GRAPH THEORY

1. Network flows

A network is an important variant of a graph in which it is

- directed (edges are given directions),
- weighted (edges assigned numeric values).

Note. Networks can be used to represent:

- computer/communication networks
- road/transportation systems
- water flow

Definition 1.1. A *directed graph*, or *digraph*, D consists of a graph $G = (V, E)$ in which each edge (or *arc*) has a direction. If $e = (u, v)$ is an arc, we call u the *initial* vertex and v the *terminal* vertex. Observe that edges (u, v) and (v, u) are the same if and only if $u = v$. The *outdegree* of a vertex v is the number of arcs with v as the initial vertex, which we denote $\operatorname{outdeg}(v)$. Similarly, the *indegree* of a vertex v is the number of arcs with v as the terminal vertex, which we denote $\operatorname{indeg}(v)$.

Example 1.2. Consider the following digraph D, where we have $\operatorname{indeg}(a) = 0$ and $\operatorname{outdeg}(a) = 2$.

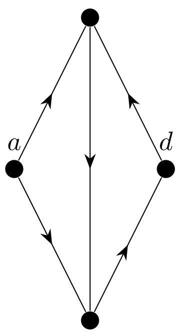

Definition 1.3. A *network* is a directed graph in which each edge e is assigned a nonnegative integer $c(e)$, called the *capacity* of e, which represents the maximum amount of 'material' that can be transported along edge e.

In a network flow problem, the network has two distinguished vertices, the *source* s and the *sink* t $(s \neq t)$ where $I(s) = 0$ and $O(t) = 0$. All other vertices are called *intermediate vertices*. A *flow* on a network N with source s and sink t is a function $f : V \times V \to \mathbb{N}$, which satisfies

1. for each edge e, $0 \leq f(e) \leq c(e)$,
2. for each intermediate vertex,

$$\sum_{e \in I(v)} f(e) = \sum_{e \in O(v)} f(e)$$

where $I(v)$ is the set of edges directed into v, and $O(v)$ is the set of edges directed out of v.

The *value* of a flow f is $\operatorname{val}(f) = \sum_{e \in I(t)} f(e) = \sum_{e \in O(s)} f(e)$.

Example 1.4. Consider the following network N.

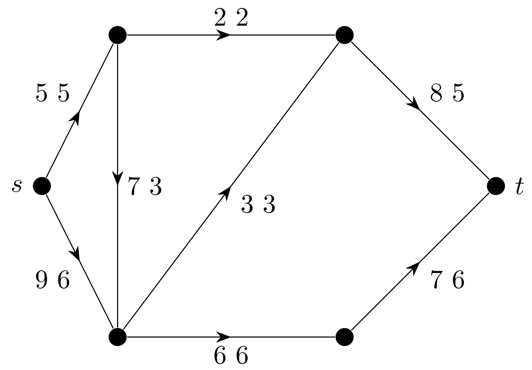

The *maximum flow problem* is to find a flow that maximises $\text{val}(f)$. A first bound on the maximum value is obtained by considering the capacity of the edges in $O(s)$ and $I(t)$. We have $\text{val}(f) \leq \sum_{e\in O(s)} c(e)$ and $\text{val}(f) \leq \sum_{e\in I(t)} c(e)$.

With our example, we can see that 11 is the maximum. It is clearly an upper bound, since the two vertices in $I(t)$ have together a maximum of $2+3+6=11$. Moreover, we provided an example with $\text{val}(f)=11$, demonstrating it is the true maximum.

1.1. **Flows and cuts.**

Definition 1.5. Let $N=(V,E)$ be a network and let $U\subset V$ be such that $s\in U$ and $t\notin U$. A *cut* in a flow network N is a pair $\big(U,\overline{U}\big)$, where $\overline{U}$ is the complement of U in V (i.e., $\overline{U}=V\setminus U$).

Definition 1.6. The *capacity of the cut* $\big(U,\overline{U}\big)$ is defined to be:

$$c\big(U,\overline{U}\big)=\sum_{e\in A(U,\overline{U})} c(e)$$

where $A\big(U,\overline{U}\big)$ denotes the set of edges directed from U to $\overline{U}$. Similarly, for a flow f on N, the *flow of the cut* $\big(U,\overline{U}\big)$ is

$$f\big(U,\overline{U}\big)=\sum_{e\in A(U,\overline{U})} f(e).$$

Example 1.7. Recall the network N.

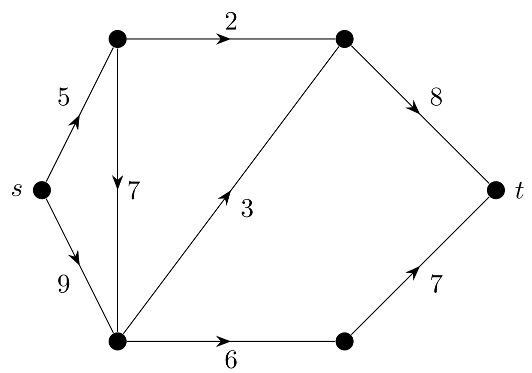

Then, for example, the cut $\big(U,\overline{U}\big)$ where $U=\{s\}$ has $c\big(U,\overline{U}\big)=14$.

Lemma 1.8. *Let N be a flow network and let f be a flow on N. If $\left(U,\overline{U}\right)$ is a cut of N, then*

$$\operatorname{val}(f) = f\left(U,\overline{U}\right) - f\left(\overline{U},U\right).$$

Proof. Suppose that $\left(U,\overline{U}\right)$ is a cut of N. Observe that

$$\begin{aligned} f(U,V) - f(V,U) &= \sum_{u\in U} f(\{u\},V) - \sum_{u\in U} f(V,\{u\}) \\ &= \sum_{u\in U} \left(f(\{u\},V) - f(V,\{U\})\right) \\ &= f(\{s\},V) - f(V,\{s\}) \\ &= \operatorname{val}(f), \end{aligned}$$

since for intermediate vertex, flow-in equals flow-out, and $f(V,\{s\}) = 0$ by definition (and $f(\{s\},V) = \operatorname{val}(f)$). But $f(U,V) - f(V,U) = f\left(U,\overline{U}\right) - f\left(\overline{U},U\right)$. □

Corollary 1.9. *Let N be a flow network and let f be a flow on N. If $\left(U,\overline{U}\right)$ is a cut of N, then*

$$\operatorname{val}(f) \leq c\left(U,\overline{U}\right).$$

Proof. Result follows by the above and $f\left(U,\overline{U}\right) \leq c\left(U,\overline{U}\right)$. □

1.2. **Flow incrementing paths.**

Definition 1.10. A *path* P is a graph $G = (V,E)$ is an alternating sequence of vertices and edges

$$v_1 e_1 v_2 e_2 \ldots e_{k-1} v_k$$

such that e_i is incident to v_i and v_{i+1} (for all i) and such that no vertex appears more than once.

Example 1.11.

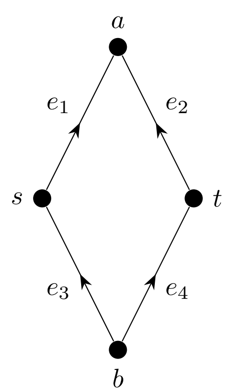

Let $P = se_1ae_2t$. Then we call e_1 a *forward edge* and e_2 a *reverse edge.*

Definition 1.12. Let N be a network with capacity c, f be a flow on N and P be a path from s to t. We say P is an *incrementing path* if:

(1) $f(e) < c(e)$ for all forward edges of P,
(2) $f(e) > 0$ for all reverse edges of P.

If P is an incrementing path, the *increment* of P is

$$i(P) = \min\left\{c(e) - f(e) \mid e \text{ is a forward edge of } P\right\} \cup \left\{f(e) \mid e \text{ is a reverse edge of } P\right\}.$$

Example 1.13.

