$$n - e + \frac{2}{3}e \geq 2,$$

thus $e \leq 3n - 6$.

When G is triangle-free, the faces have length at least 4 (except in the case of K_2). In this case $2e = \sum f_i \geq 4f$, and we obtain $e \leq 2n - 4$. □

Corollary 6.15. *If G is a planar bipartite n -vertex graph with $n \geq 3$ vertices then G has at most $2n - 4$ edges.*

Corollary 6.16. K_5 and $K_{3,3}$ are not planar.




Proof. K_5 is a non-planar graph since $e = 10 > 9 = 3n - 6$. $K_{3,3}$ is a non-planar graph since $e = 9 > 8 = 2n - 4$. \square

Remark 6.17 (Maximal planar graphs / triangulations). The proof of Theorem 6.14 shows that having $3n - 6$ edges in a simple n -vertex planar graph requires $2e = 3f$, meaning that every face is a triangle. If G has some face that is not a triangle, then we can add an edge between non-adjacent vertices on the boundary of this face to obtain a larger plane graph. Hence the simple plane graphs with $3n - 6$ edges, the triangulations, and the *maximal* plane graphs are all the same family.

6.1 Platonic Solids

Definition 6.18. A *polytope* is a solid in 3 dimensions with flat faces, straight edges and sharp corners. Faces of a polytope are joined at the edges. A polytope is *convex* if the line connecting any two points of the polytope lies inside the polytope.

Example 6.19. The tetrahedron:




Definition 6.20. A *regular* or *Platonic* solid is a convex polytope which satisfies the following:

1. all of its faces are congruent regular polygons,
2. all vertices have the same number of faces adjacent to them.

We will now characterise all Platonic solids. The first step is to convert a convex polytope into a planar graph. To do this, we place the considered polytope inside a sphere. Then we project the polytope onto the sphere (imagine that the edges of the polytope are made from wire and we place a tiny lamp in the center). This yields a graph drawn on the sphere without edge crossings.

Now let us show that planar graphs are exactly graphs that can be drawn on the sphere. This becomes quite obvious if we use the *stereographic projection*. We place the sphere in the 3-dimensional space in such a way that it touches the considered plane ρ . Let o denote the point of the sphere lying farthest from ρ , the 'north pole'.



Then the stereographic projection maps each point $x \neq o$ of the sphere to a point x' , where x' is the intersection of the line ox with the plane ρ . (For the point o , the projection is undefined.) This defines a bijection between the plane and the sphere without the point o . Given a drawing of a graph G on the sphere without edge crossings, where the point o lies on no arc of the drawing (which we may assume by a suitable choice of o), the stereographic projection yields a planar drawing of G . Conversely, from a planar drawing we get a drawing on the sphere by the inverse projection.

Corollary 6.21. *If K is a convex polytope with v vertices, e edges and f faces then $v - e + f = 2$.*

Suppose K is a Platonic solid. All its faces are congruent; assume that they have n vertices (and, thus, n edges). Let us assume moreover that each vertex is adjacent to m faces (and, thus, it has m edges adjacent to it). Since each edge is adjacent to exactly two faces,

$$2e = nf. \quad (3)$$

Moreover, each edge is adjacent to two vertices, and one vertex belongs to m edges, thus

$$mv = 2e. \quad (4)$$

Expressing v and f in terms of e , and substituting to Euler's formula, we obtain that $\frac{2e}{m} - e + \frac{2e}{n} = 2$. Rearranging, we arrive at

$$\frac{1}{m} + \frac{1}{n} = \frac{1}{2} + \frac{1}{e}.$$


Note that since K is a 3-dimensional polytope, each of its faces is a polygon and thus has at least 3 vertices, that is $n \geq 3$. Moreover, at each vertex, there are at least three faces meeting; $m \geq 3$. On the other hand, since $e \geq 1$, we must have

$$\frac{1}{m} + \frac{1}{n} > \frac{1}{2}. \quad (5)$$

These conditions do not leave too much leeway; there are only five possible (n, m) pairs for which the above inequality holds. These are $(3, 3)$, $(3, 4)$, $(3, 5)$, $(4, 3)$, $(5, 3)$.

A Platonic solid corresponds to each of these pairs. We list them below.

- **Tetrahedron.** Here $n = 3$ and $m = 3$. Thus, (5) yields that $e = 6$. By (4), $v = 4$, and by (3), $f = 4$. There are 4 vertices and 4 faces of the tetrahedron; the faces are regular triangles, and the vertices are adjacent to 3 edges.
- **Octahedron.** Here $n = 3$ and $m = 4$. Thus, (5) yields that $e = 12$. By (4), $v = 6$, and by (3), $f = 8$. There are 8 vertices and 8 faces of the octahedron; the faces are regular triangles, and the vertices are adjacent to 4 edges.
- **Icosahedron.** Here $n = 3$ and $m = 5$. Thus, (5) yields that $e = 30$. By (4), $v = 12$, and by (3), $f = 20$. There are 12 vertices and 20 faces of the icosahedron; the faces are regular triangles, and the vertices are adjacent to 5 edges.
- **Cube.** Here $n = 4$ and $m = 3$. Thus, (5) yields that $e = 12$. By (4), $v = 8$, and by (3), $f = 6$. There are 8 vertices and 6 faces of the tetrahedron; the faces are squares, and the vertices are adjacent to 3 edges.
- **Dodecahedron.** Here $n = 5$ and $m = 3$. Thus, (5) yields that $e = 30$. By (4), $v = 20$, and by (3), $f = 12$. There are 20 vertices and 12 faces of the tetrahedron; the faces are regular pentagons, and the vertices are adjacent to 3 edges.




7 Graph colouring

7.1 Vertex colouring

Definition 7.1. A k -colouring of G is a labeling $f : V(G) \rightarrow \{1, \dots, k\}$. It is a *proper k -colouring* if $(x, y) \in E(G)$ implies $f(x) \neq f(y)$. A graph G is k -colourable if it has a proper k -colouring. The *chromatic number* $\chi(G)$ is the minimum k such that G is k -colourable. If $\chi(G) = k$, then G is k -chromatic. If $\chi(G) = k$, but $\chi(H) < k$ for every proper subgraph H of G , then G is *colour-critical* or k -critical.

Example 7.2.

- $\chi(K_n) = n$



- The chromatic number of an odd cycle is 3

Remark 7.3. The vertices having a given colour in a proper colouring must form an independent set, so $\chi(G)$ is the minimum number of independent sets needed to cover $V(G)$. Hence G is k -colourable if and only if G is k -partite. Multiple edges do not affect chromatic number. Although we define k -colouring using numbers from $\{1, \dots, k\}$ as labels, the numerical values are usually unimportant, and we may use any set of size k as labels.

7.2 Some motivation

Example 7.4 (examination scheduling). 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 .

Example 7.5 (chemical storage). A company manufactures n chemicals C_1, C_2, \dots, C_n . Certain pairs of these chemicals are incompatible and would cause explosions if brought into contact with each other. As a precautionary measure, the company wishes to divide its warehouse into compartments, and store incompatible chemicals in different compartments. What is the least number of compartments into which the warehouse should be partitioned? We obtain a graph G on the vertex set $\{v_1, v_2, \dots, v_n\}$ by joining two vertices v_i and v_j if and only if the chemicals C_i and C_j are incompatible. It is easy to see that the least number of compartments into which the warehouse should be partitioned is equal to the chromatic number of G .

7.3 Simple bounds on the chromatic number


Claim 7.6. If H is a subgraph of G then $\chi(H) \leq \chi(G)$.

Proof. Note that a proper colouring of G is also a proper colouring of H . □

Corollary 7.7. $\chi(G) \geq \omega(G)$

Proof. Let $\omega(G) = t$. Then G contains a subgraph H which is isomorphic to K_t . Thus, by the claim above it follows that $\chi(G) \geq \chi(H) = t$. □

Example 7.8. Consider the following graph.



In this case we have $\chi(G) = 4$ and $\omega(G) = 3$. Thus, the chromatic number can be bigger than the clique number.

Proposition 7.9. $\chi(G) \geq \frac{|V(G)|}{\alpha(G)}$

Proof. Let $\chi(G) = k$. A k -colouring of $V(G)$ gives a partition $V(G) = V_1 \cup \dots \cup V_k$ such that every V_i is an independent set. Hence, $|V_i| \leq \alpha(G)$. Therefore, $|V(G)| = \sum_{i=1}^k |V_i| \leq k\alpha(G)$. Thus, $k = \chi(G) \geq \frac{|V(G)|}{\alpha(G)}$ as claimed. \square

Claim 7.10. For any graph $G = (V, E)$ and any $U \subseteq V$ we have $\chi(G) \leq \chi(G[U]) + \chi(G[V \setminus U])$.

Proof. Properly colour U using $\chi(G[U])$ colours and properly colour $V \setminus U$ using $\chi(G[V \setminus U])$ other colours. This gives a proper colouring of G in $\chi(G[U]) + \chi(G[V \setminus U])$ colours. \square

Claim 7.11. For any graphs G_1 and G_2 on the same vertex set, $\chi(G_1 \cup G_2) \leq \chi(G_1)\chi(G_2)$.

Proof. Let c_1 and c_2 be colourings of G_1 and G_2 with the integers in $[\chi(G_1)]$ and $[\chi(G_2)]$ respectively. We colour the vertices of $G_1 \cup G_2$ with elements of the set $[\chi(G_1)] \times [\chi(G_2)]$, with the colouring c defined by $c(v) = (c_1(v), c_2(v))$. If v is adjacent to w in $G_1 \cup G_2$ then (v, w) is an edge in one of G_1 or G_2 , so $c_1(v) \neq c_1(w)$ or $c_2(v) \neq c_2(w)$. This proves that $c(v) \neq c(w)$, so c is proper. \square

Proposition 7.12.

(i) $\chi(G)\chi(\overline{G}) \geq |G|$

(ii) $\chi(G) + \chi(\overline{G}) \leq |G| + 1$


Proof. (i) follows from Claim 7.11: we have $\chi(G)\chi(\overline{G}) \geq \chi(G \cup \overline{G}) = \chi(K_{|G|}) = |G|$.

(ii) can be proved by induction on $|G|$ (the case $|G| = 1$ is obvious). So, let $|G| = n + 1$. Let $G_0 = G \setminus v$ for some vertex v . By induction we have $\chi(G_0) + \chi(\overline{G_0}) \leq n + 1$. Let $c : V \rightarrow [k]$ be a colouring of G_0 and $f : V \rightarrow [\ell]$ be a colouring of $\overline{G_0}$, with $k + \ell = n + 1$ (we might be using more colours than are necessary). If $d_G(v) < k$ then there is a colour c_v such that v has no neighbours coloured c_v . We can then colour v with c_v to extend c to a colouring of G with k colours. This would prove $\chi(G) \leq k$, and since $\chi(\overline{G}) = \chi(\overline{G_0} \cup \{v\}) \leq \ell + 1$ we have $\chi(G) + \chi(\overline{G}) \leq k + \ell + 1 \leq n + 2$. Otherwise $d_G(v) \geq k$ so $d_{\overline{G}}(v) \leq n - k = \ell - 1$. We can then use exactly the same reasoning as before to extend f to a colouring of \overline{G} with ℓ colours, and since $\chi(G) \leq k + 1$ we are done again. \square

7.4 Greedy colouring

Definition 7.13. The *greedy colouring* with respect to a vertex ordering v_1, \dots, v_n of $V(G)$ is obtained by colouring vertices in the order v_1, \dots, v_n , assigning to v_i the smallest-indexed colour not already used on its lower-indexed neighbours.

Example 7.14.



This graph has chromatic number 2 but the greedy colouring needs 3 colours.

Definition 7.15. Let $G = (V, E)$ be 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 .

Proposition 7.16. G is k -degenerate if and only if there is an ordering v_1, \dots, v_n of the vertices of G such that each v_i has at most k neighbours among the vertices v_1, \dots, v_{i-1} .

Proof. If there is such an ordering, then for any subgraph H , consider the maximum vertex of H with respect to the ordering. This vertex has at most k neighbours in H , thus proving that G is k -degenerate.

Conversely, suppose G is k -degenerate. We prove the existence of a suitable ordering by induction on the number of vertices. If G is k -degenerate it has a vertex of degree at most k . Call this vertex v_n . Let $G' = G \setminus v_n$ and note that G' is still k -degenerate. Thus, there exists an ordering v_1, \dots, v_{n-1} of the vertices of G' satisfying the assertion of the proposition for G' . Then the ordering v_1, \dots, v_n satisfies the required conditions for G . \square

Definition 7.17. Define $\text{dg}(G)$ to be the minimum k such that G is k -degenerate.

Remark 7.18. $\delta(G) \leq \text{dg}(G) \leq \Delta(G)$.

Theorem 7.19. $\chi(G) \leq 1 + \text{dg}(G)$

Proof. Let $k = \text{dg}(G)$. Fix an ordering v_1, \dots, v_n of $V(G)$ such that each v_i has at most k neighbours among v_1, \dots, v_{i-1} . Use the greedy colouring on G with respect to this vertex ordering. This colouring uses at most $k + 1$ colours, because when one colours v_i there are at most k colours which cannot be used. \square


Corollary 7.20. $\chi(G) \leq \Delta(G) + 1$.

Remark 7.21. This bound is tight if $G = K_n$ or if G is an odd cycle.

Theorem 7.22 (Brooks 1941). *If G is a connected graph other than a clique or an odd cycle, then $\chi(G) \leq \Delta(G)$.*

Proof. Suppose G is connected but is not a clique or an odd cycle, and let $k = \Delta(G)$. We may assume $k \geq 3$, since G is a clique when $k = 1$ and G is an odd cycle or is bipartite when $k = 2$.


If G is not k -regular, choose v_n so that $d(v_n) < k$. Since G is connected, we can grow a spanning tree of G from v_n , assigning indices in decreasing order as we reach vertices. Each vertex other than v_n in the resulting ordering v_1, \dots, v_n has a higher-indexed neighbour along the path to v_n in the tree. Hence, each vertex has at most $k - 1$ lower indexed neighbours, and the greedy colouring uses at most k colours.



