From Douglas West, Intro to Graph Theory

1. alpha'(G) is the matching number, i.e the size of a maximum matching. beta'(G) is the size of the smallest edge cover. An edge cover is a set of edges L so that each edge in G is incident with at least one edge of L. *Cuong, Thursday, Jan. 30.*

3.1.22. Theorem. (Gallai [1959]) If G is a graph without isolated vertices, then $\alpha'(G) + \beta'(G) = n(G)$.

Proof: From a maximum matching M, we will construct an edge cover of size $n(G) - |M|$. Since a smallest edge cover is no bigger than this cover, this will imply that $\beta'(G) \le n(G) - \alpha'(G)$. Also, from a minimum edge cover L, we will construct a matching of size $n(G) - |L|$. Since a largest matching is no smaller than this matching, this will imply that $\alpha'(G) \ge n(G) - \beta'(G)$. These two inequalities complete the proof.

Let M be a maximum matching in G. We construct an edge cover of G by adding to M one edge incident to each unsaturated vertex. We have used one edge for each vertex, except that each edge of M takes care of two vertices, so the total size of this edge cover is $n(G) - |M|$, as desired.

Now let L be a minimum edge cover. If both endpoints of an edge e belong to edges in L other than e, then $e \notin L$, since $L - \{e\}$ is also an edge cover. Hence each component formed by edges of L has at most one vertex of degree exceeding 1 and is a star (a tree with at most one non-leaf). Let k be the number of these components. Since L has one edge for each non-central vertex in each star, we have $|L| = n(G) - k$. We form a matching M of size $k = n(G) - |L|$ by choosing one edge from each star in L. ■

3.1.23. Example. The graph below has 13 vertices. A matching of size 4 appears in bold, and adding the solid edges yields an edge cover of size 9. The dashed edges are not needed in the cover. The edge cover consists of four stars; from each we extract one edge (bold) to form the matching. ■

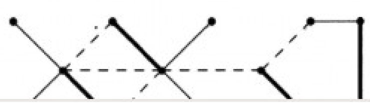

2. Dominating sets. *Beth, Thursday, Feb. 7*

3.1.26. Definition. In a graph G, a set $S \subseteq V(G)$ is a **dominating set** if every vertex not in S has a neighbor in S. The **domination number** $\gamma(G)$ is the minimum size of a dominating set in G.

3.1.27. Example. The graph G below has a minimal dominating set of size 4 (circles) and a minimum dominating set of size 3 (squares): $\gamma(G) = 3$. ■

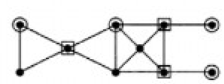

Berge [1962] introduced the notion of domination. Ore [1962] coined this terminology, and the notation $\gamma(G)$ appeared in an early survey (Cockayne–Hedetniemi [1977]). An entire book (Haynes–Hedetniemi–Slater [1998]) is devoted to domination and its variations.

3.1.28. Example. Covering the vertex set with stars may not require as many stars as covering the edge set. When a graph G has no isolated vertices, every vertex cover *is* a dominating set, so $\gamma(G) \le \beta(G)$. The difference can be large; $\gamma(K_n) = 1$, but $\beta(K_n) = n - 1$. ■

When studying domination as an extremal problem, we try to obtain bounds in terms of other graph parameters, such as the order and the minimum degree. A vertex of degree k dominates itself and k other vertices; thus every dominating set in a k-regular graph G has size at least $n(G)/(k+1)$. For every graph with minimum degree k, a greedy algorithm produces a dominating set not too much bigger than this.

3.1.30. Theorem. (Arnautov [1974], Payan [1975]) Every n-vertex graph with minimum degree k has a dominating set of size at most $n\frac{1+\ln(k+1)}{k+1}$.

Proof: (Alon [1990]) Let G be a graph with minimum degree k. Given $S \subseteq V(G)$, let U be the set of vertices not dominated by S. We claim that some vertex y outside S dominates at least $|U|(k+1)/n$ vertices of U. Each vertex in U has at least k neighbors, so $\sum_{v \in U} |N[v]| \ge |U|(k+1)$. Each vertex of G is counted at most n times by these $|U|$ sets, so some vertex y appears at least $|U|(k+1)/n$ times and satisfies the claim.