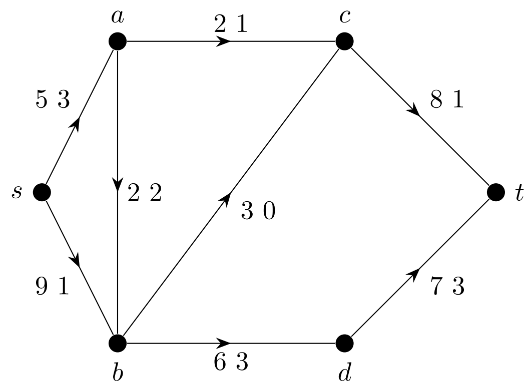

$P = sbact$ is an incrementing path, and $P' = sacbdt$ is not an incrementing path because $f(bc) = 0$. For $P = sbact$, we have that $\min\{9-1, 2, 2-1, 8-1\} = 1$ is the increment of P.

Note. Whenever we can find such a path, we can increment the flow in the overall network. Consider the example below doing this.

Example 1.14.

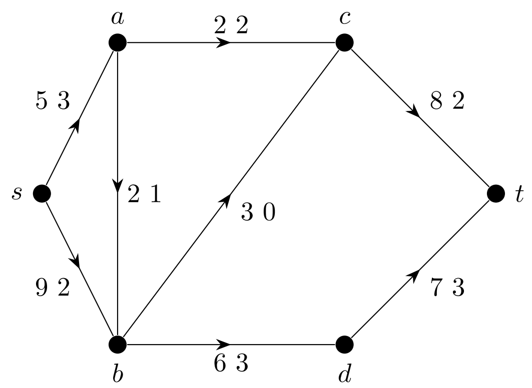

By adding 1 to each of the forward edges in P and subtracting 1 from each reverse edge, we increased the flow on N.

$$\mathrm{val}\,(f) + i(P) = 4 + 1 = 5.$$

Proposition 1.15. *Let N be a flow network and let f be a flow on N. Let P be an f-incrementing path of increment i. Then by adding i to the flow of each forward edge of P and subtracting i from the flow of each reverse edge of P, we obtain a new flow with value* $\mathrm{val}\,(f) + i$.

Corollary 1.16. *If there are no f-incrementing paths, then* $\mathrm{val}\,(f)$ *has maximum value.*

Theorem 1.17. (The max-flow min-cut theorem). *The maximum value of a flow in a network N is equal to the minimum capacity of a cut.*

Proof. Let f be a maximum flow. Since f has maximum value, N contains no f-incrementing paths from s to t. However, starting at s, we can construct a partial flow incrementing path - i.e., a path that satisfies the conditions of a flow-incrementing path, but does not end at t. Let S be the set of vertices of N that can be reached by a maximal flow incrementing path (if we have more than one path, we can take their union). Consider the cut $\big(S,\overline{S}\big)$; we will show that $\operatorname{val}(s) = c\big(S,\overline{S}\big)$.

Let $e = (u, v)$ be an edge with $u \in S$ and $v \in \overline{S}$. Then there is a partial flow-incrementing path from s to u. If $f(e) < c(e)$, then we could extend the partial flow-incrementing path to v, contradicting $v \in \overline{S}$. Hence, $f(e) = c(e)$.

Let $e = (v, u)$ be an edge with $u \in S$ and $v \in \overline{S}$. Then there is a partial flow-incrementing path from s to u. If $f(e) > 0$, then we could extend the partial flow-incrementing path to v, contradicting $v \in \overline{S}$. Hence, $f(e) = 0$.

By Lemma 1.8,

$$\operatorname{val}(f) = f\big(S,\overline{S}\big) - f\big(\overline{S},S\big).$$

But from the above, $f\big(S,\overline{S}\big) = c\big(S,\overline{S}\big)$ and $f\big(\overline{S},S\big) = 0$. Thus, $\operatorname{val}(f) = c\big(S,\overline{S}\big)$. Since $\operatorname{val}(f) \le c\big(U,\overline{U}\big)$ by Corollary 1.9, $\big(S,\overline{S}\big)$ is a minimum cut, and so the theorem follows. □

Based on the proof of the MFMC Theorem above, we get the *Ford-Fulkerson Algorithm* (1962).

OUTPUT: max flow on N

STEP 1: Begin with any flow on N.

STEP 2: Starting at s, construct a tree by adding edges that are candidates for flow-incrementing paths.

STEP 3: If the tree reaches t, then the unique path in the tree joining s to t is a flow incrementing path. Adjust the flow on the path by its increment. GOTO 2. On the other hand, if the tree does not reach t, then there is no flow incrementing-path, so the current flow is the maximum flow. END

2. Connectivity

Definition 2.1. Let $G = (V, E)$ be a connected graph. A *vertex cut* of G is a set $S \subseteq V$ such that $G - S$ has more than one component (i.e., $G - S$ is disconnected).

The *(vertex) connectivity* of G, denoted $\kappa(G)$, is the minimum size of a vertex set S such that $G - S$ is disconnected (or is a single vertex). Clearly, if G is disconnected then $\kappa(G) = 0$.

If $\kappa(G) \ge k$, we say G is *k-connected.*

G has connectivity one iff G is a graph with a *cut vertex* or is the complete graph on two vertices (i.e., K_2).

Note. One finds that

- $\kappa(K_n) = n - 1$,
- $\kappa(K_{m,n}) = \min\{m, n\}$,
- $\kappa(C_n) = 2$,
- and the connectivity of trees is one.

We also have $|V| \ge k + 1$ and $0 \le \kappa(G) \le n - 1$. Moreover, for any vertex v, we have

$$\kappa(G) - 1 \le \kappa(G - v).$$

2.1. Menger's Theorem.

In this section, we look at another measure of connectivity. The thoerem (by Menger) is one of the cornerstones of graph theory.

Definition 2.2. Given a graph $G = (V, E)$, a set $S \subseteq V$ of vertices is said to *separate* two vertices u and v iff

(1) $G - S$ is disconnected, and
(2) u, v belong to different components of $G - S$,

and we call S a *u-v separating set.* Evidently, S separates u and v iff $u, v \notin S$ and there is no path from u to v in $G - S$.

Example 2.3. Consider the graph G below.

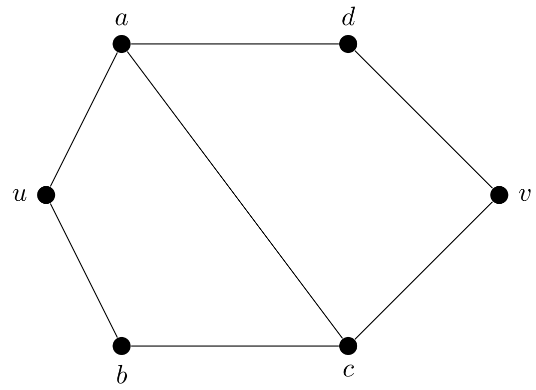

Then $S = \{a, b\}$ is a u-v separating set, where $G - S$ is shown below.

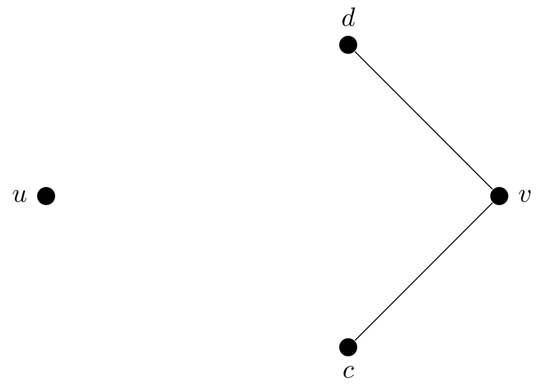

Likewise, $S' = \{a, c\}$ is a u-v separating set (where $G - S'$ is shown below).