In the remaining case, G is k -regular. If G has a cut-vertex x , let G' be a component of $G \setminus x$ together with its edges to x . The degree of x in G' is less than k , and we obtain a proper k -colouring of G' as above. By permuting the names of colours in each such subgraph we can make the colourings agree on x to complete a proper k -colouring of G .

If G has a vertex with two non-adjacent neighbours whose deletion leaves a connected subgraph, then we number these v_1, v_2 and let their common neighbour be v_n . Because $G \setminus \{v_1, v_2\}$ is connected, it has a spanning tree, and we number this tree using v_3, \dots, v_n such that labels increase along paths to the root v_n . As before, each vertex before v_n has at most $k - 1$ lower indexed neighbours. The greedy colouring also uses at most $k - 1$ colours on neighbours of v_n , since v_1 and v_2 receive the same colour.

Hence, it suffices to show that every 2-connected k -regular graph with $k \geq 3$ has such vertices. Choose a vertex x . If $\kappa(G \setminus x) \geq 2$, let v_1 be x and let v_2 be a vertex with distance two from x , which exists because G is connected and not a clique. If $\kappa(G \setminus x) = 1$, then x has a neighbour in every leaf block of $G \setminus x$. To see this, note that otherwise, the single vertex shared by a leaf block and its neighbour (recall Proposition 3.13) would be a cut-vertex, which we are assuming G does not have. Neighbours v_1, v_2 of x in two such blocks are non-adjacent. Furthermore, $G \setminus \{x, v_1, v_2\}$ is connected, since blocks have no cut-vertices (deleting the vertices v_1 and v_2 within their blocks will not disconnect $G \setminus x$). Now $k \geq 3$ tells us that $\deg(x) \geq 3$, which implies that $G \setminus \{v_1, v_2\}$ is also connected.



□

7.5 Colouring planar graphs


Claim 7.23. A (simple) planar graph G contains a vertex v of degree at most 5.

Proof. Recall that in a planar graph $|E(G)| \leq 3|V| - 6$. Thus, we have that $\sum_{v \in V(G)} d(v) \leq 6|V| - 12 < 6|V|$ and so the claim follows. □

Corollary 7.24. A planar graph G is 5-degenerate and thus 6-colourable.

Theorem 7.25 (5 colour theorem; Heawood 1890). Every planar graph G is 5-colourable.

Proof. By induction of $|V(G)|$. For $|V(G)| \leq 5$ the statement is obvious. Assume $|V(G)| > 5$. Let v be a vertex of degree at most 5 in G . By induction, $G \setminus v$ is 5-colourable. If $d(v) < 5$, then a colouring of $f : V(G) \setminus \{v\} \rightarrow \{1, \dots, 5\}$ can be extended to $V(G)$ by assigning $f(v) \in \{1, \dots, 5\} \setminus \{f(u) : uv \in E(G)\}$. Hence, we may assume that $d(v) = 5$. Fix a planar embedding of G in which the neighbours of v are coloured by f with the colours $1, \dots, 5$ in clockwise order (if f uses less than 5 colours on $N(v)$ then it can be extended to $V(G)$ as before). Let the corresponding vertices be v_1, \dots, v_5 , i.e., $f(v_i) = i$ for $i = 1, \dots, 5$.



For $1 \leq i \neq j \leq 5$ let G_{ij} be the subgraph of $G \setminus v$ induced by the colours i and j . Switching the two colours in any connected component of G_{ij} again gives a proper 5-colouring of $G \setminus v$. If the component of G_{ij} containing v_i does not contain v_j then we switch the colours in the component of G_{ij} which contains v_i , in order to remove the colour i from $N(v)$; we can then colour v with the colour i . We can therefore assume that for every pair $1 \leq i \neq j \leq 5$ the component of G_{ij} containing v_i also contains v_j . Let P_{ij} be a path in G_{ij} from v_i to v_j . Obviously, the vertices of P_{ij} are coloured alternatively by colours i and j . Consider paths P_{13} and P_{24} . By the Jordan curve theorem (look at the picture above!), they should intersect. Since the drawing is planar, they intersect in a vertex. But all the vertices of P_{13} are coloured 1 and 3 and all the vertices of P_{24} are coloured 2 and 4, contradiction! \square

Theorem 7.26 (Appel-Haken 1977; conjectured by Guthrie in 1852). *Every planar graph is 4-colourable. (the countries of every plane map can be 4-coloured so that neighbouring countries get distinct colours).*

Remark 7.27. The only known proofs heavily use computers.

8 More colouring results

Theorem 8.1 (Gallai, Roy). *If D is an orientation of G with longest path length $\ell(D)$, then $\chi(G) \leq 1 + \ell(D)$. Furthermore, equality holds for some orientation of G .*

Proof. Suppose D is an orientation of G . Let D' be a maximal acyclic subdigraph of D (this means that adding any additional edge of D to D' would create a directed cycle). An obvious way to obtain such a subgraph is to start from the empty digraph on the same set of vertices, and arbitrarily add

edges of D one-by-one as long as these new edge do not create a directed cycle. When we cannot add any further edges, we will have a suitable subdigraph D' .

Colour $V(G)$ by letting $f(v)$ be one more than the length of the longest path in D' that ends at v . Because D' is acyclic, for every edge $u \rightarrow v$, the longest path in D' ending in u cannot contain v . Therefore any path ending at u can be extended to a longer path ending at v , showing that $f(v) > f(u)$. This shows that f strictly increases along each path in D' .


The colouring f uses colours 1 through $1 + \ell(D')$ on $V(D') = V(G)$. We now prove that f is a proper colouring of G . If $(u, v) \in E(D')$, the above discussion immediately shows that $f(u) \neq f(v)$. Otherwise, if $(u, v) \in E(D)$ which is not in D' , there is a path in D' between its endpoints (since the addition of (u, v) to D' creates a cycle). This again implies $f(u) \neq f(v)$, since f increases along paths of D' .

To prove the second statement, we construct an orientation D^* such that $\ell(D^*) \leq \chi(G) - 1$. Let f be an optimal colouring of G . For each edge (u, v) in G , orient the edge as $u \rightarrow v$ in D^* if and only if $f(u) < f(v)$. Since f is a proper colouring, this defines an orientation. Since the labels used by f increase along each path in D^* , and there are only $\chi(G)$ labels in f , we have $\ell(D^*) \leq \chi(G) - 1$. \square

8.1 Large girth and large chromatic number

The bound $\chi(G) \geq \omega(G)$ can be tight, but (surprisingly) it can also be arbitrarily bad. There are graphs having arbitrarily large chromatic number, even though they do not contain K_3 . Many constructions of such graphs are known, though none are trivial. We give one here.


Example 8.2 (Mycielski's construction). Mycielski (1955) found a construction that builds from any k -chromatic triangle-free graph G a $(k+1)$ -chromatic triangle-free supergraph G' . Given G with vertex set $V = \{v_1, \dots, v_n\}$, add vertices $U = \{u_1, \dots, u_n\}$ and one more vertex w . Beginning with $G'[V] = G$, add edges to make u_i adjacent to all of $N_G(v_i)$, and then make $N(w) = U$. Note that U is an independent set in G' . From the 2-chromatic graph K_2 , one iteration of Mycielski's construction yields the 3-chromatic C_5 , and a second iteration yields the 4-chromatic *Grötzsch graph* drawn below. These graphs are the triangle-free k -chromatic graphs with fewest vertices for $k = 2, 3, 4$.



Theorem 8.3. *Mycielski's construction produces a $(k+1)$ -chromatic triangle-free graph from a k -chromatic triangle-free graph.*

Proof. Suppose G is triangle-free, $\chi(G) = k$, $V(G) = \{v_1, \dots, v_n\}$ and $V(G') = \{v_i\}_{i \in [n]} \cup \{u_i\}_{i \in [n]} \cup \{w\}$, as described above. Since $\{u_i\}_{i \in [n]}$ is independent in G' and there are no edges from w to $V(G)$, there is no triangle containing w . Moreover, any triangle containing u_i has two other vertices in $V(G)$ and these vertices are also neighbours of v_i , which completes a triangle in G . Therefore, if G is triangle-free, then G' is also triangle-free.

A proper k -colouring f of G extends to a proper $(k + 1)$ -colouring of G' by setting $f(u_i) = f(v_i)$ and $f(w) = k + 1$. Now suppose, for contradiction, that G' has a proper k -colouring g . We may assume that $g(w) = k$, which restricts g to $[k - 1]$ on $\{u_i\}_{i \in [n]}$. Let $A = \{v_i : g(v_i) = k\}$. We will change the colouring on A to obtain a proper $(k - 1)$ colouring of G , which is a contradiction of the fact $\chi(G) = k$. For each $v_i \in A$, we change the colour of v_i to $g(u_i)$. Because g properly colours G' , A is an independent set in G , so we need only to check edges of the form (v_i, v') with $v' \in V(G) \setminus A$. If $(v_i, v') \in E(G)$, then we have constructed G' so that $(u_i, v') \in E(G')$, which implies $g(u_i) \neq g(v')$. Hence our alteration does not violate the properness of the colouring on edges within G . We ignore possible conflicts between v_i and u_j , because we now delete $\{u_i\}_{i \in [n]} \cup \{w\}$ and have a proper $(k - 1)$ -colouring of G .



□

Definition 8.4. The *girth* of a graph is the length of its shortest cycle.

Theorem 8.5 (Erdős, 1959). *Given $k \geq 3$ and $g \geq 3$, there exists a graph with girth at least g and chromatic number at least k .*

This result is especially surprising because if a graph has no short cycles then it “locally” looks like a tree, and all trees have chromatic number at most 2. This shows that the chromatic number really depends on the global structure of a graph, and cannot be estimated from local considerations.

The proof of Theorem 8.5 uses the so-called *probabilistic method*.

8.1.1 Digression: the probabilistic method

The general idea of the probabilistic method is to prove the existence of certain structures by showing that they exist with positive probability in some probability space. Let us introduce some probabilistic notation and results.

Definition 8.6. The *expectation* $\mathbb{E}X$ of a random variable X is the “average” value it takes: if X takes only countably many possible values, then

$$\mathbb{E}[X] = \sum_x x \Pr[X = x].$$

Remark 8.7. Expectation has the following properties:

- Expectation is linear: $\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$, regardless of whether X and Y are independent.

- If X is a nonnegative integer valued random variable then

$$\mathbb{E}[X] = \sum_{x=0}^{\infty} x \Pr[X = x] \geq \sum_{x=1}^{\infty} \Pr[X = x] = \Pr[X \geq 1].$$

In particular, if $\mathbb{E}[X] < 1$ then $\Pr[X = 0] = 1 - \Pr[X \geq 1] \geq 1 - \mathbb{E}[X] > 0$.

- More generally, for any positive a ,

$$\mathbb{E}[X] = \sum_{x=0}^{\infty} x \Pr[X = x] \geq \sum_{x \geq a} a \Pr[X = x] = a \Pr[X \geq a].$$

Dividing by a , we get *Markov's inequality*: $\Pr[X \geq a] \leq \mathbb{E}[X]/a$.

Before we proceed to the proof of Theorem 8.5, let us illustrate the method on a slightly easier problem.

A *tournament* is an orientation of the complete graph K_n . The name comes from observing that such a digraph corresponds to an n -player round-robin tournament, where every player plays against everyone else exactly once, and there are no draws. An edge from a to b means that player a beats player b .

Theorem 8.8. *There is a tournament on n vertices where any $\frac{\log_2 n}{2}$ vertices are beaten by some other vertex.*

Proof. Set $k = \frac{\log_2 n}{2}$ and consider the random tournament where each edge of K_n is given a random orientation (each direction with probability $1/2$) independently from the others. For a set K of k vertices, let X_K be the indicator random variable of the event that K is not beaten by any other vertex. Then $\Pr[X_K = 1] = (1 - \frac{1}{2^k})^{n-k}$. The variable $X = \sum_{|K|=k} X_K$ counts the number of such “bad” k -sets. Next we bound the expectation of X using the easy estimates $\binom{n}{k} \leq n^k$ and $1 - x \leq e^{-x}$:

$$\mathbb{E}[X] = \sum_{|K|=k} X_K = \binom{n}{k} (1 - \frac{1}{2^k})^{n-k} \leq n^k \cdot e^{-\frac{n-k}{2^k}} = \exp(k \log n - \frac{n-k}{\sqrt{n}}).$$

Here the exponent tends to $-\infty$ and hence $\mathbb{E}[X] \rightarrow 0$ as $n \rightarrow \infty$. So for n large enough, $\mathbb{E}[X] < 1$. Applying the observation above, we get $\Pr[X = 0] > 0$, i.e., the random tournament satisfies the desired property with positive probability. This proves the existence of such a tournament. \square

We will look at tournaments in more depth later in the course (Section 10.2).

8.1.2 Proof of Theorem 8.5

Proof. Given a large value of n (large enough to satisfy inequalities we will use later), we randomly generate a graph with vertex set $[n]$. Simply let each pair $(x, y) \in [n]^2$ be an edge with probability p , independently. A graph with no large independent set has large chromatic number since $\chi(G) \geq$

$n(G)/\alpha(G)$. We therefore choose p large enough to make the existence of large independent sets unlikely. We also want to choose p small enough that the expected number of short cycles (length less than g) is small. When we have such a graph satisfying both conditions, we can delete a vertex from each short cycle to obtain the desired graph.

To make it unlikely that the graph generated will have more than $n/2$ short cycles, we let $p = n^{t-1}$, where $t = 1/(2g)$. Let X denote the number of cycles of length less than g in the resulting random graph. For each possible cycle C , define a random variable X_C that takes the value 1 if C is present in our random graph, and 0 otherwise. We then have $X = \sum_C X_C$, where the sum is over all possible cycles C of length less than g .