We iteratively select a vertex that dominates the most of the remaining undominated vertices. We have proved that when r undominated vertices remain, after the next selection at most $r(1-(k+1)/n)$ undominated vertices remain. Hence after $n\frac{\ln(k+1)}{k+1}$ steps the number of undominated vertices is at most

$$n(1-\tfrac{k+1}{n})^{n\ln(k+1)/(k+1)} < ne^{-\ln(k+1)} = \tfrac{n}{k+1}$$

The selected vertices and these remaining undominated vertices together form a dominating set of size at most $n\frac{1+\ln(k+1)}{k+1}$. ■

3. Chromatic number of the Carthesian product of graphs. *Nora, Thursday, Jan. 6.*

5.1.9. Definition. The **cartesian product** of G and H, written $G \square H$, is the graph with vertex set $V(G) \times V(H)$ specified by putting (u, v) adjacent to (u', v') if and only if (1) $u = u'$ and $vv' \in E(H)$, or (2) $v = v'$ and $uu' \in E(G)$.

5.1.10. Example. The cartesian product operation is symmetric; $G \square H \cong H \square G$. Below we show $C_3 \square C_4$. The hypercube is another familiar example: $Q_k = Q_{k-1} \square K_2$ when $k \ge 1$. The m-by-n **grid** is the cartesian product $P_m \square P_n$.

In general, $G \square H$ decomposes into copies of H for each vertex of G and copies of G for each vertex of H (Exercise 10). We use $\square$ instead of $\times$ to avoid confusion with other product operations, reserving $\times$ for the cartesian product of vertex sets. The symbol $\square$, due to Rödl, evokes the identity $K_2 \square K_2 = C_4$. ■

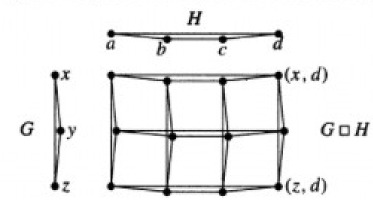

5.1.11. Proposition. (Vizing [1963], Aberth [1964]) $\chi(G \square H) = \max\{\chi(G), \chi(H)\}$.

Proof: The cartesian product $G \square H$ contains copies of G and H as subgraphs, so $\chi(G \square H) \ge \max\{\chi(G), \chi(H)\}$.

Let $k = \max\{\chi(G), \chi(H)\}$. To prove the upper bound, we produce a proper k-coloring of $G \square H$ using optimal colorings of G and H. Let g be a proper $\chi(G)$-coloring of G, and let h be a proper $\chi(H)$-coloring of H. Define a coloring f of $G \square H$ by letting $f(u, v)$ be the congruence class of $g(u) + h(v)$ modulo k. Thus f assigns colors to $V(G \square H)$ from a set of size k.

We claim that f properly colors $G \square H$. If (u, v) and (u', v') are adjacent in $G \square H$, then $g(u) + h(v)$ and $g(u') + h(v')$ agree in one summand and differ by between 1 and k in the other. Since the difference of the two sums is between 1 and k, they lie in different congruence classes modulo k. ■

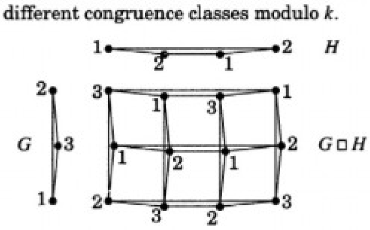

The cartesian product allows us to compute chromatic numbers by computing independence numbers, because a graph G is m-colorable if and only if the cartesian product $G \square K_m$ has an independent set of size $n(G)$ (Exercise 31).

4. Colouring with orientations. *Kyle, Tuesday, Feb. 4*

5.1.21. Theorem. Gallai–Roy–Vitaver Theorem (Gallai [1968], Roy [1967], Vitaver [1962]) If D is an orientation of G with longest path length $l(D)$, then $\chi(G) \le 1 + l(D)$. Furthermore, equality holds for some orientation of G.

Proof: Let D be an orientation of G. Let D' be a maximal subdigraph of D that contains no cycle (in the example below, uv is the only edge of D not in D'). Note that D' includes all vertices of G. Color $V(G)$ by letting $f(v)$ be 1 plus the length of the longest path in D' that ends at v.

Let P be a path in D', and let u be the first vertex of P. Every path in D' ending at u has no other vertex on P, since D' is acyclic. Therefore, each path ending at u (including the longest such path) can be lengthened along P. This implies that f strictly increases along each path in D'.

The coloring f uses colors 1 through $1+l(D')$ on $V(D')$ (which is also $V(G)$). We claim that f is a proper coloring of G. For each $uv \in E(D)$, there is a path in D' between its endpoints (since uv is an edge of D' or its addition to D' creates a cycle). This implies that $f(u) \ne f(v)$, since f increases along paths of D'.