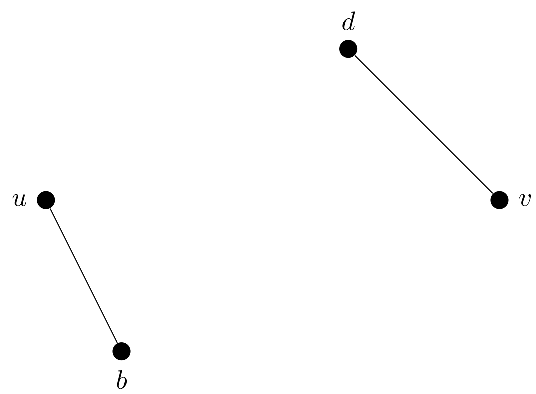

As an aside, $\kappa(G) = 2$ and consequently G is $0, 1, 2$-connected.

Note. If S separates u and v, then

(1) u and v are non-adjacent, and
(2) S is a vertex cut of G.

It follows from (2) that $|S| \geq \kappa(G)$.

Definition 2.4. A collection $\{P_1, P_2, \ldots, P_k\}$ of u-v paths is *internally disjoint* iff every pair of paths only have u and v in common.

Example 2.5. Consider the graph G from Example 2.3. Then the u-v 'paths'

- $P_1 = uadv$,
- $P_2 = ubcv$,
- $P_3 = uacv$,
- $P_4 = ubcadv$,

have P_1 and P_2 as internally disjoint.

Theorem 2.6. *Let u and v be two distinct vertices of a graph $G = (V, E)$. A set $S \subseteq V$ separates u from v iff every u-v path has at least one vertex belonging to S.*

Proof. ($\Longrightarrow$) Suppose there exists a u-v path P with no vertices belonging to S. Then u and v belong to the same component of $G - S$, and hence S does not separate u from v.

($\Longleftarrow$) Suppose S does not separate u from v (we assume $u, v \in S$, for otherwise trivially follows). Then there exists a path from u to v in $G - S$. Thus, there exists a u-v path with no vertices belonging to S. □

Example 2.7. Consider the network N below.

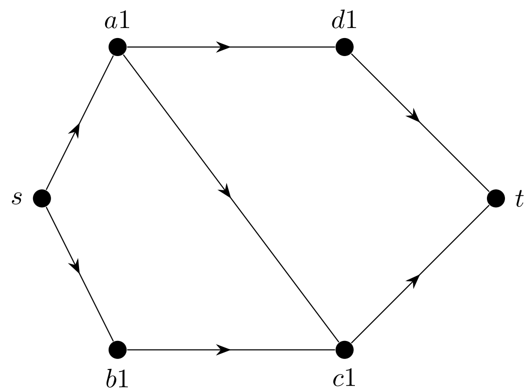

We give each of the vertices except s and t capacity 1. We can then construct N' (shown below) by converting each vertex x into x_- and x_+ allows us to model network N in terms of edge capacities.

The method:

(1) Replace each vertex $x \in V \setminus \{s, t\}$ by two vertices, say x_- and x_+.
(2) Send each incoming edge to x_- and each outgoing edge from x_+.
(3) For each x, add an edge from x_- to x_+ with capacity $c(x_-, x_+) = c(x)$.
(4) For a maximum flow, we get that it will be equal to the internally disjoint paths. Since u-v separating sets can be seen as cuts, it is clear where the following theorem comes from.

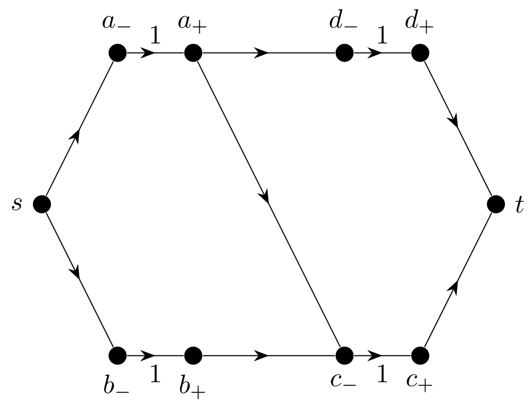

Theorem 2.8. (Menger's Theorem). *Let u and v be non-adjacent vertices in a graph G. Then the minimum number of vertices in a u-v separation set equals the maximum number of internally disjoint u-v paths.*

Proof. Let $P_1, \ldots, P_m$ be a set of m internally disjoint u-v paths, and assume m is the maximum size of such a set. Let S be a u-v separating set of minimum size and let $|S| = n$.

By Theorem 2.6, each of the paths $P_1, \ldots, P_m$ contains at least one vertex from S. Since the paths are internally disjoint, no member of S can occur in more than one path. It follows that we have at least as many vertices in S as we have internally disjoint paths. Hence, $m \leq n$.

Now, we show that $m \geq n$. To do so, we convert $G = (V, E)$ into a flow network N:

- Make u the source and v the sink. Direct all incident edges away from u and towards v.
- Replace each edge xy ($x, y \in V \setminus \{u, v\}$) by two directed edges xy and yx.
- Assign each vertex (other than u and v) a capacity of 1.

We can convert N to N', an equivalent network in which only edges have capacities. Each internal (x_- to x_+) edge has capacity 1 and each external has infinite capacity.

Consider a max flow f on N' (which we can find using the Ford-Fulkerson Algorithm). We will show that $m = \text{val}(f)$. We start by noting that every vertex in N' (except u, v) is in the form of x_- or x_+, and is the start or end of exactly one internal edge (which has capacity 1). It follows that the flow into or out of any vertex in N' is either 0 or 1. Hence, f assigns a value of 0 or 1 to any edge in N'.

This means that f must use $\text{val}(f)$ internally disjoint u-v paths in N'. Each path contributes 1 to the flow value. Since each of these internally disjoint u-v paths correspond to internally disjoint u-v paths in G, we get that $\text{val}(f) \leq m$. Similarly, since there are m internally disjoint u-v paths which can carry a flow of 1 in N', $m \leq \text{val}(f)$. Hence, $m = \text{val}(f)$.

By Theorem 1.17 (max-flow min-cut theorem), there is a minimum cut $\big(U, \overline{U}\big)$ in N' such that $c\big(U, \overline{U}\big) = \text{val}(f)$. We will show that $c\big(U, \overline{U}\big) \geq n$.

Since $c\big(U, \overline{U}\big) = \text{val}(f)$ and $\text{val}(f)$ is finite, each edge from U to $\overline{U}$ must be an internal edge. Moreover, since $\big(U, \overline{U}\big)$ is a cut, each directed path from u to v must use at least one of these edges. Hence, deleting either all of the 'heads' (x_+) or all of the 'tails' (x_-) of these edges removes all possible u-v directed paths in N'. Therefore, the set of vertices in G corresponding to these 'heads' or 'tails' must be a u-v separating set, of size $\text{val}(f)$. Thus, $m = \text{val}(f) \geq n$ and so it follows $m = n$, as desired. □

Theorem 2.9. (Whitney). *A simple graph G is k-connected if and only if for each pair of distinct vertices, there are at least k internally disjoint paths.*

Note. $\kappa(G) = \max k$ for which every pair of distinct non-adjacent vertices in G are connected by k internally disjoint paths.

Example 2.10. Consider the graph $K_{2,3}$ (shown below).

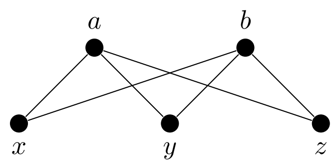

Consider a and b:

- there are three internally disjoint a-b paths (axb, ayb and azb);
- $S = \{x, y, z\}$ is an a-b separating set.