Since there are at most n^i potential cycles of length i , linearity of expectation yields $\mathbb{E}[X] \leq \sum_{i=3}^{g-1} n^i p^i = \sum_{i=3}^{g-1} n^{ti} < 2n^{tg} = 2\sqrt{n}$. This implies that $\mathbb{E}[X]/n \rightarrow 0$ as $n \rightarrow \infty$. Since $\mathbb{E}[X]/n \rightarrow 0$, we also have $\Pr[X \geq n/2] \rightarrow 0$ as $n \rightarrow \infty$ by Markov's inequality. In particular, we can take n large enough so that $\Pr[X \geq n/2] < 1/2$.

Since we will retain at least $n/2$ vertices, it suffices to show that there will be such graphs with $\alpha(G) \leq n/(2k)$; $\alpha(G)$ cannot grow when we delete vertices, and hence at least k independent sets will be needed to cover the remaining vertices. With $r = \lfloor n/(2k) \rfloor$, we have


$$\Pr[\alpha(G) > r] \leq \binom{n}{r+1} (1-p)^{\binom{r+1}{2}} < \left(ne^{-pr/2} \right)^{r+1}.$$

As $n \rightarrow \infty$ we have $ne^{-pr/2} = ne^{-n^t/(4k)} \rightarrow 0$ (and $r \rightarrow \infty$), so $\Pr[\alpha(G) > r] \rightarrow 0$. If we also make n large enough that $\Pr[X \geq n/2] < 1/2$ and $\Pr[\alpha(G) > r] < 1/2$, there will exist an n -vertex graph G such that $\alpha(G) \leq n/(2k)$ and G has fewer than $n/2$ cycles of length less than g . We delete a vertex from each short cycle and retain a graph with girth at least g and chromatic number at least k . \square

8.2 Chromatic number and clique minors

Definition 8.9. Let $e = (x, y)$ be an edge of a graph $G = (V, E)$. By G/e we denote the graph obtained from G by *contracting* the edge e into a new vertex v_e , which becomes adjacent to all the former neighbours of x and of y . Formally, G/e is a graph (V', E') with vertex set $V' := (V \setminus \{x, y\}) \cup \{v_e\}$ (where v_e is the 'new' vertex, i.e. $v_e \notin V \cup E$) and edge set

$$E' := \{(v, w) \in E \mid \{v, w\} \cap \{x, y\} = \emptyset\} \cup \{(v_e, w) \mid (x, w) \in E \setminus \{e\} \text{ or } (y, w) \in E \setminus \{e\}\}$$



H is a *minor* of G if it can be obtained from G by deleting vertices and edges, and contracting edges.

Question 8.10. We know that if G has high chromatic number then it might not have a large clique. However, must G necessarily have a subgraph which is similar to a clique?

Conjecture 8.11 (Hadwiger). *If $\chi(G) \geq t$ then G has a t -clique minor.*

Note that large chromatic number implies large maximum degree.

Theorem 8.12 (Mader). *If the average degree of G is at least 2^{t-2} then G has a K_t minor.*

Proof. We proceed by induction on $t + |V(G)|$. If $t = 2$, then the average degree of G is at least 1. This means G has an edge; this is a K_2 minor.

Suppose now that $\bar{d}(G) \geq 2^{t-1}$, i.e., $e(G) \geq 2^{t-2}n$. Let v be any vertex and let $u \in N(v)$. If u and v have at most $2^{t-2} - 1$ common neighbours, contract the edge $e = (u, v)$. Note that the number of vertices of G/e is one less than the number of vertices of G and that the number of edges in G/e decreased by at most $(2^{t-2} - 1) + 1 = 2^{t-2}$ (note that when we contract e we lose the edge e , and one edge for every common neighbour of u and v). Thus,

$$\bar{d}(G/e) = \frac{2e(G/e)}{n-1} \geq \frac{2(e(G) - 2^{t-2})}{n-1} \geq \frac{2^{t-1}n - 2^{t-1}}{n-1} = 2^{t-1}.$$

Therefore, by induction G/e has a K_{t+1} minor and so G also has one. We may assume now that for every $u \in N(v)$ the vertices u and v have at least 2^{t-2} common neighbours. In that case it is clear that $G' = G[N(v)]$ has $\bar{d}(G') \geq 2^{t-2}$. Thus, by induction G' has a K_t minor which, together with v , gives a K_{t+1} minor in G . \square


Remark 8.13. It is known that $\bar{d}(G) \geq ct\sqrt{\log t}$ already implies the existence of a K_t minor in G , for some constant $c > 0$.

8.3 Edge-colourings

Definition 8.14. A k -edge-colouring of G is a labeling $f : E(G) \rightarrow [k]$; the labels are ‘‘colours’’. A *proper k -edge-colouring* is a k -edge-colouring such that edges sharing a vertex receive different colours; equivalently, each colour class is a matching. A graph G is *k -edge-colourable* if it has a proper k -edge-colouring. The *edge-chromatic number* or *chromatic index* $\chi'(G)$ is the minimum k such that G is k -edge colourable.

Remark 8.15.

- (i) An edge-colouring of a graph G is the same as a vertex-colouring of its line graph $L(G)$.
- (ii) A graph G with maximum degree Δ has $\chi'(G) \geq \Delta$ since the edges incident to a vertex of degree Δ must have different colours.
- (iii) If G has maximum degree Δ then $L(G)$ has maximum degree at most $2(\Delta - 1)$ (see the picture below for the basic idea of why this is true).



So it follows from (i) that $\chi'(G) \leq 2\Delta - 1$.

Theorem 8.16 (Kőnig 1916). *If G is a bipartite multigraph, then $\chi'(G) = \Delta(G)$.*

Proof. An edge colouring is actually a partition of the edge set of a graph into matchings. We therefore make use of the theory of matchings from Section 5.

Corollary 5.9 states that every regular bipartite graph H has a perfect matching. This also holds for multigraphs, with the same proof. By induction on $\Delta(H)$, this yields a proper $\Delta(H)$ -edge-colouring. It therefore suffices to show that every bipartite multigraph G with maximum degree k has a k -regular bipartite supergraph $H \supset G$.

We construct such a supergraph. If G does not have the same number of vertices in each partite set, add vertices to the smaller set to equalize the sizes. If the resulting G' is not regular, then each partite set has a vertex with degree less than $\Delta(G') = \Delta(G)$. Add an edge consisting of this pair. Continue adding such edges until the graph becomes regular. \square

Theorem 8.17 (Vizing). *Let G be a simple graph with maximum degree Δ . Then $\Delta(G) \leq \chi'(G) \leq \Delta(G) + 1$.*

Proof. The inequality $\Delta(G) \leq \chi'(G)$ being trivial, we show $\chi'(G) \leq \Delta(G) + 1$. We prove this by induction on the number of vertices of G , by applying the following claim with $k = \Delta + 1$ and any vertex v .

Claim. Let v be a vertex such that v and all its neighbours have degree at most k , while at most one neighbour has degree precisely k . Then if $G \setminus v$ is k -edge-colourable, G is also k -edge-colourable.

We prove the claim by induction on k . We can assume that each neighbour u of v has degree $k - 1$, except for one of degree k , since otherwise we can add a new vertex w and an edge (u, w) without violating the conditions in the claim. We can do this until all neighbours of v have degree $k - 1$, except for one having degree k . Similarly we can assume v has degree k , since otherwise we can repeatedly add new vertices adjacent only to v .

Consider any k -edge-colouring of $G \setminus v$. For $i = 1, \dots, k$, let X_i be the set of neighbours of v that are missed by colour i . All but one of the neighbours of v have degree $k - 2$ in $G \setminus v$ (these neighbours are missing two colors on their edges), and one neighbor of v has degree $k - 1$ in $G \setminus v$ so (this neighbour misses one color on its edges). So all but one neighbour of v are in precisely two of the X_i , and one neighbour is in precisely one X_i . Hence

$$\sum_{i=1}^k |X_i| = 2(k - 1) + 1 < 2k. \quad (6)$$

We can assume that we have chosen the colouring such that $\sum_{i=1}^k |X_i|^2$ is minimized. Then for any $i, j = 1, \dots, k$:

$$||X_i| - |X_j|| \leq 2. \quad (7)$$

For if, say, $|X_1| > |X_2| + 2$, consider the subgraph H made by all edges of colours 1 and 2. Each component of H is a path or an even cycle. At least one component of H contains more vertices in X_1 than in X_2 . This component is a path P starting in X_1 and not ending in X_2 . Exchanging colours 1 and 2 on P reduces X_1 by two and increases X_2 by two. This reduces $|X_1|^2 + |X_2|^2$, contradicting our minimality assumption. This proves (7).

It follows that there exists an i with $|X_i| = 1$, since otherwise by (6) and (7) each $|X_i|$ is 0 or 2, while their sum is odd, a contradiction.

So we can assume $|X_k| = 1$, say $X_k := \{u\}$. Let G' be the graph obtained from G by deleting the edge (v, u) and deleting all edges of colour k . So $G' \setminus v$ is $(k-1)$ -edge-coloured. Moreover, in G' , vertex v and all its neighbours have degree at most $k-1$, and at most one neighbour has degree $k-1$. So by the induction hypothesis, G' is $(k-1)$ -edge-colourable. Restoring colour k , and giving edge (v, u) colour k , gives a k -edge-colouring of G . (Giving (v, u) colour k does not violate the properness of the colouring: v has no other adjacent edges coloured k by construction, and since $u \in X_k$, u also has no other adjacent edges coloured k). \square

Remark 8.18. In general it is a hard problem to decide whether $\chi'(G) = \Delta(G)$.

8.4 List colouring

List colouring is a more general version of the usual colouring problem. We still want to pick a single colour for each vertex, but the set of colours available at each vertex may be restricted.

Definition 8.19. For each vertex v in a graph G , let $L(v)$ denote a list of colours available for v . A *list colouring* or *choice function* from a given collection of lists is a proper colouring f such that $f(v)$ is chosen from $L(v)$. A graph G is *k -choosable* or *k -list-colourable* if it has a proper list colouring from every assignment of k -element lists to the vertices. The *list chromatic number* or *choosability* $\chi_l(G)$ is the minimum k such that G is k -choosable.

Since the lists can be chosen to be identical, $\chi_l(G) \geq \chi(G)$. If the lists have size at least $1 + \Delta(G)$, then colouring the vertices in succession leaves an available colour at each vertex. This argument is analogous to the greedy colouring algorithm and proves that $\chi_l(G) \leq \Delta(G) + 1$. It is not possible, however, to place an upper bound on $\chi_l(G)$ in terms of $\chi(G)$. For example, there are bipartite graphs with arbitrarily large list-chromatic number.

Theorem 8.20 (Erdős, Rubin, Taylor 1979). *If $m = \binom{2k-1}{k}$, then $K_{m,m}$ is not k -choosable.*

Proof. Let X, Y be the bipartition of $G = K_{m,m}$. For each of X and Y , let the lists of the vertices be all the k -subsets of $[2k-1]$. Suppose G has a choice function f . If f uses fewer than k distinct choices in X , then there is a k -set $K \subset [2k-1]$ not used, which means that no colour was chosen for the vertex of X having K as its list. If f uses at least k colours on vertices of X , then there is a k -set $K \subset [2k-1]$ of colours used in X , and no colour can be properly chosen for the vertex Y with list K . \square

Definition 8.21. Let $L(e)$ denote the list of colours available for e . A *list edge-colouring* is a proper edge-colouring f with $f(e)$ chosen from $L(e)$ for each e . The *list chromatic index* or *edge-choosability* $\chi'_l(G)$ is the minimum k such that G has a proper list edge-colouring for each assignment of lists of size k to the edges. Equivalently, $\chi'_l(G) = \chi_l(L(G))$, where $L(G)$ is the line graph of G .

Conjecture 8.22. $\chi'_l(G) = \chi'(G)$, and in particular $\chi'_l(G) \leq \Delta(G) + 1$.

Theorem 8.23 (Galvin,1995). $\chi'_l(K_{n,n}) = n$.

Note that the graph $L(K_{n,n})$ is essentially an $n \times n$ grid. That is, the vertices of $L(K_{n,n})$ can be arranged in an $n \times n$ grid, where vertices are adjacent if and only if they share a row or share a column. (This graph is known as the *Cartesian product* $K_n \times K_n$). Theorem 8.23 therefore has the following appealing equivalent formulation. If we have an $n \times n$ grid where each cell is assigned a list of n colours, then we can choose a colour for each cell such that each row and column have distinct colours.

Definition 8.24. A *kernel* of a digraph is an independent set S having an edge to every vertex outside S . A digraph is *kernel-perfect* if every induced sub-digraph has a kernel. Given a function $f : V(G) \rightarrow \mathbb{N}$, the graph G is *f-choosable* if a proper list colouring can be chosen whenever the lists satisfy $|L(x)| \geq f(x)$ for each x .

Lemma 8.25. If D is a kernel-perfect orientation of G and $f(x) = d_D^-(x)$ for all $x \in V(G)$, then G is $(1 + f)$ -choosable.

Proof. By induction on $|G|$; trivial for $|G| = 1$. Suppose $|G| > 1$, and consider an assignment of lists, with the list $L(x)$ for each x having size $1 + f(x)$. Choose a colour c appearing in some list. Let $U = \{v : c \in L(v)\}$. Let S be the kernel of the induced subdigraph $D[U]$. Assign the colour c to all of S , which is permissible since S is independent. Delete c from $L(v)$ for each $v \in U \setminus S$. Although the list of every vertex in $U \setminus S$ lost color c , the indegree of these vertices in $D \setminus S$ also decreased by at least one. Hence we still have at least $d_{D \setminus S}^-(u) + 1$ colors left in the list for every vertex in $U \setminus S$. Therefore for each $x \in V(D) \setminus S$ its list has size at least $1 + f'(x)$, where $f'(x) = d_{D \setminus S}^-(x)$. By the induction hypothesis, $D \setminus S$ is $(1 + f')$ -choosable, so we can complete a list colouring for the original graph from the remaining colours at $V(G) \setminus S$. \square

Proof of Theorem 8.23. First note that $\chi'_l(K_{n,n}) \geq \chi'(K_{n,n}) \geq n$.