To prove the second statement, we construct an orientation D^* such that $l(D^*) \le \chi(G) - 1$. Let f be an optimal coloring of G. For each edge uv in G, orient it from u to v in D^* if and only if $f(u) < f(v)$. Since f is a proper coloring, 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 $l(D^*) \le \chi(G) - 1$. ∎

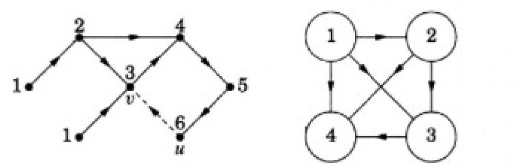

D' and D D^*

5. Turan graphs. Note that the lemma corresponds to a problem on assignment 2. *Thomas, Tuesday, Feb. 4.*

5.2.6. Definition. A **complete multipartite graph** is a simple graph G whose vertices can be partitioned into sets so that $u \leftrightarrow v$ if and only if u and v belong to different sets of the partition. Equivalently, every component of $\overline{G}$ is a complete graph. When $k \ge 2$, we write $K_{n_1,\dots,n_k}$ for the complete k-partite graph with partite sets of sizes $n_1, \dots, n_k$ and complement $K_{n_1} + \cdots + K_{n_k}$.

We use this notation only for $k > 1$, since K_n denotes a complete graph. A complete k-partite graph is k-chromatic; the partite sets are the color classes in the only proper k-coloring. Also, since a vertex in a partite set of size t has degree $n(G) - t$, the edges can be counted using the degree-sum formula (Exercise 18). Which distribution of vertices to partite sets maximizes $e(G)$?

5.2.7. Example. *The Turán graph.* The **Turán graph** $T_{n,r}$ is the complete r-partite graph with n vertices whose partite sets differ in size by at most 1. By the pigeonhole principle (see Appendix A), some partite set has size at least $\lceil n/r \rceil$ and some has size at most $\lfloor n/r \rfloor$. Therefore, differing by at most 1 means that they all have size $\lfloor n/r \rfloor$ or $\lceil n/r \rceil$.

Let $a = \lfloor n/r \rfloor$. After putting a vertices in each partite set, $b = n - ra$ remain, so $T_{n,r}$ has b partite sets of size $a + 1$ and $r - b$ partite sets of size a. Thus the defining condition on $T_{n,r}$ specifies a single isomorphism class. ∎

5.2.8. Lemma. Among simple r-partite (that is, r-colorable) graphs with n vertices, the Turán graph is the unique graph with the most edges.

5.2.9. Theorem. (Turán [1941]) Among the n-vertex simple graphs with no $r+1$-clique, $T_{n,r}$ has the maximum number of edges.

Proof: The Turán graph $T_{n,r}$, like every r-colorable graph, has no $r+1$-clique, since each partite set contributes at most one vertex to each clique. If we can prove that the maximum is achieved by an r-partite graph, then Lemma 5.2.8 implies that the maximum is achieved by $T_{n,r}$. Thus it suffices to prove that if G has no $r+1$-clique, then there is an r-partite graph H with the same vertex set as G and at least as many edges.

We prove this by induction on r. When $r = 1$, G and H have no edges. For the induction step, consider $r > 1$. Let G be an n-vertex graph with no $r+1$-clique, and let $x \in V(G)$ be a vertex of degree $k = \Delta(G)$. Let G' be the subgraph of G induced by the neighbors of x. Since x is adjacent to every vertex in G' and G has no $r+1$-clique, the graph G' has no r-clique. We can thus apply the induction hypothesis to G'; this yields an $r-1$-partite graph H' with vertex set $N(x)$ such that $e(H') \geq e(G')$.

Let H be the graph formed from H' by joining all of $N(x)$ to all of $S = V(G) - N(x)$. Since S is an independent set, H is r-partite. We claim that $e(H) \geq e(G)$. By construction, $e(H) = e(H') + k(n-k)$. We also have $e(G) \leq e(G') + \sum_{v \in S} d_G(v)$, since the sum counts each edge of G once for each endpoint it has outside $V(G')$. Since $\Delta(G) = k$, we have $d_G(v) \leq k$ for each $v \in S$, and $|S| = n-k$, so $\sum_{v \in S} d_G(v) \leq k(n-k)$. As desired, we have

$$e(G) \leq e(G') + (n-k)k \leq e(H') + k(n-k) = e(H) \qquad \blacksquare$$

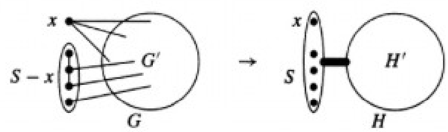

6. Connectivity of k-critical graphs. *Melanie, Thursday, Jan. 30*

5.2.15. Lemma. (Kainen) Let G be a graph with $\chi(G) > k$, and let X, Y be a partition of $V(G)$. If $G[X]$ and $G[Y]$ are k-colorable, then the edge cut $[X, Y]$ has at least k edges.