Consider x, y, z:

- there are two internally disjoint paths between each pair (for example xay and xby for pair x and y);
- $S = \{a, b\}$ is a x-y separating set (in fact, S separates every pair between x, y and z).

The connectivity of $K_{2,3}$ is 2.

2.2. **Edge connectivity.**

Definition 2.11. An *edge cut* of $G = (V, E)$ is an edge set of the form $S = \left(U, \overline{U}\right)$ where $U \subset V$ such that U is non-empty and $G - S$ has more than one component.

Definition 2.12. The *edge connectivity* of G, denoted $\lambda(G)$, is the minimum size of an edge cut. If $\lambda(G) \geq k$, then G is *k-edge-connected*. We will denote $\delta(G)$ as the minimum degree and $\Delta(G)$ as the maximum degree.

Note. $\lambda(G) = 0$ iff G is disconnected and $\lambda(G) = 1$ iff G has a bridge.

Comparing edge and vertex connectivity (and minimum degree):

- $\kappa(K_n) = n - 1 = \lambda(K_n) = \delta(K_n)$
- $\kappa(C_n) = 2 = \lambda(C_n) = \delta(C_n)$
- $\kappa(P_n) = 1 = \lambda(P_n) = \delta(P_n)$
- $\kappa(T_n) = 1 = \lambda(T_n) = \delta(T_n)$
- $\kappa(K_{m,n}) = \min(m, n) = \lambda(K_{m,n}) = \delta(K_{m,n})$

Example 2.13. We provide an example showing that vertex connectivity and edge connectivity are not the same. Consider the graph G below.

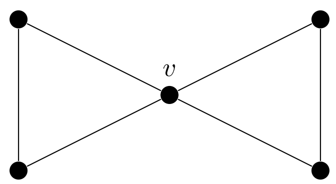

We have that v is a vertex cut such that $\kappa(G) = 1$. Yet, $\lambda(G) = 2$.

Now, consider the graph G' (shown below).

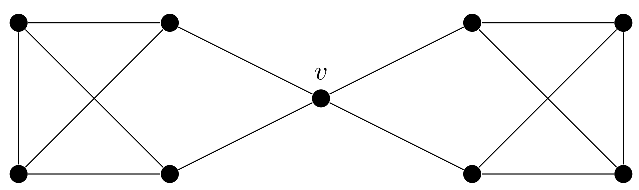

Once again, vertex connectivity and edge connectivity are not the same. Yet, we also get that edge connectivity does not equal minimal degree:

$$\kappa(G') = 1 < \lambda(G') = 2 < \delta(G') = 3.$$

Theorem 2.14. (Whitney, 1932). *For every graph G, with minimal degree $\delta(G)$, we have*

$$\kappa(G) \leq \lambda(G) \leq \delta(G).$$

Proof. The edges incident to a vertex v of minimal degree form an edge cut. Hence, $\lambda(G) \leq \delta(G)$. It remains to show $\kappa(G) \leq \lambda(G)$.

Let $\left(U,\overline{U}\right)$ be a minimum edge cut of G, and recall that $\kappa(G) \leq n-1$. If every vertex in U is adjacent to every vertex in $\overline{U}$, then

$$\left|\left(U,\overline{U}\right)\right| = |U|\left|\overline{U}\right| \geq n-1 \geq \kappa(G)$$

and hence $\lambda(G) \geq \kappa(G)$. Otherwise, we choose $x \in U$ and $y \in \overline{U}$ with x non-adajcent to y.

Let T be the set of all vertices adjacent to x in $\overline{U}$ and all vertices in $U - x$ adjacent to vertices in U. Every x-y path passes through T, so T is a separating set. It follows $|T| \geq \kappa(G)$.

If we choose the edges from x to $T \cap \overline{U}$ and one edge from each vertex of $T \cap U$ to $\overline{U}$, we have $|T|$ edges of $\left(U,\overline{U}\right)$. Hence, $|T| \leq \left|\left(U,\overline{U}\right)\right|$ and thus

$$\kappa(G) \leq |T| \leq \left|\left(U,\overline{U}\right)\right| = \lambda(G),$$

as desired. □

Theorem 2.15. (Menger's Theorem for edges). *Let u and v be distinct vertices in a graph G. The* minimum *number of edges separating u and v equals the* maximum *number of edge-disjoint u-v paths.*

Theorem 2.16. *A simple graph G is k-edge-connected iff for each pair of distinct vertices u and v in G, there are at least k-edge-disjoint u-v paths in G.*

3. Ramsey Theory

Ramsey theorey is a branch of mathematics named after British mathematician Frank Ramsey.

3.1. The Party Problem.

Six people are at a party. We will show that there are three people who all know each other, or three people who are all strangers.

For example, we can consider the graph below. We can consider the three vertices coloured blue as people who are all strangers.

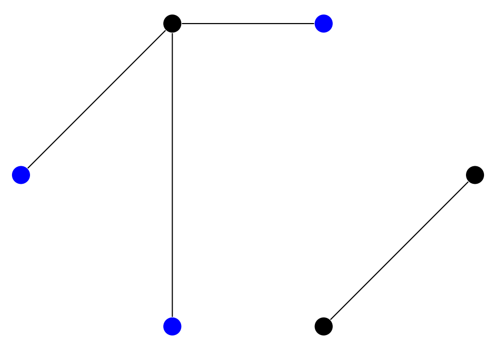

In a simple graph on 6 vertices, there are 3 vertices which are mutually adjacent or 3 vertices which are mutually non-adjacent.

Example 3.1. Consider the graph G_1, which those coloured blue are three mutually adjacent vertices:

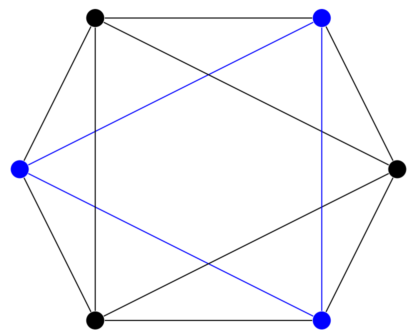

Consider the graph G_1, which those coloured blue are three pairwise non-adjacent vertices:

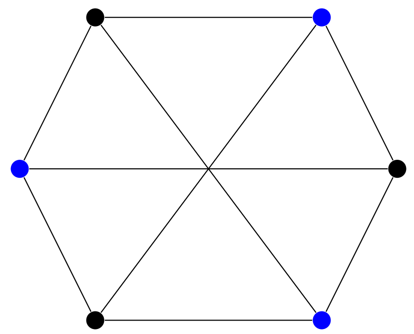

Definition 3.2. A *clique* in a graph G is a set of vertices which are mutually adjacent (they form a complete subgraph of G).

Definition 3.3. A *stable set* in a graph G is a set of vertices which are pairwise non-adjacent (i.e., there are no edges between them).

Note. Cliques and stable sets are complementary notions. That is to say, a set of vertices is a clique in G iff it is a stable set in $\overline{G}$. Note that a single vertex is both a clique and stable set.

Theorem 3.4. *Every simple graph G on 6 vertices contains either a clique of size 3 or a stable set of size 3. Equivalently, if G is a simple graph on 6 vertices, then G or $\overline{G}$ contains a clique of size 3.*

Note. There are 156 non-isomorphic graphs on 6 vertices, so this could be proved by an exhaustive search.