It suffices by the lemma to prove that $D = K_n \times K_n$ has a kernel-perfect orientation with each vertex having indegree and outdegree $n - 1$. Label the vertices with labels $1, 2, \dots, n$ so that vertex (r, s) has label $r + s - 1 \pmod n$. Note that the labels of vertices in every row and column are just permutations of $\{1, \dots, n\}$. Define an orientation of $K_n \times K_n$ by directing edges from vertex (r, s) to the vertices in row r with higher labels and to the the vertices in column s with lower labels. Since label j is higher than $j - 1$ other labels, a vertex with label j has $j - 1$ predecessors in its row and $n - j$ predecessors in its column. Hence each vertex has indegree $n - 1$ and outdegree $n - 1$.

1	2	3	...	n
2	3	...	n	1
3	\vdots
\vdots	n	$n - 2$
n	1	...	$n - 2$	$n - 1$

We prove that for every set of vertices $U \subseteq D$, the subdigraph $D[U]$ has a kernel, by induction on $|U|$. The case $|U| = 1$ is trivial, because U itself is trivially an independent set. Otherwise, consider the set S_0 consisting of the vertices which have minimum label in their row (among the vertices in U). This means that each $u \in S_0$ has no incoming arcs from vertices of U , in its row. If no two vertices of S_0 fall in the same column then S_0 is an independent set, therefore a kernel. Otherwise, there is a column C with multiple vertices of S_0 . Let u be the vertex of S_0 with the minimum label in C . Then $D[U \setminus u]$ has a kernel S by induction. S must contain a vertex from C , because the non- u vertices in $C \cap S_0$ must have incoming arcs from S , and these incoming arcs cannot come from their rows. By the choice of u , any such vertex in C has an edge to u , so we conclude that S has an edge to every vertex in U , and is a kernel, as required. \square

9 The Matrix Tree Theorem

In Section 2.3 we gave two proofs of the following classical result.

Theorem 9.1 (Cayley's formula). *There are n^{n-2} labeled trees on n vertices.*

Here we will show yet another proof that we deduce from a very general theorem about counting spanning trees in graphs. Indeed, we can think of the number of labeled trees as the number of spanning trees in the complete graph K_n .

Now consider an arbitrary connected simple graph G on vertex set $[n]$, and denote the number of spanning trees by $t(G)$. The following celebrated result is Kirchhoff's *matrix tree theorem*. To formulate it, consider the incidence matrix B of G (as in Definition 1.13), and replace one of the two 1's by -1 in an arbitrary manner to obtain the matrix C (we say C is the incidence matrix of an orientation of G). $M = CC^T$ is then a symmetric $n \times n$ matrix, which is

$$\begin{pmatrix} d(1) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & d(n) \end{pmatrix} - A_G.$$

Theorem 9.2 (Matrix tree theorem). *We have $t(G) = \det M_{ii}$ for all $i = 1, \dots, n$, where M_{ii} results from M by deleting the i -th row and the i -th column.*

The key to the proof is the following theorem of Binet and Cauchy:

Theorem 9.3. *If P is an $r \times s$ matrix and Q is an $s \times r$ matrix with $r \leq s$, then*

$$\det(PQ) = \sum_Z (\det P_Z)(\det Q_Z),$$

where P_Z is the $r \times r$ submatrix of P with column set Z , and Q_Z is the $r \times r$ submatrix of Q with the corresponding rows Z , and the sum is over all r -sets $Z \subseteq [s]$.

We will prove Theorem 9.3 in the next section (Section 9.1), but for now we use it to give the proof of Theorem 9.2.

Proof of Theorem 9.2. Note that C has at least $n - 1$ columns, because the connected graph G has at least $n - 1$ edges. This means that we can apply Theorem 9.3 to M_{ii} and get

$$\det M_{ii} = \sum_N \det N \cdot \det N^T = \sum_N (\det N)^2,$$

where N runs through all $(n - 1) \times (n - 1)$ submatrices of $C \setminus \{\text{row } i\}$. The $n - 1$ columns of N correspond to a subgraph of G with $n - 1$ edges on n vertices, and it remains to show that

$$\det N = \begin{cases} \pm 1 & \text{if these edges span a tree} \\ 0 & \text{otherwise.} \end{cases}$$

Suppose the $n - 1$ edges do not span a tree. Then there exists a component that does not contain the vertex i . Since the corresponding rows of this component add to 0, we infer that they are linearly dependent, and hence $\det N = 0$.

Assume now that the columns of N span a tree. Then there is a vertex $j_1 \neq i$ of degree 1; let e_1 be the incident edge. Deleting j_1, e_1 we obtain a tree with $n - 2$ edges. Again there is a vertex $j_2 \neq i$ of degree 1 with incident edge e_2 . Continue in this way until j_1, \dots, j_{n-1} and e_1, \dots, e_{n-1} with $j_i \in e_i$ are determined. Now permute the rows and columns to bring j_k into the k 'th row and e_k into the k 'th column. Since by construction $j_k \notin e_l$ for $k < l$, we see that the new matrix N' is lower triangular with all elements on the main diagonal equal to ± 1 . Thus $\det N = \pm \det N' = \pm 1$, and we are done. \square

Third proof of Cayley's formula. In the special case $G = K_n$, we have

$$M = \begin{pmatrix} n-1 & -1 & \cdots & -1 \\ -1 & n-1 & & -1 \\ \vdots & & \ddots & \vdots \\ -1 & -1 & \cdots & n-1 \end{pmatrix}.$$

(The $(n - 1)$ s on the diagonal correspond to the fact that the degree of each vertex is $n - 1$, and the -1 s in the other positions i, j correspond to the fact that there is an edge between each i and j).

The $(n - 1) \times (n - 1)$ matrices M_{ii} have the same form as M . To compute $\det M_{ii}$, add all the columns to the first column to get

$$\det M_{ii} = \det \begin{pmatrix} 1 & -1 & \cdots & -1 \\ 1 & n-1 & & -1 \\ \vdots & & \ddots & \vdots \\ 1 & -1 & \cdots & n-1 \end{pmatrix},$$

then add the first column to all the rest to get

$$\det M_{ii} = \det \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 1 & n & & 0 \\ \vdots & & \ddots & \vdots \\ 1 & 0 & \cdots & n \end{pmatrix}.$$

This is a lower-triangular matrix, so we can see $\det M_{ii} = n^{n-2}$. Then Theorem 9.2 finishes the proof. \square

9.1 Lattice paths and determinants


In this section we will introduce some theory connecting graphs and determinants. In particular we will establish the Cauchy-Binet Theorem (Theorem 9.3).

The starting point of our discussion is the usual permutation representation of the determinant of a matrix. Let $M = (m_{ij})$ be a real $n \times n$ matrix. Then

$$\det M = \sum_{\sigma} \text{sign}(\sigma) m_{1\sigma(1)} m_{2\sigma(2)} \cdots m_{n\sigma(n)}, \quad (8)$$

where σ runs through all permutations of $[n]$, and the sign of σ is 1 or -1 , depending on whether σ is the product of an even or an odd number of transpositions.

Now we pass to graphs, more precisely to *weighted directed bipartite* graphs. Let the vertices a_1, \dots, a_n stand for the rows of M , and b_1, \dots, b_n for the columns. For each pair i, j , draw an arrow from a_i to b_j and give it the weight m_{ij} , as in the figure.



In terms of this graph, the formula (8) has the following interpretation:

- The left-hand side is the determinant of the *path matrix* M , whose (i, j) entry is the *weight* of the (unique) directed path from a_i to b_j .
- The right-hand side is the weighted (signed) sum over all *vertex-disjoint path systems* from $A = \{a_1, \dots, a_n\}$ to $B = \{b_1, \dots, b_n\}$. Such a system \mathcal{P}_σ is given by paths

$$a_1 \rightarrow b_{\sigma(1)}, \dots, a_n \rightarrow b_{\sigma(n)},$$

and the *weight* of the path system \mathcal{P}_σ is the product of the weights of the individual paths:

$$w(\mathcal{P}_\sigma) = w(a_1 \rightarrow b_{\sigma(1)}) \cdots w(a_n \rightarrow b_{\sigma(n)}).$$

In this interpretation formula (8) reads

$$\det M = \sum_{\sigma} \text{sign}(\sigma) w(\mathcal{P}_\sigma).$$

A natural generalization of (8) from bipartite to arbitrary graphs was found by Gessel and Viennot. This widely applicable result has a very simple and elegant proof, as we will see below.

Let us first collect the necessary concepts. We are given a finite acyclic directed graph $G = (V, E)$, i.e. one that has no directed cycles in it. In particular, there are only finitely many directed paths between any two vertices a and b in G , where we include all trivial paths $a \rightarrow a$ of length 0. Every edge e carries a weight $w(e)$. If P is a directed path from a to b , written shortly $P : a \rightarrow b$, then we define the weight of P as

$$w(P) = \prod_{e \in P} w(e),$$

which is defined to be 1 if P is a path of length 0.

Now let $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ be two sets of n vertices, where A and B are not necessarily disjoint. To A and B we associate the *path matrix* $M = (m_{ij})$ with

$$m_{ij} = \sum_{P: a_i \rightarrow b_j} w(P).$$

A *path system* \mathcal{P} from A to B consists of a permutation σ together with n paths $P_i : a_i \rightarrow b_{\sigma(i)}$ for $i = 1, \dots, n$. We write $\text{sign}(\mathcal{P}) = \text{sign}(\sigma)$. The *weight* of \mathcal{P} is the product of the path weights

$$w(\mathcal{P}) = \prod_{i=1}^n w(P_i), \tag{9}$$

which is the product of the weights of all the edges of the path system.

Finally, we say that the path system $\mathcal{P} = (P_1, \dots, P_n)$ is *vertex-disjoint* if the paths of \mathcal{P} are pairwise vertex-disjoint.

Lemma 9.4 (Gessel-Viennot). *Let $G = (V, E)$ be a finite weighted acyclic directed graph, $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ two n -sets of vertices, and M the path matrix from A to B . Then*

$$\det M = \sum_{\substack{\mathcal{P} \text{ vertex-disjoint} \\ \text{path system}}} \text{sign}(\mathcal{P})w(\mathcal{P}). \tag{10}$$

Proof. A typical summand of $\det(M)$ is $\text{sign}(\sigma)m_{1\sigma(1)} \cdots m_{n\sigma(n)}$, which can be written as

$$\text{sign}(\sigma) \left(\sum_{P_1: a_1 \rightarrow b_{\sigma(1)}} w(P_1) \right) \cdots \left(\sum_{P_n: a_n \rightarrow b_{\sigma(n)}} w(P_n) \right)$$

Summing over σ we immediately find from (9) that

$$\det M = \sum_{\mathcal{P}} \text{sign}(\mathcal{P})w(\mathcal{P}),$$

where \mathcal{P} runs through all path systems from A to B (vertex-disjoint or not). Hence to arrive at (10), all we have to show is


$$\sum_{\mathcal{P} \in \mathcal{N}} \text{sign}(\mathcal{P})w(\mathcal{P}) = 0, \tag{11}$$

where N is the set of all path systems that are *not* vertex-disjoint. To show this, we exhibit an involution $\pi : N \rightarrow N$ (without fixed points) such that for \mathcal{P} and $\pi\mathcal{P}$

$$w(\pi\mathcal{P}) = w(\mathcal{P}) \quad \text{and} \quad \text{sign}(\pi\mathcal{P}) = -\text{sign}(\mathcal{P}).$$

This will clearly imply (11) and thus the formula (10) of the Lemma.

The involution π is defined in a natural way. Let $\mathcal{P} \in N$ with paths $P_i : a_i \rightarrow b_{\sigma(i)}$. By definition, some pair of paths will intersect:



- Let i_0 be the minimal index such that P_{i_0} shares some vertex with another path.
- Let x be the first such common vertex on the path P_{i_0} .
- Let j_0 be the minimal index ($j_0 > i_0$) such that P_{j_0} has the vertex x in common with P_{i_0} .

Now we construct the new system $\pi\mathcal{P} = (P'_1, \dots, P'_n)$ as follows:

- Set $P'_k = P_k$ for all $k \neq i_0, j_0$.
- The new path P'_{i_0} goes from a_{i_0} to x along P_{i_0} , and then continues to $b_{\sigma(j_0)}$ along P_{j_0} . Similarly, P'_{j_0} goes from a_{j_0} to x along P_{j_0} and continues to $b_{\sigma(i_0)}$ along P_{i_0} .


Clearly $\pi(\pi\mathcal{P}) = \mathcal{P}$, since the index i_0 , the vertex x and the index j_0 are the same as before. In other words, applying π twice we switch back to the old paths P_i . Next, since $\pi\mathcal{P}$ and \mathcal{P} use precisely the same edges, we certainly have $w(\pi\mathcal{P}) = w(\mathcal{P})$. And finally, since the new permutation σ' is obtained by multiplying σ with the transposition (i_0, j_0) , we find that $\text{sign}(\pi\mathcal{P}) = -\text{sign}(\mathcal{P})$, and we are done. \square

The Gessel-Viennot Lemma can be used to derive basic properties of determinants, just by looking at appropriate graphs. For us, it will serve as a tool to deduce the Cauchy-Binet formula.

Proof of Theorem 9.3. As before, let A and B be vertex sets corresponding to the rows and columns, respectively, of the matrix P . Similarly let B and C correspond to the matrix Q . Consider now

the concatenated graph as indicated in the figure, and observe that the (i, j) -entry m_{ij} of the path matrix M from A to C is precisely $m_{ij} = \sum_k p_{ik}q_{kj}$, thus $M = PQ$.