Proof: Let $X_1, \ldots, X_k$ and $Y_1, \ldots, Y_k$ be the partitions of X and Y formed by the color classes in proper k-colorings of $G[X]$ and $G[Y]$. If there is no edge between X_i and Y_j, then $X_i \cup Y_j$ is an independent set in G. We show that if $|[X, Y]| < k$, then we can combine color classes from $G[X]$ and $G[Y]$ in pairs to form a proper k-coloring of G.

Form a bipartite graph H with vertices $X_1, \ldots, X_k$ and $Y_1, \ldots, Y_k$, putting $X_iY_j \in E(H)$ if in G there is *no edge* between the set X_i and the set Y_j. If $|[X, Y]| < k$, then H has more than $k(k-1)$ edges. Since m vertices can cover at most km edges in a subgraph of $K_{k,k}$, $E(H)$ cannot be covered by $k-1$ vertices. By the König–Egerváry Theorem, H therefore has a perfect matching M.

In G, we give color i to all of X_i and all of the set Y_j to which it is matched by M. Since there are no edges joining X_i and Y_j, doing this for all i produces a proper k-coloring of G, which contradicts the hypothesis that $\chi(G) > k$. Hence we conclude that $|[X, Y]| \geq k$. $\blacksquare$

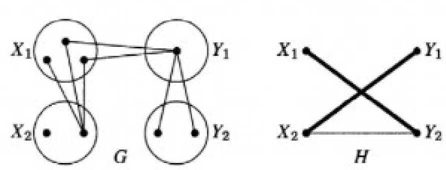

5.2.16. Theorem. (Dirac [1953]) Every k-critical graph is $k-1$-edge-connected.

Proof: Let G be a k-critical graph, and let $[X, Y]$ be a minimum edge cut. Since G is k-critical, $G[X]$ and $G[Y]$ are $k-1$-colorable. Applied with $k-1$ as the parameter, Lemma 5.2.15 then states that $|[X, Y]| \geq k-1$. $\blacksquare$

7. Theorem by Berge that a matching is maximum if and only if there is no augmenting path. For proof see Notes3. - *Poppy, Thursday, Jan. 30.*

8. Stable marriage algorithm and proof (*Kyle, Feb 11)*

STABLE MATCHINGS (optional)

Instead of optimizing total weight for a matching, we may try to optimize using preferences. Given n men and n women; we want to establish n "stable" marriages. If man x and woman a are paired with other partners, but x prefers a to his current partner and a prefers x to her current partner, then they might leave their current partners and switch to each other. In this situation we say that the unmatched pair (x, a) is an **unstable pair**.

3.2.15. Definition. A perfect matching is a **stable matching** if it yields no unstable unmatched pair.

3.2.16. Example. Given men x, y, z, w, women a, b, c, d, and preferences listed below, the matching $\{xa, yb, zd, wc\}$ is a stable matching. ∎

Men $\{x, y, z, w\}$	Women $\{a, b, c, d\}$
$x : a > b > c > d$	$a : z > x > y > w$
$y : a > c > b > d$	$b : y > w > x > z$
$z : c > d > a > b$	$c : w > x > y > z$
$w : c > b > a > d$	$d : x > y > z > w$

3.2.17. Algorithm. (Gale–Shapley Proposal Algorithm)
Input: Preference rankings by each of n men and n women.
Idea: Produce a stable matching using proposals by maintaining information about who has proposed to whom and who has rejected whom.
Iteration: Each man proposes to the highest woman on his preference list who has not previously rejected him. If each woman receives exactly one proposal, stop and use the resulting matching. Otherwise, every woman receiving more than one proposal rejects all of them except the one that is highest on her preference list. Every woman receiving a proposal says "maybe" to the most attractive proposal received. ∎

3.2.18. Theorem. (Gale–Shapley [1962]) The Proposal Algorithm produces a stable matching.

Proof: The algorithm terminates (with some matching), because on each non-terminal iteration, the total length of the lists of potential mates for the men decreases. This can happen only n^2 times.