Proof. Let G be a simple graph on 6 vertices, and let v be a vertex of G. If $\deg(v) < 3$, then we consider the following argument for $\overline{G}$. Therefore, without loss of generality, assume $\deg(G) \geq 3$. Consider the graph G_1, which those coloured blue are three mutually adjacent vertices:

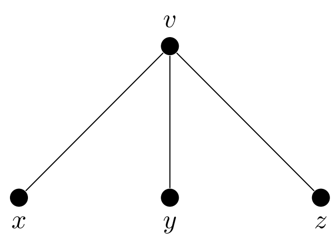

Let x, y, z be three neighbourhoods of v. If any pair of x, y, z are adjacent, then we have a clique of size 3. Otherwise, if x, y, z are pairwise non-adjacent, then we have a stable set of size 3. □

Theorem 3.5. (Ramsey's Theorem). *For positive integers s and t and for sufficiently large n, every simple graph on n vertices contains a clique of size s or a stable set of size t.*

A natural question follows: what is the smallest value of n for which Theorem 3.5 holds?

Definition 3.6. The *Ramsey number*, denoted $R(s,t)$, is the minimum positive integer n such that a graph on n vertices contains a clique of size s or a stable set of size t. When $s = t$, we have $R(s,s) = R(t,t)$ (which we may denote $R(s)$), and these are called the *diagonal Ramsey numbers*.

Example 3.7. Theorem 3.4 tells us that $R(3,3) \leq 6$. Consider the cycle C_5, shown below.

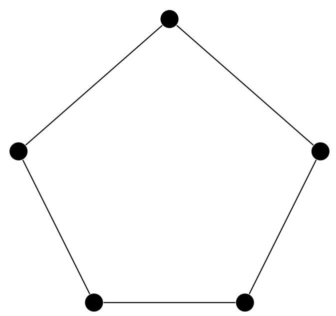

Since C_5 contains neither a clique of size 3 or a stable set of size 3, it follows that $R(3,3) = 6$.

There are several different ways to define Ramsey numbers. $R(s,t)$ is the minimum positive integer n such that:

(1) If G is a simple graph on n vertices, then G or $\overline{G}$ contains a clique of size t or size s.
(2) For every positive colouring of the edges of K_n with colours red and blue, there exists either red clique of size s or a blue clique of size t.

Proposition 3.8. *Ramsey numbers are symmetric (i.e., $R(s,t) = R(t,s)$ for all $s, t \geq 1$).*

Proof. This follows from the definition of $R(s,t)$ using edge coloured K_n, since the colours are interchangeable. □

Proposition 3.9. $R(s,1) = 1$ *for all* $s \geq 1$.

Proof. Let G be a simple graph on $n \geq 1$ vertices. It suffices to choose any arbitrary vertex of G, since it will form a stable set of size 1. □

Proposition 3.10. $R(s,2) = s$ *for all* $s \geq 1$.

Proof. The statement is trivially true for K_1, so assume $s \geq 2$. Let G be a simple graph with $n = s$ vertices. The statement is true for $G = K_s$ (since it contains a clique of size s). So, suppose that G is not complete. It follows that there are at least two non-adjacent vertices which form a stable set of size two. Hence, $R(s,2) \leq s$. Consider K_{s-1}, which neither has a clique of size $s-1$ nor a stable set of size 2. Thus, $R(s,2) = s$. □

Theorem 3.11. (Ramsey's Theorem). *For positive integers s and t, there exists a (minimal) positive integer $R(s,t)$ such that if we colour the edges of $K_{R(s,t)}$ red and blue, this graph will either have a K_s subgraph with only red edges, or a K_t subgraph with only blue edges.*

Proof. We use induction on s and t (double induction). The base case is clear, since $R(s,1) = 1$, $R(s,2) = s$, $R(1,t) = 1$, and $R(2,t) = t$ (so statement holds for $s + t \leq 3$).

Now, we assume that the statement is true for $s + t - 1 \geq 2$, i.e., $R(s-1,t)$ and $R(s,t-1)$ exists. We prove the induction step by showing that $R(s,t) \leq R(s,t-1) + R(s-1,t)$, for $s,t \geq 2$.

Consider a 2-edge coloured complete graph G on $R(s,t-1)+R(s-1,t)$ vertices. Each edge is coloured either red or blue.

Let v be a vertex of G. Since G is complete, v has $R(s,t-1) + R(s-1,t) - 1$ neighbours. Let N be the number of red edges incident to v and M be the number of blue edges incident to v. Then $0 \leq N, M \leq R(s,t-1) + R(s-1,t) - 1$. it follows that

$$M + N + 1 = R(s,t-1) + R(s-1,t).$$

This means that either $R(s,t-1) \leq M$ or $R(s-1,t) \leq N$.

If $M \geq R(s,t-1)$, then the induced subgraph on these M vertices adjacent to v in G has either a red K_s or a blue K_{t-1}. If the former is true, then we are done. If it is the latter, then we can form a blue K_t by adding vertex v to K_{t-1}, along with the M blue edges, and we are done.

On the other hand, if $N \geq R(s-1,t)$, then the induced subgraph on those N vertices adjacent to v contains either a red K_{s-1} or a blue K_t. In any case, by similar argument by the above, we are done.

Hence,

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

and the result follows. □

Corollary 3.12. *For all $s,t \geq 2$,*

$$R(s,t) \leq R(s-1,t) + R(s,t-1).$$

Moreover, if $R(s-1,t)$ and $R(s,t-1)$ are both even, then the equality is strict.

Proposition 3.13. $R(3,4) = 9$.

Proof. We found that $R(3,3) = 6$ and $R(2,4) = 4$. Hence, the Corollary above,

$$R(3,4) < R(2,4) + R(3,3) = 10$$

and so $R(3,4) \leq 9$. Hence, it remains to show that $R(3,4) > 8$. Consider the graph G on 8 vertices, which has neither a K_3 subgraph nor stable set of size 4.

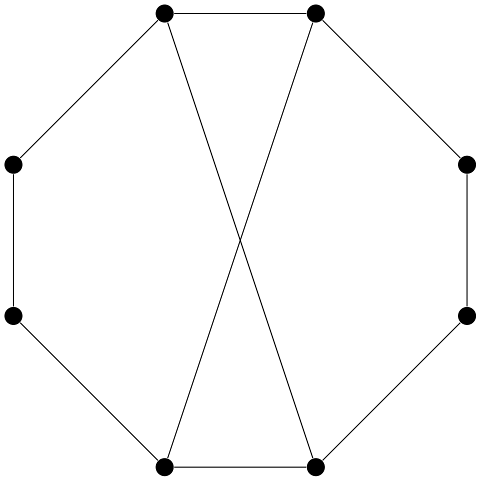

□

Proposition 3.14. $R(3,5) = 14$.

Proposition 3.15. $R(4,4) = 18$.

Proposition 3.16. $R(4,5) = 25$ *(proved in* 1995*)*.

Conjecture 3.17. $R(5,5) = 43$.

Proposition 3.18. $43 \leq R(5,5) \leq 48$ *(and* $R(5,5) \leq 46$ *soon)*.

Theorem 3.19. *For positive integers $s, t \geq 2$, we have*

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

Proof. We provide a proof by induction on s and t. We have that $R(s,2) = s = \binom{s}{1}$ and $R(2,t) = t = \binom{t}{1}$. Hence, the statement holds for $s + t \leq 4$.

Now, assume the statement holds for some positive integers $s, t \geq 2$ such that $s + t \geq 5$. That is to say, we have bounds for $R(s-1,t)$ and $R(s,t-1)$. Hence,

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

where the final equality follows from Pascal's identity, and so the result follows. □

Corollary 3.20. *For all positive integers s and t, we have*

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

with equality if $s = t = 1$.

Proof. Trivially $R(1,1) = 1 = 2^0$, so assume $s, t \geq 2$. Observe that $\binom{s+t-2}{s-1}$ is the number of $(s-1)$-subsets of a $(s+t-2)$-set, and 2^{s+t-2} is the total number of subsets of this set. □

Note. It follows that $R(s,s) \le 2^{2s-2} = 4^{s-1}$.