Te vertex-disjoint path systems from A to C in the concatenated graph correspond to pairs of systems from A to Z and Z to C , for r -element subsets Z of B . The result follows immediately from Lemma 9.4, by noting that $\text{sign}(\sigma\tau) = \text{sign}(\sigma) \text{sign}(\tau)$.



□


10 More Theorems on Hamiltonicity

Definition 10.1. The (Hamiltonian) closure of a graph G , denoted $C(G)$, is the subgraph of G on $V(G)$ obtained by iteratively adding edges between pairs of nonadjacent vertices whose degree sum is at least n , until no such pair remains. In the example below, $n = 6$.



Theorem 10.2 (Bondy Chvátal 1976). *A simple n -vertex graph is Hamiltonian if and only if its closure is Hamiltonian.*

Sketch of Proof. Suppose $G \cup \{(u, v)\}$ is Hamiltonian and $d(u) + d(v) \geq n$. Then G has a Hamilton path from u to v . Since $|N(u)| + |N(v)| \geq n$, on the path there is a vertex $v_i \sim v$ such that the next vertex $v_{i+1} \sim u$ (as in the proof of Dirac's Theorem). This gives a Hamilton cycle in G .



□

Theorem 10.3 (Chvátal 1972). *Suppose G has vertex degrees $d_1 \leq \dots \leq d_n$. If $i < n/2$ implies that $d_i > i$ or $d_{n-i} \geq n - i$, then G is Hamiltonian.*


Proof. Because adding edges to form the closure reduces no entry in the degree sequence, and a graph is Hamiltonian if and only if its closure is Hamiltonian, it suffices to consider the special case where G is closed. In this case, we prove that the condition implies $G = K_n$. We prove the contrapositive; if $G = C(G) \neq K_n$, then we construct a value of i less than $n/2$ for which Chvátal's condition is violated. Here violation means that at least i vertices have degree at most i and at least $n - i$ vertices have degree less than $n - i$, implying that $d_i \leq i$ and $d_{n-i} < n - i$.

If $G \neq K_n$, then among the pairs of nonadjacent vertices we choose a pair u, v with maximum degree sum. Because G is closed, $u \approx v$ implies $d(u) + d(v) < n$, and we may assume $d(u) \leq d(v)$. We then have $d(u) < n/2$; let $i = d(u)$.

Because we chose a nonadjacent pair with maximum degree sum, every vertex of $V \setminus v$ that is not adjacent to v has degree at most $d(u) = i$; furthermore, there are at least $n - 1 - d(v) \geq d(u) = i$ of these vertices. Similarly, every vertex of $V \setminus u$ that is not adjacent to u has degree at most $d(v) < n - d(u) = n - i$, and there are $n - 1 - d(u) = n - 1 - i$ of these. Since $d(u) \leq d(v)$, we can also add u to the set of vertices with degree at most $d(v)$, so we obtain $n - i$ vertices with degree less than $n - i$. Hence we have proved $d_i \leq i$ and $d_{n-i} < n - i$ for this specially chosen i , which contradicts the hypothesis. \square

Theorem 10.4 (Chvátal-Erdős 1972). *If $\kappa(G) \geq \alpha(G)$, then G has a Hamiltonian cycle (unless $G = K_2$).*

Proof. The theorem is trivially true for complete graphs. Otherwise, the condition requires $\kappa(G) \geq \alpha(G) > 1$. We will prove that if G has no Hamilton cycle, then $k = \kappa(G) < \alpha(G)$. Let C be a longest cycle (which we are assuming is not Hamiltonian) in G . Let H be a component of $G \setminus C$. Since $\delta(G) \geq \kappa(G)$ and every graph with $\delta(G) \geq 2$ has a cycle of length at least $\delta + 1$ (Proposition 1.32), C has at least $k + 1$ vertices. Also C has at least k vertices with edges to H , else the vertices of C with edges to H contradict $\kappa(G) = k$. Let u_1, \dots, u_ℓ be the vertices of C with edges to H ($\ell \geq k$), in clockwise order. For $i = 1, \dots, k$, let a_i be the vertex immediately following u_i on C . If any two of these vertices are adjacent, say $a_i \sim a_j$, then we construct a longer cycle by using (a_i, a_j) , the portions of C from a_i to u_j and a_j to u_i , and a u_i, u_j -path through H (see first illustration). This argument includes the case $a_i = u_{i+1}$ (see second illustration), so we also conclude that no a_i has a neighbour in H . Hence $\{a_1, \dots, a_k\}$ plus a vertex of H forms an independent set of size $k + 1$, proving that $\alpha(G) > k$ as desired.




\square

10.1 Pósa's Lemma

Let P be a path in a graph G , say from u to v . Given a vertex $x \in P$, we write x^- for the vertex preceding x on P , and x^+ for the vertex following x on P (whenever these exist). Similarly, for $X \subseteq V(P)$ we put

$$X^- := \{x^- : x \in X\} \quad \text{and} \quad X^+ := \{x^+ : x \in X\}.$$

If $x \in P \setminus u$ is a neighbour of u in G , then $P \cup \{(u, x)\} \setminus \{(x, x^-)\}$ (which is a path in G with vertex set $V(P)$) is said to have been obtained from P by a *rotation* fixing v (see figure below). A path obtained from P by a (possibly empty) sequence of rotations fixing v is a path *derived* from P . The set of starting vertices of paths derived from P , including u , will be denoted by $S(P)$. As all paths derived from P have the same vertex set as P , we have $S(P) \subseteq V(P)$.



Remark 10.5. If some sequence of rotations can delete the edge (x, x^-) , call this edge a *broken* edge. Note that every interval of the original path not containing broken edges is traversed by all derived paths as a whole piece (however, the direction can change).

Definition 10.6. For a graph G and a subset $S \subseteq V(G)$, let $\partial S = \{v \in G \setminus S : \exists y \in S, v \sim y\}$.

Lemma 10.7. Let G be a graph, let $P = u \dots v$ be a longest path in G , and put $S := S(P)$. Then $\partial S \subseteq S^- \cup S \cup S^+$.

Proof. Let $x \in S$ and $y \in N(x)$ be given; we show that $y \in S^- \cup S \cup S^+$. As $x \in S$ there is a path $Q = x \dots v$ derived from P . Then $y \in V(Q) = V(P)$, because Q (like P) is a longest path and cannot be extended. Let z denote the predecessor of y on Q .

Suppose that $y \notin S^- \cup S \cup S^+$. Then each of the (one or two) edges $e \in P$ at y lies on every path derived from P (and in particular on Q), because in any rotation in which e is first deleted, its two ends (including y) would have become members of S and of $S^- \cup S^+$. Hence $z \in \{y^-, y^+\}$. But $Q \cup \{(x, y)\} \setminus \{(y, z)\}$ is obtained from Q by a rotation, which puts z in S and y in $S^- \cup S^+$. \square

Lemma 10.8. Let G be a graph, let $P = u \dots v$ be a longest path in G , and put $S := S(P)$. If $\deg(u) \geq 2$ then G has a cycle containing $S \cup \partial S$.

Proof. Let y be the last vertex of P in ∂S . Then all the vertices from $S \cup \partial S$ lie on the subpath P_1 of P between u and y , because any vertex in S which is after y in P would differ from v and hence have its successor on P in ∂S . Let $x \in S$ be a neighbour of y in G , and let $Q = x \dots v$ be derived from P . As in the proof of Lemma 10.7, all the edges of P after y are still edges of Q , so the segments of Q and P after y are the same (denote this common path segment by $P_2 = Q_2$). Let Q_1 be the subpath of Q between x and y . We have

$$S \cup \partial S \subseteq V(P_1) = V(P \setminus P_2) = V(P) \setminus V(P_2) = V(Q) \setminus V(Q_2) = V(Q_1).$$

and $Q_1 \cup \{(y, x)\}$ is a cycle in G . For this argument it is important that Q_1 contains a vertex other than x and y . Indeed, since $\deg(u) \geq 2$ there is a nontrivial rotation so $|S| > 1$ and $|Q_1| = |P_1| > 3$. \square

Corollary 10.9. Fix $k \geq 2$ and let G be a graph such that for all $S \subseteq V(G)$ with $|S| \leq k$, we have $|\partial S| \geq 2|S|$. Then G has a cycle of length at least $3k$.

Proof. Let $S = S(P)$ be the set of potential starting points of paths derived from a longest path $P = u \dots v$. Note that every vertex of G has degree at least 2 (for a singleton set S we have $|\partial S| \geq 2|S|$).

If $|S| < k$ then because u has no predecessor, we have $|S^- \cup S^+| < 2|S| \leq |\partial S|$ and therefore $\partial S \not\subseteq S^- \cup S^+$. This is a contradiction of Lemma 10.7, and shows that $|S| \geq k$. By Lemma 10.8, G has a cycle containing $S \cup \partial S$. Since S has a subset of size k , we have $|S \cup \partial S| \geq 3k$. \square

10.2 Tournaments


Definition 10.10. A *tournament* is a directed graph obtained by assigning a direction to every edge of the complete graph. That is, it is an orientation of K_n .

In the context of tournaments, it only makes sense to consider directed paths and cycles, because every edge is present. So, a Hamilton cycle is a spanning directed cycle.

Theorem 10.11. Every tournament has a Hamilton path.

Proof. We proceed by induction on n . The case $n = 2$ is clear (there is only one tournament on two vertices, namely a single directed edge, which is already a Hamilton path).

Suppose the claim holds for all tournaments on n vertices, and let T have $n + 1$ vertices. Let v be any vertex of T . Then $T \setminus v$ is also a tournament and hence has a Hamilton path $u_1 \rightarrow \dots \rightarrow u_n$. Note that if $v \rightarrow u_1$, we are done. So suppose $u_1 \rightarrow v$. Let i be the maximum index such that $u_i \rightarrow v$. If $i = n$ we are done, with the Hamilton path $u_1 \rightarrow \dots \rightarrow u_n \rightarrow v$. If not then $v \rightarrow u_{i+1}$ and we can take the Hamilton path $u_1 \rightarrow \dots \rightarrow u_i \rightarrow v \rightarrow u_{i+1} \rightarrow \dots \rightarrow u_n$.




\square

Definition 10.12. A tournament is *strongly connected* if for all u, v there is a directed path from u to v .

Theorem 10.13. A tournament T is strongly connected if and only if it has a Hamilton cycle.


Proof. If T is Hamiltonian, then it is immediately clear that it is strongly connected: for any u and v simply consider the portion of a Hamilton path between u and v .

Now, suppose T is not Hamiltonian. Let C be a longest cycle in T and let $v \notin C$. If C has 2 consecutive vertices u, u' such that $u \rightarrow v$ and $v \rightarrow u'$, then there is a longer cycle on the vertex set $C \cup \{v\}$:



Otherwise all edges between v and C go in only one direction, and we can assume this holds for all $v \notin C$. Let A be the set of $v \notin C$ such that edges go from v to C , and let B be the set of $v \notin C$ such that edges go from C to v . If one of A or B are empty, it immediately follows that T is not strongly connected: if say B is empty then there is no path from any vertex of C to any vertex of A .

So, suppose both of A and B are nonempty. There is no edge oriented from a vertex $b \in B$ to a vertex $a \in A$, because if there was we could extend C , contradicting maximality (we could replace any edge $x \rightarrow y$ in C with the path $x \rightarrow b \rightarrow a \rightarrow y$). But then there is no path from any vertex in B to any vertex in A , so again T is not strongly connected.




□

11 Kuratowski's Theorem

In Corollary 6.16 we saw that K_5 and $K_{3,3}$ are not planar. In this section we will prove a remarkable theorem of Kuratowski, roughly saying that K_5 and $K_{3,3}$ are in some sense the *only* obstructions to being planar. We will also simultaneously prove Fáry's theorem, that any planar graph can in fact be drawn using only straight lines.

Definition 11.1. A *subdivision* of a graph H is a graph obtained from H by replacing the edges of H by internally vertex disjoint paths of non-zero length with the same endpoints.

Example 11.2.



Remark 11.3. If G contains a subdivision of H , it also contains an H -minor.

Definition 11.4. A *Kuratowski graph* is a graph which is a subdivision of K_5 or $K_{3,3}$. If G is a graph and H is a subgraph of G which is a Kuratowski graph then we say that H is a *Kuratowski subgraph* of G .

Theorem 11.5 (Kuratowski 1930). *A graph is planar if and only if it has no Kuratowski subgraph.*

It is straightforwardly true that no Kuratowski graph is planar. Indeed, given a planar drawing (by polygonal curves) of a Kuratowski subgraph, we can consider the subdivided paths to be long polygonal curves in a planar drawing of K_5 or $K_{3,3}$, which is impossible. So, to prove Theorem 11.5 we need only prove that every graph with no Kuratowski subgraph has a plane drawing. As promised, we prove a stronger theorem, combining Kuratowski's Theorem with Fáry's Theorem on straightline drawings.

Definition 11.6. A *straightline drawing* of a planar graph G is a drawing in which every edge is a straight line.

Theorem 11.7. *If G is a graph with no Kuratowski subgraph then G has a straightline drawing in the plane.*

The proof of Theorem 11.7 is broken down into two parts. In the first part, we prove a version of the theorem for 3-connected graphs, then in the second part we show how the general case can be reduced to the 3-connected case.

11.1 Convex drawings of 3-connected graphs

Definition 11.8. A *convex drawing* of G is a straightline drawing in which every non-outer face of G is a convex polygon, and the outer face is the complement of a convex polygon. (That is, the boundary of each face is the boundary of a convex polygon).

Theorem 11.9 (Tutte 1960). *If G is a 3-connected graph which has no Kuratowski subgraphs then G has a convex drawing in the plane with no three vertices on a line.*

We will need some lemmas to prove Theorem 11.9.