Key Observation: the sequence of proposals made by each man is nonincreasing in his preference list, and the sequence of men to whom a woman says "maybe" is nondecreasing in her preference list, culminating in the man assigned. This holds because men propose repeatedly to the same woman until rejected, and women say "maybe" to the same man until a better offer arrives.

If the result is not stable, then there is an unstable unmatched pair (x, a), with x matched to b and y matched to a. By the key observation, x never proposed to a during the algorithm, since a received a mate less desirable than x. The key observation also implies that x would not have proposed to b without earlier proposing to a. This contradiction confirms the stability of the result. ∎

9. List chromatic number of complete bipartite graphs *Emma*

8.4.25. Proposition. (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 $G = K_{m,m}$. Assign the distinct k-subsets of $[2k-1]$ as the lists for the vertices of X, and do the same for Y. Consider a choice function f. If f uses fewer than k distinct choices in X, then there is a k-set $S \subseteq [2k-1]$ not used, which means that no color was chosen for the vertex of X having S as its list. If f uses at least k colors on vertices of X, then there is a k-set $S \subseteq [2k-1]$ of colors used in X, and no color can be properly chosen for the vertex of Y with list S. ∎

10. List colouring of planar graphs *Hoa, Thursday, Feb. 13*

8.4.32. Theorem. (Thomassen [1994b]) Planar graphs are 5-choosable.

Proof: Adding edges cannot reduce the list chromatic number, so we may restrict our attention to plane graphs where the outer face is a cycle and every bounded face is a triangle. By induction on $n(G)$, we prove the stronger result that a coloring can be chosen even when two adjacent external vertices have distinct lists of size 1 and the other external vertices have lists of size 3. For the basis step ($n = 3$), a color remains available for the third vertex.

Now consider $n > 3$. Let v_p, v_1 be the vertices with fixed colors on the external cycle C. Let $v_1, \ldots, v_p$ be $V(C)$ in clockwise order.

Case 1: C has a chord v_iv_j with $1 \le i \le j-2 \le p-2$. We apply the induction hypothesis to the graph consisting of the cycle $v_1, \ldots, v_i, v_j, \ldots, v_p$ and its interior. This selects a proper coloring in which v_i, v_j receive some fixed colors. Next we apply the induction hypothesis to the graph consisting of the cycle $v_i, v_{i+1}, \ldots, v_j$ and its interior to complete the list coloring of G.

Case 2: C has no chord. Let $v_1, u_1, \ldots, u_m, v_3$ be the neighbors of v_2 in order ($3 = p$ is possible). Because bounded faces are triangles, G contains the path P with vertices $v_1, u_1, \ldots, u_m, v_3$. Since C is chordless, $u_1, \ldots, u_m$ are internal vertices, and the outer face of $G' = G - v_2$ is bounded by a cycle C' in which P replaces v_1, v_2, v_3.

Let c be the color assigned to v_1. Since $|L(v_2)| \ge 3$, we may choose distinct colors $x, y \in L(v_2) - \{c\}$. We reserve x, y for possible use on v_2 by forbidding x, y from $u_1, \ldots, u_m$. Since $|L(u_i)| \ge 5$, we have $|L(u_i) - \{x, y\}| \ge 3$. Hence we can apply the induction hypothesis to G', with $u_1, \ldots, u_m$ having lists of size at least 3 and other vertices having the same lists as in G. In the resulting coloring, v_1 and $u_1, \ldots, u_m$ have colors outside $\{x, y\}$. We extend this coloring to G by choosing for v_2 a color in $\{x, y\}$ that does not appear on v_3 in the coloring of G'. ■

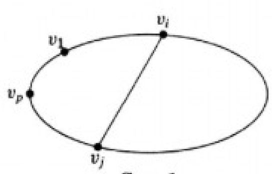

Case 1

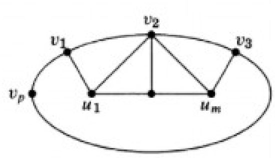

Case 2

Next round: graph models

Copy model -- Krapitsky and Redner
Geographic Threshold Model – Bradonjic, Hagberg, Percus
Kronecker graph model – Leskovec et al.
Ranking model -- Menczer
Random dot product graphs Young and Scheinerman
HOT model – Fabrikant, Koutsoupias, Papadimitriou
Geo-protean – Bonato, Janssen, Pralat
Grid plus long links -- Kleinberg
Scalefree percolation – Dreyfus, v.d. Hofstad, Hooghiemstra
Geometric Preferential Attachment – Flaxman, Frieze, Vera