Theorem 3.21. *For any positive integer s, $R(s,s) \ge 2^{\frac{s}{2}}$.*

Proof. We know that $R(1,1) = 1$ and $R(2,2) = 2$, so we may assume that $s \ge 3$.

Let $\mathscr{G}_n$ denote the collection of simple graphs with vertex set $\{v_1, v_2, \dots, v_n\}$. Let $\mathscr{G}_n^s$ be a subcollection of $\mathscr{G}_n$, which contains these labelled graphs which have a clique of s vertices.

Observe that $|\mathscr{G}_n| = 2^{\binom{n}{2}}$ since each subset of the $\binom{n}{2}$ possible edges determines a graph in $\mathscr{G}_n$. Similarly, the number of graphs in $\mathscr{G}_n^s$ which have a *particular* set of s vertices has a clique is $2^{\binom{n}{2}-\binom{s}{2}}$. Because there are $\binom{n}{s}$ places in a graph in which a K_s subgraph can appear, we have

$$|G_n^s| \le \binom{n}{s} 2^{\binom{n}{2}-\binom{s}{2}}.$$

This inequality arises because the graphs in G_n^s which have more than one K_s subgraph are counted more than once. Hence,

$$\begin{aligned}\frac{|G_n^s|}{|G_n|} &\le \frac{\binom{n}{s} 2^{\binom{n}{2}-\binom{s}{2}}}{2^{\binom{n}{2}}} \\ &= \binom{n}{s} 2^{-\binom{s}{2}} \\ &< \frac{n^s 2^{-\binom{s}{2}}}{s!}.\end{aligned}$$

Now, suppose that $n < 2^{\frac{s}{2}}$. Then

$$\begin{aligned}\frac{|G_n^s|}{|G_n|} &< \frac{\left(2^{\frac{s}{2}}\right)^s 2^{-\binom{s}{2}}}{s!} \\ &= \frac{2^{\frac{s^2}{2}-\binom{s}{2}}}{s!} \\ &= \frac{2^{\frac{s}{2}}}{s!} \\ &< \frac{1}{2}.\end{aligned}$$

In other words, if $n < 2^{\frac{s}{2}}$, then fewer than half of the graphs in G_n contain a clique of size s. Likewise, fewer than half of the graphs in G_n contain a stable set of size s. Hence, there exists a graph in G_n that contains neither a clique of size s nor a stable set of size s. Thus, because this holds for $n < 2^{\frac{s}{2}}$, it follows that $R(s,s) \ge 2^{\frac{s}{2}}$. □

Note. Hence,

$$2^{\frac{s}{2}} \le R(s,s) \le 4^{s-1}.$$

These are the 'best' bounds on the Ramsey numbers.

4. Extremal Graph Theory

We now ask the following: given a graph of order n, what is the maximum size (number of edges) it can have without having a given property (for example, Eulerian, Hamiltonian, or clique of certain size).

Recall the result that every simple graph on 6 vertices contains either a clique of size 3 or a stable set of size 3. We now ask what is the maximum number of edges of a graph of order 6 we can have without containing a triangle (K_3).

Example 4.1. Every simple graph of order 6 and size 10 contains a triangle. One can consider $K_{3,3}$ and the graph below, both of which are simple graphs of order 6 and size 9 with no triangles.

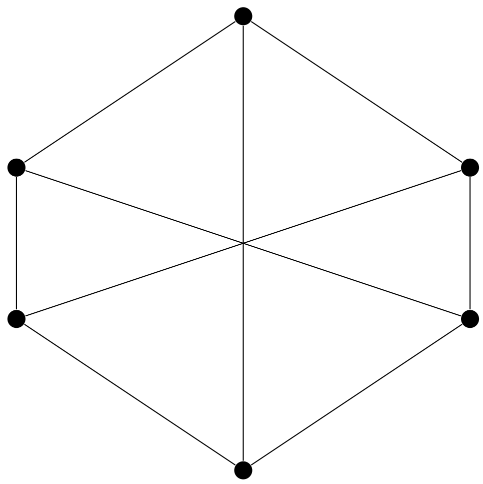

Theorem 4.2. (Mantel's Theorem, 1907). *If G is a graph of order $n \geq 3$ and size $m > \left\lfloor \frac{n^2}{4} \right\rfloor$, then G contains a triangle.*

Proof. To derive a contradiction, suppose the statement is false. Then there exists a smallest integer n for which the statement is false, so there is some graph $G = (V, E)$ of order n and size $\left\lfloor \frac{n^2}{4} \right\rfloor + 1$ that contains no triangle. Let $uv \in E$ and let $v_1, v_2, \ldots, v_{n-2}$ be the remaining vertices of G. Since G contains no triangle, it is clear u and v cannot share any neighbours. That is to say, at most one of u or v is adjacent to v_i $(1 \leq i \leq n-2)$. Hence,

$$\deg(u) + \deg(v) \leq n.$$

Let $H = G - v - u$. Then H has order $n-2$ and size m, where

$$\begin{aligned} m &= \left\lfloor \frac{n^2}{4} \right\rfloor + 1 - \left(\overbrace{\deg(u) + \deg(v)}^{\leq n} - 1 \right) \\ &\geq \left\lfloor \frac{n^2}{4} \right\rfloor - n + 2 \\ &= \left\lfloor \frac{(n-2)^2}{4} \right\rfloor + 1. \end{aligned}$$

Since H has fewer vertices than G and by our assumption, it follows that H must contain a triangle. But H does not contain a triangle, a contradiction. □

Note. If n is even, say $n = 2k$, then every graph of size $m > k^2$ contains a triangle.

The bound cannot be improved because the graph $K_{k,k}$ has order $2k$, size k^2 and contains no triangles. This is the only of this order/size up to (graph) isomorphism.

Theorem 4.3. *Let G be a graph of order $n \geq 3$ and size m. If $m > \left\lfloor \frac{n^2}{4} \right\rfloor$ or $m = \left\lfloor \frac{n^2}{4} \right\rfloor$ and $G \neq K_{\lfloor \frac{n}{2} \rfloor, \lceil \frac{n}{2} \rceil}$, then G contains a triangle.*

We can generalise this and consider the problem for arbitrary subgraphs.

Definition 4.4. Let F be a simple graph, we denote by $ex(n, F)$ the maximum number of edges in a graph G on n vertices which does not contain F.

Such a graph is called an *extremal graph* (for this property) and the collection of graphs is denoted $Ex(n, F)$.

Example 4.5. We have $ex(n, K_3) = \left\lfloor \frac{n^2}{4} \right\rfloor$ and $Ex(n, K_3) = \left\{ K_{\lfloor \frac{n}{2} \rfloor, \lceil \frac{n}{2} \rceil} \right\}$ (i.e., there is only one extremal graph for K_3). For instance, $ex(6, K_3) = 9$ and $Ex(6, K_3) = \{K_{3,3}\}$.

Note. This example tells us that the bound in Mantel's theorem cannot be improved because $K_{\lfloor \frac{n}{2} \rfloor, \lceil \frac{n}{2} \rceil}$ contains no triangles and it has $\lfloor \frac{n}{2} \rfloor \lceil \frac{n}{2} \rceil \leq \frac{n^2}{4}$ edges (equality when n is even, so bound is tight).

The graph $K_{\lfloor \frac{n}{2} \rfloor, \lceil \frac{n}{2} \rceil}$ is an example of a *Turán graph*.

Definition 4.6. A *Turán graph*, denoted $T_{n,r}$ or $T(n, r)$, is a complete r-partite graph on n vertices ($n \geq r$) whose partition sets differ in size by at most 1.