Lemma 11.10 (Thomassen 1980). *Every 3-connected graph G with at least five vertices has an edge e such that G/e is 3-connected.*

Proof. We proceed by contradiction. Suppose G has no edge whose contraction yields a 3-connected graph. For every edge $(x, y) \in E(G)$, the graph G is 3-connected but $G/(x, y)$ is not 3-connected. A minimal (2 vertex) separating set in $G/(x, y)$ must therefore contain the newly contracted vertex, so there exists a vertex $z \in V(G) \setminus \{x, y\}$ such that $\{x, y, z\}$ is a separating set in G . We call such a vertex z a *mate* of the edge (x, y) .

Choose an edge $(x, y) \in E(G)$ and a mate z and a connected component H in $G \setminus \{x, y, z\}$, in such a way that H has maximum possible size. Let H' be another connected component of $G \setminus \{x, y, z\}$. Since $\{x, y, z\}$ is a separating set of minimal size, each of x, y and z have a neighbour in both of H and H' . Let u be a neighbour of z in H' and let v be a mate of the edge (u, z) .

We claim now that the induced subgraph F of G with vertex set $V(H) \cup \{x, y\} \setminus v$ is connected. If that is the case then we reach a contradiction because $|F| > |H|$ and F is a connected subgraph of $G \setminus \{u, v, z\}$.

To prove the claim, we consider two cases: $v \notin V(H)$ or $v \in V(H)$. The first case is straightforward: since H is connected and since x and y each have a neighbour in H , the graph F must also be connected. Consider now the second case. The following picture gives an illustration of the situation.



Observe that $G \setminus \{v, z\}$ is connected since G is 3-connected. Thus, for any vertex $w \in V(H) \setminus \{v, z\}$ there exists a path in $G \setminus \{v, z\}$ from w to $\{x, y\}$. If we can find such a path for any $w \in V(H) \setminus \{v\}$ inside of F then it follows that F is connected.

We claim that we can just choose these paths to be the shortest paths (in $G \setminus \{v, z\}$) from each w to $\{x, y\}$. Indeed, since $\{x, y, z\}$ was the separating set with which we defined H , any path from w avoiding z must pass through x or y before it leaves H . \square


Lemma 11.11. *If G has no Kuratowski subgraphs, then G/e has no Kuratowski subgraph, for any edge $e \in E(G)$.*

Proof. We prove the contrapositive: if G/e contains a Kuratowski subgraph then so does G .

Let z be the vertex of G/e obtained by contracting $e = (x, y)$. Let H be a Kuratowski subgraph in G/e . If $z \notin V(H)$ then it is clear that H is also a subgraph of G . Otherwise, write $N_H(z) = A \cup B$, where the vertices in A were adjacent to x while the vertices in B were adjacent to y , before the contraction.

If $|A| = 0$ (respectively, $|B| = 0$) then we obtain a copy of H in G by replacing z with x (respectively y), in the copy of H in G/e . If $|A| = 1$ (respectively $|B| = 1$) we can replace z with y , add x and include the edge (x, y) and the edge from x to A , to obtain a copy of H with one of its edge subdivided (specifically, the edge from z to A , respectively from z to B).

Since H is a subdivision of $K_{3,3}$ or K_5 , every vertex has degree 2,3 or 4. The only possibility that we have not accounted for is that $d_H(z) = |N_H(z)| = 4$, and $|A| = |B| = 2$. We can then obtain a subdivision of $K_{3,3}$ in G from H by replacing z with the edge (x, y) and deleting the inner vertices of two of the paths in H . This is illustrated below; the thicker lines indicate subdivided paths, and the four vertices directly above z are in A , A , B and B respectively.



\square


Proof of Theorem 11.9. Let G be a 3-connected graph which has no Kuratowski subgraphs. We prove that G has a convex drawing in the plane with no three vertices on a line, by induction on $|G|$.

The base case is $|G| = 4$. In that case, since G is 3-connected, the only possibility is that $G = K_4$ and it is easy to find a suitable drawing (see for example Example 6.2).

Consider now $|G| \geq 5$. The induction hypothesis says that every 3-connected graph on at most $|G| - 1$ vertices without Kuratowski subgraphs has a convex drawing in the plane with no three vertices on a line.

By Lemma 11.10 there exists an edge $e = (x, y) \in E(G)$ such that G/e is 3-connected. Moreover, by Lemma 11.11 G/e does not have any Kuratowski subgraphs. From the induction hypothesis it follows that G/e has a convex drawing in the plane with no three vertices on a line. Fix such a convex drawing and let z be the vertex of G/e obtained by contracting the edge e in G . Note that if we delete z from G/e we get a graph which is 2-connected and hence its faces are cycles. Let C be the cycle formed by the boundary of the face of $(G/e) \setminus z$ which contains z (this could be the outer face).

Observe that all the neighbours of z in G/e (and thus, also of x and y in G) lie in C . Moreover, since G is 3-connected, x and y each have at least 2 neighbours in C . Let z_1, \dots, z_k be the neighbours of z in G/e ordered cyclically (according to an orientation of C) and let $z_{i_1}, \dots, z_{i_\ell}$ denote the neighbours of x in G (in the same order). If the neighbours of y all fall in between two neighbours of x (that is, all are in $\{z_{i_j}, z_{i_{j+1}}, \dots, z_{i_{j+1}}\}$ for some $j \in \{1, \dots, \ell\}$) then we can obtain a convex drawing of G with no three vertices on a line by putting x at z and putting y at a point close to z in the wedge formed by (x, z_{i_j}) and $(x, z_{i_{j+1}})$. Note that it is possible for the angle between (x, z_{i_j}) and $(x, z_{i_{j+1}})$ to be a reflex angle, as illustrated below (the vertices marked \otimes correspond to neighbours of y). In this case we need to put our point in the region formed by continuing the lines z_{i_j}, z and $z_{i_{j+1}}, z$, shaded in the picture below. This is so that we can draw the triangles x, y, z_{i_j} and $x, y, z_{i_{j+1}}$ without intersecting existing lines, if necessary.




In general, the reason it is always possible to choose an appropriate position for y very close to z is that if we take a convex polygon with no three points on a line, and we perturb one of the vertices slightly, then the result is still a convex polygon. Note also that it is always possible to choose our new point such that still no three vertices are on a line, because the set of bad positions has measure zero.

We claim the neighbours of y are always situated in between two neighbours of x , so the above argument is always valid. Indeed, otherwise, one of the following two cases holds:

- y shares three neighbours z_{j_1}, z_{j_2} and z_{j_3} with x . In this case, the cycle C together with (x, y) and the edges from $\{x, y\}$ to $\{z_{j_1}, z_{j_2}, z_{j_3}\}$ forms a subdivision of K_5 in G (left image below).

- y has neighbours z_{j_1} and z_{j_3} which alternate with two neighbours z_{j_2} and z_{j_4} of x in the cycle C . In that case, the cycle C together with (x, y) , (x, z_{j_2}) , (x, z_{j_4}) , (y, z_{j_1}) and (y, z_{j_3}) forms a subdivision of $K_{3,3}$ in G (right image below):




Since G has no Kuratowski subgraph, both of these cases are impossible. □

11.2 Reducing the general case to the 3-connected case


Definition 11.12. Given a subdivision H' of H , we call the vertices of the original graph *branch vertices*.

Example 11.13. In the following picture, the branch vertices are marked with an \otimes .



Fact. We make three observations. See the picture below.

1. In a Kuratowski subgraph, there are three internally vertex-disjoint paths connecting any two branch vertices. For K_5 -subdivisions, we even have four such paths.
2. In a Kuratowski subgraph, there are four internally vertex-disjoint paths between any two pairs of branch vertices.
3. Any cycle in a subdivision contains at least three branch vertices.



Proposition 11.14. Let G be a graph with at least 4 vertices which has no Kuratowski subgraph, and suppose that adding an edge joining any pair of non-adjacent vertices creates a Kuratowski subgraph. Then G is 3-connected.

Theorem 11.7 is a corollary of Proposition 11.14, as follows.


Proof of Theorem 11.7. Start with a graph G which has no Kuratowski subgraph, and add edges as long as possible without creating a Kuratowski subgraph. Let the final graph be $G' \supseteq G$. Then G' is 3-connected so by Theorem 11.9 it has a planar convex drawing. Therefore G has a planar straightline drawing. \square

Proof of Proposition 11.14. The proof is by induction on $|G|$. If $|G| = 4$ then G is complete so we are done. So assume $|G| \geq 5$.


First, note that G is connected. Otherwise, we could add an edge between connected components, and this cannot create a Kuratowski subgraph because Kuratowski graphs have no bridges.

Next, we prove that G is 2-connected. Suppose otherwise, that we can write $G = G_1 \cup G_2$ in such a way that $G_1 \cap G_2 = v$ for some vertex v , each G_i contains more vertices than just v , and there are no edges from $G_1 \setminus v$ to $G_2 \setminus v$. We claim that the addition of any edge to each G_i creates a Kuratowski subgraph, so we can apply the induction hypothesis to each. Indeed, consider $G_i \cup \{e\}$ for some edge e between vertices of G_i . By assumption, $G \cup \{e\}$ contains a Kuratowski subgraph, which is 2-connected. Therefore deleting v (which separates G_1 from G_2) cannot disconnect the Kuratowski subgraph, so this subgraph must have been entirely contained in G_i .


By the induction hypothesis and Theorem 11.9, each G_i has a planar convex drawing (if either has fewer than 4 vertices, then it trivially has such a drawing). Now, consider any $x \in G_1 \setminus v$, $y \in G_2 \setminus v$. The addition of $e = (x, y)$ to G creates a Kuratowski subgraph.



By the first observation above, a branch vertex cannot be disconnected from the other branch vertices by deleting an edge and a vertex. Hence, the branch vertices of this subdivision are either all in G_1 or all in G_2 (say G_1). Therefore the only part of the Kuratowski subgraph in $G \cup \{e\}$ which is not contained in G_1 , is a path between v and x .



Hence, if we start with G_1 and add any path from v to any x , then the resulting graph also has a Kuratowski subgraph. But if we choose x to be a neighbour of v , then we can add such a path to a plane drawing of G_1 while preserving planarity.




This is impossible because no planar graph has a Kuratowski subgraph. It follows that G is 2-connected.

The final step is to show that G is 3-connected. Suppose otherwise, so that we can write $G = G_1 \cup G_2$ with $G_1 \cap G_2 = \{u, v\}$, where each G_i contains more than just the vertices u, v , and there are no edges from $G_1 \setminus \{u, v\}$ to $G_2 \setminus \{u, v\}$. We know that G is 2-connected so we can assume each G_i is connected (each of u and v must have a neighbour in each component of each $G_i \setminus \{u, v\}$, otherwise deleting one of u and v would disconnect G).


Now, we claim that $(u, v) \in E(G)$. If not, then $G \cup \{(u, v)\}$ contains a Kuratowski subgraph. The branch vertices of this subgraph must be entirely in G_1 or in G_2 (say G_1) because, by the first observation, deleting 2 vertices from a Kuratowski graph cannot disconnect a pair of branch vertices. Thus $G_2 \cup \{(u, v)\}$ contains at most 2 branch vertices, so by the third observation, the intersection of the Kuratowski subgraph with $G_2 \cup \{(u, v)\}$ is cycle-free, i.e. it is just the edge (u, v) . On the other hand, G_2 is connected, so it contains a u - v path (disjoint from the Kuratowski subgraph). We can view that path as a subdivision of (u, v) showing that G contains a subdivision of $G_1 \cup \{(u, v)\}$, whence G contains a Kuratowski subgraph as well.

Next, we claim that adding any edge to each G_i creates a Kuratowski subgraph, so we can apply the induction hypothesis. Indeed, consider $G_1 \cup \{e\}$ for some edge e between vertices of G_1 . By assumption, $G \cup \{e\}$ contains a Kuratowski subgraph, but deleting u and v disconnects $G \cup \{e\}$. Therefore $G_2 \setminus \{u, v\}$ contains no branch vertices, and since $(u, v) \in E(G)$, a Kuratowski subgraph is also contained in $G_1 \cup \{e\}$.

By the induction hypothesis and Theorem 11.9 (and by considering separately the case $|G_i| = 3$), each G_i has a convex drawing. For each i choose $w_i \in V(G_i)$ such that u, v and w_i belong to the boundary of some convex face of G_i . By hypothesis the graph $G \cup \{(w_1, w_2)\}$ contains a Kuratowski subgraph K . Note that, by the second observation, deleting two vertices and an edge cannot disconnect these pairs from one another, and it follows that all but possibly one of the branch vertices are contained in some G_i (say G_1). Note that we can use the same reasoning as in the proof of 2-connectedness to see that in fact G_2 has exactly one branch vertex x . Indeed, otherwise G_1 contains all the branch vertices, so the other vertices of K can only be on a path from w_1 to u or from w_1 to v . But we can add such a path to our convex drawing of G_1 (as illustrated below), which contradicts the nonplanarity of K .



Note that x is separated from the other branch vertices of K by the 3-element set $\{u, v, w_1\}$; by the first observation, K must be a subdivision of $K_{3,3}$. It also follows that if in $G \cup \{(w_1, w_2)\}$ we identify all vertices of $V(G_2) \setminus V(G_1)$ into a single vertex, the new graph also contains a subdivision of $K_{3,3}$. But we can construct a planar drawing of this graph, which is a contradiction! Indeed, u , v and w_1 are on the boundary of some convex face of G_1 . Put the new vertex anywhere inside this face and connect it to u , v and w_1 with straight lines, as illustrated below.




We conclude that G is 3-connected, as required. □

12 Ramsey Theory

Ramsey theory refers to a large body of deep results in mathematics concerning partitions of large collections. An informal way of describing this is Motzkin’s statement that “Complete disorder is impossible”. Although this does not always seem to be the case in the world around us, in mathematics this principle appears in great generality.