Note. In general: $K_{\lfloor \frac{n}{r} \rfloor, \lfloor \frac{n}{r} \rfloor, \ldots, \lceil \frac{n}{r} \rceil, \lceil \frac{n}{r} \rceil}$.

Example 4.7. The Turán graph $T(8, 3)$ is shown below (which is equivalently $K_{2,3,3}$ and does not contain K_4).

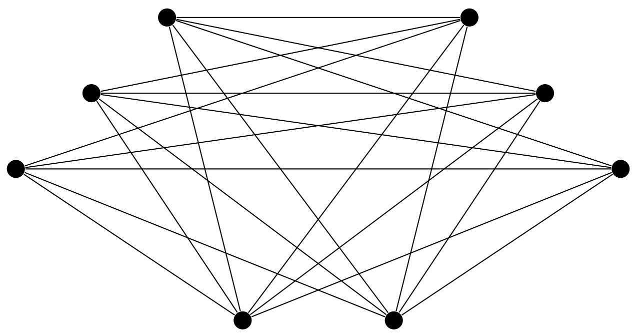

Intuitively, if we want to avoid K_{r+1}, we construct $T(n, r)$.

Theorem 4.8. *Among all r-partite graphs of order n, the Turán graph $T(n, r)$ is the unique graph of maximum size.*

Note. Turán graphs are *dense* - they have n^2 edges. More specifically, let $t(n, r)$ denote the number of edges in $T(n, r)$. Then we have

$$t(n, r) \leq \frac{n^2}{2}\left(1 - \frac{1}{r}\right),$$

with equality whenever r divides n.

We now meet Turán's Theorem, which generalises Mantel's Theorem from triangles to cliques of any given size. Turán asked the question: Suppose G is a simple graph that does not contain an $(r + 1)$-clique. What is the maximum number of edges that G can have?

Theorem 4.9. *If a graph $G = (V, E)$ on n vertices has no $(r + 1)$-clique ($r \geq 1$), then*

$$|E| \leq \frac{n^2}{2}\left(1 - \frac{1}{r}\right).$$

Equivalently: if a graph $G = (V, E)$ on n vertices has no $(r + 1)$-clique ($r \geq 1$) and $ex(n, K_{r+1})$ edges, then $G = T(n, r)$.

Note. For $r = 1$, this is trivial and when $r = 2$, we have $ex(n, K_3) = \left\lfloor \frac{n^2}{4} \right\rfloor$ by Mantel's theorem (giving us that Mantel's theorem is a special case of this one).

Proof. We provide a proof by (strong) induction on n. When $n \geq r$, the statement trivially holds. We can achieve equality when $n = r$ and $G = K_r$.

Assume $n > r$ and that the statement holds for all graphs with r to $n-1$ vertices. Let G be a graph on vertex set $V = \{v_1, \ldots, v_n\}$ without an $(r+1)$-clique with a maximum number of edges, where $n > r$. We know that G must contain at least one K_r clique (otherwise it would not have maximum size).

Let A be one of the K_r subgraphs in G and let B be the induced subgraph on the vertices $V \setminus A$.

Subgraph A contains $e_A = \binom{r}{2}$ edges. We now estimate the number of edges in B, e_B, and the number of edges between A and B, e_{AB}.

By our induction hypothesis, we have $e_B \leq \frac{(n-r)^2}{2}\left(1 - \frac{1}{r}\right)$. Since G has no K_{r+1} subgraph, every vertex v_i of B $(1 \leq i \leq n-r)$ is adjacent to at most $r-1$ vertices of A, giving us

$$e_{AB} \leq (r-1)(n-r).$$

Thus, we have

$$\begin{aligned} |E| &= e_A + e_{AB} + e_B \\ &\leq \binom{r}{2} + (r-1)(n-r) + \frac{(n-r)^2}{2}\left(1 - \frac{1}{r}\right) \\ &= \frac{n^2}{2}\left(1 - \frac{1}{r}\right). \end{aligned}$$

□

5. Random Graphs

A *random graph* is a graph in which the number of vertices is specified and their adjacencies are determined at random.

The theory of random graphs uses probabilistic methods to establish the existence of certain types of graphs, and to determine some properties of 'almost all' graphs.

Erdős-Renyi (the simplest model) $1940 - 50s$: There are two models, both originate from a simple model introduced in a 1947 paper by Erdős.

Let $N = \binom{n}{2}$.

(1) The uniform model $\mathscr{G}(n, m)$. For $0 \leq m \leq N$, the probability space $\mathscr{G}(n, m)$ consists of all $\binom{N}{m}$ subgraphs of K_n with m edges and uniform probability function
$$\mathcal{P}(G) = \frac{1}{\binom{N}{m}}$$
for all $G \in \mathscr{G}(n, m)$.

(2) The Binomial Model $\mathscr{G}(n, p)$. The probability space $\mathscr{G}(n, p)$ is obtained by fixing a probability p between 0 and 1, and choosing each edge with probability p (these choices are independent). Since $(1-p)$ is the probability an edge is not chosen, we have probability function
$$\mathcal{P}(G) = p^m (1-p)^{N-m}$$
for each $G \in \mathscr{G}(n, p)$.

Note. These two models are equivalent when n is large and $p = \frac{m}{\binom{n}{2}}$. In practice, the binomial theorem model is the most-commonly used because it is easier/neater to work with.

Example 5.1. The probability space $\mathscr{G}(3, p)$:

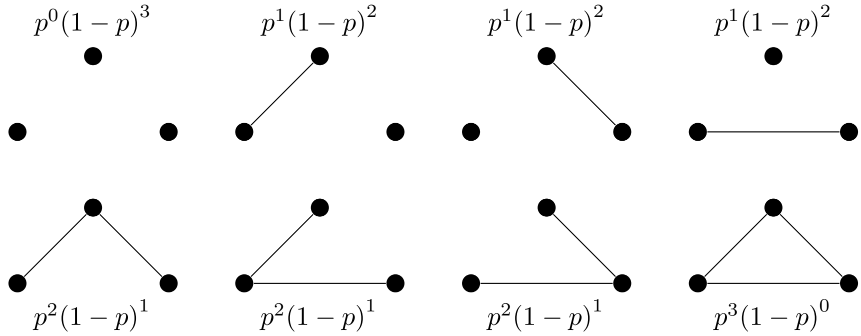

One can find the probability that G has $0 \leq m \leq 3$ edges as $\binom{3}{m} p^m (1-p)^{3-m}$. Using this, one could also find the probability that G is connected (which is $p^2(3-2p)$).

Some questions: what happens as we change p? What happens when p is a function of n?

We are interested in the properties of our probability spaces as $n \to \infty$.

We will say that a typical element of our space has property Q when the probability that a random graph $G \in \mathscr{G}(n, p)$ on n vertices has Q tends to 1 as $n \to \infty$. We say *almost all* or *almost every* (a.e.) graph has property Q or property Q *almost always* holds.

Consider $\mathscr{G}(n, p)$, we want to know about $\lim_{n \to \infty} \mathscr{G}(n, p)$.

Some questions:

- Is it connected?
- Does it contain a triangle? (think of Ramsey theory)
- Is it Hamiltonian?
- Is it Eulerian?
- What are the expected number of edges?

These questions are most interesting when $p = \frac{c}{n}$, where c is constant.

Lemma 5.2. *When p is constant, for every graph H, almost every $G \in \mathscr{G}(n, p)$ contains an induced copy of H.*

Proof. Let H have k vertices. If $n \geq k$ and $U \subseteq \{0, \ldots, n-1\}$ ($|U| = k$) is a fixed set of vertices of G, then $G[U]$ is isomorphic to H with a certain probability $r > 0$.