Proposition 12.1. *Among six people it is possible to find three mutual acquaintances or three mutual non-acquaintances.*

Proof. Consider the complete graph whose vertices are the six people. Colour an edge between two people red if those people know each other, and blue otherwise. Single out some vertex u . Out of the five edges incident to u , at least three are blue or at least three are red. Without loss of generality say there are three red edges (u, v_1) , (u, v_2) , (u, v_3) . If there is a red edge (v_i, v_j) , then u , v_i and v_j are mutual acquaintances. Otherwise v_1 , v_2 and v_3 are mutual non-acquaintances.



As we shall see, given a natural number s , there is an integer R such that if $n \geq R$ then every colouring of the edges of K_n with red and blue contains either a red K_s or a blue K_s . The assertion about a party of six people claims that for $s = 3$, $R = 6$ will do. More generally, we define the *Ramsey number* $R(s, t)$ as the smallest value of N for which every red-blue colouring of K_N yields a red K_s or a blue K_t . In particular, $R(s, t) = \infty$ if there is no such N such that in every red-blue colouring of K_N there is a red K_s or a blue K_t . It is obvious that

$$R(s, t) = R(t, s)$$

for every $s, t \geq 2$ and $R(s, 2) = R(2, s) = s$, since in a red-blue colouring of K_s either there is a blue edge or else every edge is red. The following result, due to Erdős and Szekeres, states that $R(s, t)$ is finite for every s and t , and at the same time it gives a bound on $R(s, t)$. This bound is considerably better than that given by Ramsey.

Theorem 12.2. *The function $R(s, t)$ is finite for all $s, t \geq 2$. Quantitatively, if $s > 2$ and $t > 2$ then*

$$R(s, t) \leq R(s - 1, t) + R(s, t - 1) \tag{12}$$

and

$$R(s, t) \leq \binom{s + t - 2}{s - 1}. \tag{13}$$

Proof. When proving (12) we may assume that $R(s - 1, t)$ and $R(s, t - 1)$ are finite. Let $N = R(s - 1, t) + R(s, t - 1)$ and consider a colouring of the edges of K_N with red and blue. We have to show that in this colouring there is either a red K_s or a blue K_t . To this end, let x be a vertex of K_N . Since $d(x) = N - 1$, either there are at least $N_1 = R(s - 1, t)$ red edges incident with x or there are at least $N_2 = R(s, t - 1)$ blue edges incident with x . Without loss of generality assume that the first case holds. Consider a subgraph K_{N_1} of K_N spanned by N_1 vertices joined to x by red edges. If K_{N_1} has a blue K_t , we are done. Otherwise, by the definition of $R(s - 1, t)$, the graph K_{N_1} contains a red K_{s-1} which forms a red K_s with x .

Inequality (13) holds if $s = 2$ or $t = 2$ (in fact, we have equality). Assume that $s > 2$, $t > 2$, and that (13) holds for every pair (s', t') with $2 \leq s' + t' < s + t$. Then by (12) we have

$$\begin{aligned} R(s, t) &\leq R(s - 1, t) + R(s, t - 1) \\ &\leq \binom{s + t - 3}{s - 2} + \binom{s + t - 3}{s - 1} = \binom{s + t - 2}{s - 1}. \end{aligned}$$

□

Remark. Note that $R(s, t) \leq N$ means that every graph on N vertices has either $\omega(G) \geq s$ or $\alpha(G) \geq t$.

It is customary to distinguish *diagonal Ramsey numbers* $R(s) = R(s, s)$ and *off-diagonal Ramsey numbers* $R(s, t)$ for $s \neq t$. It is not surprising that the diagonal Ramsey numbers are of greatest interest, and they are also the hardest to estimate.

We see from Theorem 12.2 that

$$R(s) \leq \binom{2s-2}{s-1} \leq \frac{2^{2s-2}}{\sqrt{s}}. \quad (14)$$

Although the proof above is very simple, the bound (14) was hardly improved over the last 50 years.

The result easily extends to colourings with any finite number of colours.

Theorem 12.3. *Given k and s_1, s_2, \dots, s_k , if N is sufficiently large, then every colouring of K_N with k colours is such that for some i , $1 \leq i \leq k$, there is a K_{s_i} coloured with the i -th colour. The minimal value of N for which this holds is usually denoted by $R_k(s_1, \dots, s_k)$, and it satisfies*

$$R_k(s_1, \dots, s_k) \leq R_{k-1}(R(s_1, s_2), s_3, \dots, s_k).$$

Proof. Assume the result is true for $k-1$ colours. The case $k=2$ is resolved by Theorem 12.2.

In a k -colouring of K_N we replace the first two colours by a single new colour (say, purple). If $N = R_{k-1}(R(s_1, s_2), s_3, \dots, s_k)$, then either there is a K_{s_i} coloured with the i -th colour for some i , $3 \leq i \leq k$, or else for $m = R(s_1, s_2)$ there is a purple K_m . In the original colouring this K_m is coloured with the first two original colours. In the first case we are done, and in the second, for $i=1$ or 2 we can find a K_{s_i} in K_m coloured with the i -th colour. This shows that $R_k(s_1, \dots, s_k) \leq R_{k-1}(R(s_1, s_2), s_3, \dots, s_k)$. \square

In fact, Theorem 12.2 also extends to hypergraphs, that is, to colourings of the set $X^{(r)}$ of all r -tuples of a finite set X with k colours. This is one of the theorems proved by Ramsey. We now turn our attention to this.

Denote by $R^{(r)}(s, t)$ the minimal value of N for which every red-blue colouring of $X^{(r)}$ yields a red s -set or a blue t -set, provided that $|X| = N$. By analogy with graphs, a set $Y \subseteq X$ is called *red* (*blue*) if every element of $Y^{(r)}$ is red (blue). Note that $R(s, t) = R^{(2)}(s, t)$. Just as for Theorem 12.2, we can show that each $R^{(r)}(s, t)$ is finite and give simple bounds. As an example, we demonstrate the case $r=3$. Note that if $\min\{s, t\} < 3$, then $R^{(3)}(s, t) = \min\{s, t\}$, and if $\min\{s, t\} = 3$ then $R^{(3)}(s, t) = \max\{s, t\}$.

Theorem 12.4. *Let $\min\{s, t\} > 3$. Then*

$$R^{(3)}(s, t) \leq R(R^{(3)}(s-1, t), R^{(3)}(s, t-1)) + 1.$$


Proof. Let X be a set with $R(R^{(3)}(s-1, t), R^{(3)}(s, t-1)) + 1$ elements. Given any red-blue colouring c of $X^{(3)}$, arbitrarily pick an $x \in X$. Let K_x be the complete graph on the vertex set $X \setminus \{x\}$ and define an edge colouring c' of K_x by $c'((x, y)) = c(\{x, y, z\})$. Without loss of generality, suppose that $X \setminus \{x\}$ has a red (in c') subset Z with $R^{(3)}(s-1, t)$ elements.

Now let us look at the restriction of c to $Z^{(3)}$. If it has a blue t -set, we are done, since a blue t -set of Z is also a blue t -set of X . On the other hand, if there is no blue t -set of Z then there is a red $(s - 1)$ -set. The union of this red $(s - 1)$ -set with $\{x\}$ is then a red s -set of X , because $\{x\} \cup e$ is red for every edge e in $K_x[Z]$. \square


12.1 Applications

Theorem 12.5 (Erdős-Szekeres 1935). *Given an integer m , there exists a (least) integer $N(m)$ such that every set of at least $N(m)$ points in the plane, with no three collinear, contains an m -subset forming a convex m -gon.*

Proof. We need two facts. (1) *Among five points in the plane, there are four which determine a convex quadrilateral* (if no three are collinear). To see this, construct the convex hull of the five points. If it is a pentagon or a quadrilateral, the result follows immediately. If it is a triangle, the other two points lie inside. By the pigeonhole principle, two of the vertices of the triangle are on one side of the line through the two inside points. These two vertices together with the two points inside form a convex quadrilateral, as illustrated below.



In a convex m -gon, any four points determine a convex quadrilateral; we need the converse: (2) *If every 4-subset of m points in the plane form a convex quadrilateral, then the m points form a convex m -gon.* For, if the claim fails, then the convex hull of the m points consists of t points, for some $t < m$. The remaining points lie inside the t -gon. When we triangulate the t -gon, as illustrated, a point inside lies in one of the triangles. With the vertices of that triangle, it forms a 4-set that does not determine a convex quadrilateral.



To prove the theorem, let $N = R^{(4)}(m, 5)$. Given N points in a plane, colour each 4-set by convexity: red if it determines a convex quadrilateral, blue if it does not. By fact (1), there cannot be 5 points whose 4-subsets are all blue. By Ramsey's Theorem, this means that there are m points whose 4-subsets are all red. By fact (2), they form a convex m -gon. Hence $N(m)$ exists and is at most $R^{(4)}(m, 5)$. \square

12.2 Bounds on Ramsey numbers

Theorem 12.6 (Erdős 1947). *We have $R(p, p) > 2^{p/2}$.*

Proof. Consider the graphs with vertex set $[n]$. Each particular p -clique occurs in $2^{\binom{n}{2} - \binom{p}{2}}$ of these $2^{\binom{n}{2}}$ graphs. Similarly, each particular set of p vertices occurs as an independent set in $2^{\binom{n}{2} - \binom{p}{2}}$ of these graphs. Discarding these leaves only graphs having no p -clique or independent p -set. Since there are $\binom{n}{p}$ ways to choose p vertices, the inequality $2^{\binom{n}{2} - \binom{p}{2}} < 1$ would imply $R(p, p) > n$. Rough approximations yield $\binom{n}{p} 2^{1 - \binom{p}{2}} < 1$ whenever $n < 2^{p/2}$. \square

Theorem 12.7. *We have*

$$R_k(3) \stackrel{\text{def}}{=} R_k(3, 3, \dots, 3) \leq \lfloor e \cdot k! \rfloor + 1.$$

Proof. We use induction on k . For $k = 2$, we have already shown $R_2(3, 3) \leq 6$ (here equality holds).

Suppose $k \geq 3$ and let x be any point of a k -coloured complete graph on $\lfloor e \cdot k! \rfloor + 1$ points. There are $\lfloor e \cdot k! \rfloor$ edges adjacent to this vertex, which are split into k colour classes. Note that

$$\lfloor e \cdot k! \rfloor = \left\lfloor \sum_{j=0}^{\infty} \frac{k!}{j!} \right\rfloor = \sum_{j=0}^k \frac{k!}{j!} = 1 + k \sum_{j=1}^{k-1} \frac{(k-1)!}{j!} = 1 + k \lfloor e(k-1)! \rfloor.$$

So, one of the the k colours, say red, contains at least $\lfloor e \cdot (k-1)! \rfloor + 1$ edges adjacent to x . Let S be the set of those vertices joined to x by a red edge.

If S spans a red edge this forms a red triangle together with x , and we are finished. If S spans no red edge it spans a complete graph with $\lfloor e \cdot (k-1)! \rfloor + 1$ points, whose edges are $(k-1)$ -coloured; thus by induction hypothesis, one of these $k-1$ colours contains a triangle and we are done again. \square

12.3 Ramsey theory for integers

The following theorem, proved by Schur in 1916, became the starting point of an area that is still very active today.

Theorem 12.8. *For every $k \geq 1$ there is an integer m such that every k -colouring of $[m]$ contains integers x, y, z of the same colour such that $x + y = z$.*

Proof. We claim that we can choose $m = R_k(3) - 1$. Let $n = R_k(3) = m + 1$ and let $c : [m] \rightarrow [k]$ be a k -colouring. This induces a k -edge-colouring c' of the complete graph with vertex set $[n]$, as follows. For $(i, j) \in E(K_n)$ set $c'(i, j) = c(|i - j|)$. By the definition of $n = R_k(3)$, there is a monochromatic triangle, say with vertex set $\{h, i, j\}$, so that $1 \leq h < i < j \leq n$ and $c'(h, i) = c'(i, j) = c'(j, h)$ ($= \ell$, say). But then $x = i - h, y = j - i$ and $z = j - h$ are such that $c(x) = c(y) = c(z) = \ell$ and $x + y = z$. \square

12.4 Graph Ramsey numbers

Definition 12.9. Let G_1, G_2 be graphs. $R(G_1, G_2)$ is the minimal N such that any red/blue colouring of K_N contains either a red copy of G_1 , or a blue copy of G_2 .

Remark 12.10. Note that $R(G_1, G_2) \leq R(|G_1|, |G_2|)$.

Theorem 12.11 (Chvátal 1977). *If T is any m -vertex tree, then $R(T, K_n) = (m - 1)(n - 1) + 1$.*

Proof. For the lower bound, colour $K_{(m-1)(n-1)}$ with $n - 1$ disjoint red $(m - 1)$ -cliques and every other edge blue. With red components of order $m - 1$, there is no red m -vertex tree. The blue edges form an $n - 1$ -partite graph and hence cannot contain K_n .

For the upper bound, we proceed by induction on n (the case $n = 1$ is trivial). Given a two-colouring of $K_{(m-1)(n-1)+1}$, consider a vertex x . If x has more than $(m - 1)(n - 2)$ neighbours along blue edges, then among them is a red T or a blue K_{n-1} , which yields a blue K_n with x . Otherwise, every vertex has at most $(m - 1)(n - 2)$ incident blue edges and at least $m - 1$ incident red edges. That is, the red subgraph R has minimum degree at least $m - 1$. But this means we can embed a copy of T . Indeed, we claim that if a graph R has minimum degree $m - 1$ then it contains every m -vertex tree T .

We prove this claim by induction. If $m \leq 1$ the claim is obvious, so suppose $T \geq 2$ and that the claim holds for every tree smaller than T . Let ℓ be a leaf of T ; by the induction hypothesis R contains $T \setminus \ell$ (which has $m - 1$ vertices). Since every vertex has degree at least $m - 1$, every vertex of $T \setminus \ell$ has a neighbour that is not in $T \setminus \ell$. Hence we can “reattach” ℓ to T using such a neighbour. \square

13 Extremal problems

Question 13.1. Let H be a fixed graph, and G be a graph on n vertices that contains no copy of H . How many edges can G have?