The probability r depends on p but not n. G contains a collection of $\left\lfloor \frac{n}{k} \right\rfloor$ disjoint such sets U. The probability that none of the corresponding graphs $G[u]$ are isomorphic to H is: $(1-r)^{\left\lfloor \frac{n}{k} \right\rfloor}$ since these events are independent (as the edge sets of the vertices in U are disjoint). Thus,

$$\mathcal{P}(H \ncong G[u]) \leq (1-r)^{\left\lfloor \frac{n}{k} \right\rfloor},$$

since there are more choices than this. Because $(1-r)^{\left\lfloor \frac{n}{k} \right\rfloor} \to 0$ as $n \to \infty$, the result holds. □

Some other properties of almost all graphs are below.

Theorem 5.3. *In almost all graphs, every pair of vertices are connected by a path of length* 2.

Theorem 5.4. *Almost every graph is connected.*

Theorem 5.5. *For $k \geq 0$, almost every graph is k-connected.*

5.1. **Threshold Functions.** The most intriguing discover made by Erdős and Renyi in the course at investigating random graphs in the phenomenon of *thresholds.*

For many graph properties, the limiting probability that a random graph possesses them jumps from 0 to 1 (or vice versa) very rapidly (i.e., with a small increase in the number of edges).

The number of triangles in $G \in \mathscr{G}(n,p)$ changes abruptly at the threshold $p = \frac{1}{n}$.

- If $p << \frac{1}{n}$, then almost every $G \in \mathscr{G}(n,p)$ has no triangle.
- If $p >> \frac{1}{n}$, then almost every $G \in \mathscr{G}(n,p)$ has at least one triangle.

The function $f(n) = \frac{1}{n}$ is called a *threshold function* for the property of G containing a triangle.

More generally, a threshold function for a graph property Q is a function $f : \mathbb{N} \to [0,1]$, such that

- If $p << f(n)$, then almost every $G \in \mathscr{G}(n,p)$ does not have Q.
- If $p >> f(n)$, then almost every $G \in \mathscr{G}(n,p)$ have Q.

Bollobes and Thomason (1987) showed that (trivial exceptions aside), all *monotone* graph properties have threshold functions. *Monotone* properties are properties which are preserved when edges are added.

5.2. **The Evolution of Random Graphs.** In a random graph $G \in \mathscr{G}(n,p)$ (where n is arbitrarily large), when we increase p from 0, it is possible to determine at what stage in the 'evolution' a particular graph property Q is likely to arise.

The idea: we start with graph $G \in \mathscr{G}(n,p)$ with n vertices and no edges ($p = 0$). As p increases, the graph 'evolves' as it acquires more edges (and therefore more structure). Properties appear in 'jumps' as p crosses the threshold at that property.

A rough timeline:

- $p = 0$: G consists of isolated vertices.
- $p = \frac{1}{n^2}$: the expected number (average) of edges in G is 1.
- $p = \frac{d}{n^2}$ $(d > 1)$: the expected number of edges in G is almost surely d and the components are all of size 1 or 2 (K_2 or path length 2).
- $p = \frac{1}{n^{\frac{k}{k-1}}}$: the components of G are all trees of size at most k.
- $p = \frac{1}{n}$: cycles are born (i.e., triangles).
- $p = \frac{d}{n}$ $(d < 1)$: there are a constant number of components with cycles in them. Almost surely all components are trees or unicycle and are size at most $\log n$.
- 'Soon' after : G is non-planar.
- $p = \frac{d}{n}$ $(d \sim 1)$: the largest component has size $n^{\frac{2}{3}}$
- $p = \frac{d}{n}$ $(d > 1)$: the largest component has size cn, where $0 < c < 1$. This is called the 'double jump' and the largest component is called the 'giant component'.
- Soon after: G is almost surely Hamiltonian.
- Soon after: has diameter 2.

5.3. **The Regularity Theorem.** Paradoxically, the behaviour of random graphs is often highly predictable. For example, when p is constant, many properties of the random graph $G \in \mathscr{G}(n,p)$ hold almost surely.

One of the difficulties in saying things about 'concrete' graphs is that they are less homogenous: their edges can be scattered about the graph the graph in unpredictable ways, even if basic information (like connectivity) is known about the graph.

Fortunately, it turns out that any sufficiently large, dense graph can be split such that the resulting (approximately equal size) subgraphs are joined to one another in an essentially random manner.

This allows us to establish many properties of such graphs. This remarkable and surprising fact is known as the *Regularity Lemma* (Szemeredi, 1978).

First, some definitions.

Let $G = (V, E)$ be a graph and let $X, Y \subseteq V$ be disjoint. Let $e(X, Y)$ be the number of edges between X and Y. Then the *density* (X, Y) is

$$d(X,Y) = \frac{e(X,Y)}{|X||Y|}.$$

Let $\epsilon > 0$ be arbitrary. We refer to a subset S of X as *small* if $|S| \leq \epsilon|X|$, and *large* subset of X if $|S| \geq \epsilon|X|$.

Given some $\epsilon > 0$, we call pair (X, Y) ϵ-*regular* if

$$|d(X',Y') - d(X,Y)| \leq \epsilon$$

whenever X' and Y' are large subsets of X and Y, respectively. That is, $|X'| \geq \epsilon|X|$ and $|Y'| \geq \epsilon|Y|$. This tells us that the edges between ϵ-regular pairs are distributed fairly uniformly.

Consider a partition $\{V_0, V_1, \ldots, V_k\}$ of V in which V_0 is singled out as an 'exceptional set' (may be empty). We call such a partition ϵ-*regular* if:

(1) $|V_0| \leq \epsilon|V|$ (in other words, V_0 is small in V)
(2) $|V_1| = \ldots = |V_k|$
(3) all but at most ϵk^2 of the pairs (V_i, V_j) $1 \leq i < j \leq k$ are ϵ-regular.

Theorem 5.6. (Regularity Lemma). *For every $\epsilon > 0$ and every integer $m \geq 1$, there exists an integer M such that every graph of at least order m has ϵ-regular partitions $\{V_0, V_1, \ldots, V_k\}$ with $m \leq k \leq M$.*

What the lemma tells us:

- Every sufficiently large graph has an ϵ-regular partition into not too many parts.
- We can specify a lower bound (m) on the number of parts (ϵ-regularity is trivial when the parts are singletons and powerful when they are large).
- We can specify an upper bound (M) on the number of parts. This ensures that when G is large, the parts are also large.

We now apply the lemma.

Recall the best bounds we have on $R(s, s)$ are

$$2^{\frac{s}{2}} \leq R(s,s) \leq 4^{s-1}.$$

Define $R(G, G)$, the Ramsey number of a simple graph G, as the smallest n for which a 2-edge-colouring of K_n yields a monochrome copy of G.

Theorem 5.7. *For any graph G of maximum degree Δ,*

$$R(G,G) \leq cn$$

where c depends only on Δ.

Corollary 5.8. *For any graph of maximum degree Δ,*

$$R(s,s) \leq cn,$$

where c depends only on Δ.

Recall that Turan's theorem states that $T(n, r)$ has no K_{r+1} subgraphs. Denote by K_{r+1}^s, the $(r+1)$-partite graph with $r+1$ parts of size s.

Theorem 5.9. *For all integers r, $s \geq 1$, every $\epsilon > 0$, there exists an integer n_0 such that for every graph with $n \geq n_0$ vertices, and at least $t(n,r) + \epsilon n^2$ edges, contains a K^s_{r+1} subgraph.*

This says that if we add just ϵn^2 edges to $T(n,r)$, we get not only K_{r+1}, but K^s_{r+1} for any given integers. Thus the graph has many K_{r+1} subgraphs.