Definition 13.2. $\text{ex}(n, H)$ is the maximal value of $e(G)$ among graphs G with n vertices containing no H as a subgraph.

Example 13.3. Consider the case where H is a triangle. Recall that bipartite graphs contain no triangles. So $K_{\frac{n}{2}, \frac{n}{2}}$ gives a triangle-free graph with $\frac{n^2}{4}$ edges.

13.1 Turán’s theorem

As a generalization Example 13.3, notice that dense graphs not having K_{r+1} as a subgraph can be obtained by dividing the vertex set V into r pairwise disjoint subsets $V = V_1 \cup \dots \cup V_r$, $|V_i| = n_i$, $n = n_1 + \dots + n_r$, joining two vertices if and only if they lie in distinct sets V_i, V_j . We denote the resulting graph by K_{n_1, \dots, n_r} . It has $\sum_{i < j} n_i n_j$ edges. Assuming n is fixed, we get the maximal number of edges among such graphs when we divide the numbers n_i as evenly as possible, that is $|n_i - n_j| \leq 1$ for all i, j . Indeed, suppose $n_1 \geq n_2 + 2$. By shifting one vertex from V_1 to V_2 , we

obtain $K_{n_1-1, n_2+1, \dots, n_r}$, and that contains $(n_1 - 1)(n_2 + 1) - n_1 n_2 = n_1 - n_2 - 1 \geq 1$ more edges than K_{n_1, \dots, n_r} . In particular, if r divides n , then we may choose $n_i = \frac{n}{r}$ for all i , obtaining

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

edges. Turán's theorem states that this number is an upper bound for the edge-number of *any* graph on n vertices without an $(r + 1)$ -clique.

Definition 13.4. We call the graph K_{n_1, \dots, n_r} with $|n_i - n_j| \leq 1$ the *Turán graph*, denoted by $T_{n,r}$.

Theorem 13.5 (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.*

First Proof. Let $G = (V, E)$ have the maximum number of edges among graphs containing no $(r + 1)$ -cliques. Let $v_m \in V$ be a vertex of maximal degree $\Delta(G)$. Denote the set of neighbors of v_m by S , $|S| = d_m$, and set $T = V \setminus S$. As G contains no $(r + 1)$ -clique and v_m is adjacent to all vertices of S , we note that S contains no r -clique.

We now construct the following graph H on V . H corresponds to G on S and contains all edges between S and T , but no edges within T . In other words, T is an independent set in H , and it follows that H again has no r -cliques. If $v \in S$, then we certainly have $d_H(v) \geq d_G(v)$ by the construction of H , and for $v \in T$ we see $d_H(v) = |S| = \Delta(G) \geq d_G(v)$ by the choice of v_m . We infer $|E(H)| \geq |E(G)|$ and find that among all graphs with a maximal number of edges, there must be one of the form of H . By induction, the graph induced by S has at most as many edges as a suitable graph $K_{n_1, \dots, n_{r-1}}$ on S . So $|E(G)| \leq |E(H)| \leq E(K_{n_1, \dots, n_r})$ with $n_r = |T|$. We have established that $E(K_{n_1, \dots, n_r})$ is maximized for $|n_i - n_j| \leq 1$, which implies that $T_{n,r}$ has the maximum possible number of edges.

To prove uniqueness, note that equality can only hold in our previous bound if S induces the complete $r - 1$ -partite graph $K_{n_1, \dots, n_{r-1}}$ and T touches exactly Δn_r edges in G . But the latter can only happen if T is an independent set in G . Indeed, the sum of the degrees of the vertices in T counts each edge spanned by T twice, and each edge connecting T and S once. As Δ is the maximum degree in G , the sum of degrees is at most Δn_r , so T 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, $G = T_{n,r}$. \square

Second Proof. In this second proof we only prove maximality (not uniqueness), and only for the case where r divides n .

Consider a weighting $w = (w_1, \dots, w_n)$ of the vertices such that $w_i \geq 0$ for every vertex v_i , and $\sum_{i=1}^n w_i = 1$. Our goal is to maximize the function

$$f(w) = \sum_{v_i \sim v_j} w_i w_j.$$

(This maximum is called the *Lagrangian* of G). Suppose w is any such weighting, and let v_i and v_j be a pair of non-adjacent vertices with positive weights w_i and w_j . Let s_i be the sum of the weights of all vertices adjacent to v_i , and define s_j similarly for v_j , where we may assume that $s_i \geq s_j$. Now

we move the weight from v_j to v_i ; that is, the new weight of v_i is $w_i + w_j$, while the weight of v_j drops to 0. For the new weighting w' we find

$$f(w') = f(w) + w_j s_i - w_j s_j \geq f(w).$$

We repeat this (reducing the number of vertices with a positive weight by one in each step) until there are no non-adjacent vertices of positive weight anymore. Thus we conclude that there is an optimal weighting whose nonzero weights are concentrated on a clique, say on a k -clique. Now if, say, $w_1 > w_2 > 0$, then choose ε with $0 < \varepsilon < w_1 - w_2$ and change w_1 to $w_1 - \varepsilon$ and w_2 to $w_2 + \varepsilon$. The new weighting w' satisfies $f(w') = f(w) + \varepsilon(w_1 - w_2) - \varepsilon^2 > f(w)$, and we infer that the maximal value of $f(w)$ is attained for $w_i = \frac{1}{k}$ on a k -clique and $w_i = 0$ otherwise. Since a k -clique contains $\frac{k(k-1)}{2}$ edges, we obtain

$$f(w) = \frac{k(k-1)}{2} \cdot \frac{1}{k^2} = \frac{1}{2} \left(1 - \frac{1}{k}\right).$$

Since this expression is increasing in k , the best we can do is to set $k = r$ (since G has no $(r+1)$ -cliques). So we conclude

$$f(w) \leq \frac{1}{2} \left(1 - \frac{1}{r}\right)$$

for *any* weighting w . In particular, this inequality holds for the uniform weighting w^* given by $w_i = \frac{1}{n}$ for all i . Thus we find

$$|E| = f(w^*)n^2 \leq \left(1 - \frac{1}{r}\right) \frac{n^2}{2} = e(T_{n,r})$$

when r divides n . □

As an application of Turán's theorem we answer the following question.

Question 13.6. Let $a_1, \dots, a_n \in \mathbb{R}^d$ be vectors such that $|a_i| \geq 1$ for each $i \in [n]$. What is the maximum number of pairs satisfying $|a_i + a_j| < 1$?

Claim 13.7. There are at most $\lfloor \frac{n^2}{4} \rfloor$ such pairs.

Proof. Define the graph G on $[n]$ where $i \sim j$ 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 $i, j, k \in [n]$,

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

so at least one of $|a_i + a_j|^2, |a_j + a_k|^2, |a_k + a_i|^2$ is at least 1. □


Definition 13.8. For some fixed graph H , we define $\pi(H) = \lim_{n \rightarrow \infty} \text{ex}(n, H) / \binom{n}{2}$. (One can prove that $\text{ex}(n, H) / \binom{n}{2}$ is monotone decreasing in n , hence the limit exists for any H).

Turán's theorem says that $\pi(K_r) = 1 - \frac{1}{r-1}$. More generally, we have the following theorem, which we present without proof.


Theorem 13.9 (Erdős-Stone). *Let H be a graph of chromatic number $\chi(H) = r + 1$. Then for every $\varepsilon > 0$ and large enough n ,*

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

Example 13.10. The following graph H has $\chi = 3$, so the maximum number of edges in an H -free graph is $(1 + o(1))\frac{n^2}{4}$.



If we take H to be the graph pictured below (the isocahedron) then $\chi(H) = 4$, so the maximum number of edges in an H -free graph is $(1 + o(1))\frac{n^2}{3}$.




Note that if $\chi(H) = 2$ (that is, if H is bipartite) then the Erdős-Stone theorem simply says that $\text{ex}(n, H) \leq \varepsilon n^2$ for any ε , and does not give precise asymptotics. It is interesting to ask what more we can say in this case.

13.2 Bipartite Turán Theorems

Theorem 13.11. *If a graph G on n vertices contains no 4-cycles, 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. We count the following set S in two ways: S is the set of pairs $(u, \{v, w\})$ where u is adjacent to v and to w , with $v \neq w$. In other words, we count all occurrences of



Summing over u , we find $|S| = \sum_{u \in V} \binom{d(u)}{2}$. On the other hand, every pair $\{v, w\}$ has at most one common neighbor (because G is C_4 -free). Hence $|S| \leq \binom{n}{2}$, and we conclude

$$\sum_{u \in V} \binom{d(u)}{2} \leq \binom{n}{2}$$

or equivalently,

$$\sum_{u \in V} d(u)^2 \leq n(n-1) + \sum_{u \in V} d(u). \quad (15)$$

Next (and this is quite typical for this sort of extremal problem), we apply the Cauchy-Schwarz inequality to the vectors $(d(u_1), \dots, d(u_n))$ and $(1, \dots, 1)$, obtaining

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

and hence, by (15),

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

Using the fact that the sum of the degrees is $2|E|$, we find

$$4|E|^2 \leq n^2(n-1) + 2n|E|.$$

Equivalently,


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

Solving this quadratic equation yields the theorem. \square

Example 13.12. Let G_0 be the graph on the vertex set $\mathbb{Z}_p \times \mathbb{Z}_p$ ($p \geq 3$ prime) where $(x, y) \sim (x_1, y_1)$ if and only if $x + x_1 = yy_1$. (Technically this is a multigraph as it has loops).

Note that G_0 is p -regular. Indeed, for every x, y, y_1 there is a unique choice of x_1 such that $(x, y) \sim (x_1, y_1)$. Also, loops correspond to solutions of the equation $2x = y^2$. There is therefore one loop for every choice of y , comprising p loops total. Let G be the graph we obtain by deleting the loops from G_0 .

Now, G has $n = p^2$ vertices, and $\frac{1}{2}(np - p) = (\frac{1}{2} + o(1))n^{3/2}$ edges. It also has no C_4 s. Indeed, for any (x_1, y_1) and (x_2, y_2) , a vertex (x, y) adjacent to both of them satisfies $x + x_1 = yy_1$ and $x + x_2 = yy_2$, so $x_1 - x_2 = y(y_1 - y_2)$. If $y_1 = y_2$ then $x_1 = x_2$, so if our chosen vertices (x_1, y_1) and (x_2, y_2) were distinct then $y_1 - y_2 \neq 0$ and there is a unique (x, y) satisfying the equations.



Now we prove the following more general theorem.

Theorem 13.13 (Kővári-Sós-Turán). *For any integers $r \leq s$, there is a constant c such that every $K_{r,s}$ -free graph on n vertices contains at most $cn^{2-\frac{1}{r}}$ edges. In other words,*

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

We will prove this theorem using (a special case of) Jensen's inequality, which says that if f is a convex function then for any x_1, \dots, x_n we have

$$f(x_1) + \dots + f(x_n) \geq nf\left(\frac{x_1 + \dots + x_n}{n}\right).$$

Proof of Theorem 13.13. Let $m = |E|$, $V = \{x_1, \dots, x_n\}$ and $d(x_i) = d_i$. Since G contains no $K_{r,s}$, for any given r -tuple in V there are at most $s-1$ vertices whose neighbourhood contains that r -tuple. The neighbourhood of x_i contains $\binom{d_i}{r}$ r -tuples, so

$$\sum_{i=1}^n \binom{d_i}{r} \leq (s-1) \binom{n}{r}.$$

We are going to estimate the left-hand side using Jensen's inequality. Note that the function $f_0 : x \mapsto \binom{x}{r}$ is convex for $x \geq r-1$. Indeed, by the product rule,

$$f_0''(x) = \frac{2}{r!} \sum_{0 \leq i < j \leq r-1} \prod_{\ell=0}^{r-1} (x-\ell) \frac{1}{(x-i)(x-j)} \geq 0.$$

We can therefore define the convex function

$$f(x) = \begin{cases} \binom{x}{r} & \text{if } x \geq r-1 \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Since d_i is an integer, we have

$$\sum_{i=1}^n \binom{d_i}{r} = \sum_{i=1}^n f(d_i) \geq n \cdot f\left(\frac{d_1 + \dots + d_n}{n}\right) = nf\left(\frac{2m}{n}\right).$$

Now, if $\frac{2m}{n} < r-1$, we have nothing to prove. Suppose $\frac{2m}{n} \geq r-1$, then we have

$$(s-1) \binom{n}{r} \geq \sum_{i=1}^n \binom{d_i}{r} \geq nf\left(\frac{2m}{n}\right) = n \binom{2m/n}{r}.$$

Note that

$$(s-1) \frac{n^r}{r!} > (s-1) \binom{n}{r} \quad \text{and} \quad \binom{2m/n}{r} > \frac{\left(\frac{2m}{n} - r + 1\right)^r}{r!},$$

so we obtain $((2m/n) - r + 1)^r < (s-1)n^{r-1}$. But then

$$2m < (s-1)^{1/r} n^{2-1/r} + (r-1)n < c \cdot n^{2-1/r}$$

for some c , finishing the proof. □

Theorem 13.14. *There is c depending on k such that if G is a graph on n vertices that contains no copy of C_{2k} , then G has at most $cn^{1+\frac{1}{k}}$ edges.*

This theorem has been proved to be tight for $k = 2, 3$ and 5 .

Consider the following question, which we will see can be answered with the Kővári-Sós-Turán theorem.

Question 13.15. Given n points in the plane, how many pairs can be at distance 1?

Theorem 13.16 (Erdős). *There are at most $cn^{3/2}$ pairs.*

Proof. Define a graph G on the point set, where two points x and y are connected by an edge if their distance is 1. It is easy to see that G does not contain $K_{2,3}$ as a subgraph. Indeed, the neighbours of each point lie on a circle of radius 1, and since two circles have at most two points in common, two of our vertices can have at most two common neighbours. Now an application of the Kővári-Sós-Turán theorem on the graph G gives the desired bound on the number of edges. □