Graphs and Digraphs

Ideas about what to cover, what to skip, and why: Cover Chapter 1 (Introduction to Graphs), Chapter 2 (Trees and Connectivity), Chapter 4 (Digraphs), Chapter 5 (The Automorphism Group of a Graph), Chapter 6 (Planar Graphs), Chapter 8 (Vertex Colorings), Chapter 12 (Extremal Graph Theory). Reason: Interesting, and I know something about them. Skipped chapters: let students pick something to study and present in class. If we still have time and I need to present more material, do Polya counting. Another possibility: the probabilistic method. (Note: In Fall 2013 I skipped Chapter 5 and had just enough time to get through the following topics in Chapters 6, 8, and 12: Euler's formula and applications, Kuratowski's Theorem, Brook's Theorem, and Ramsey Theory.)

Section 1.1: Graphs and Subgraphs

Theorem 1.4, p. 6: The sum of vertex degrees is twice the edges.

Proof: We have

$$\sum_{v\in V} deg(v) = \sum_{v\in V}\sum_{e\in E}\chi(v\in e) = \sum_{e\in E}\sum_{v\in V}\chi(v\in e) = \sum_{e\in E} 2 = 2|E|.$$

Corollary 1.5, p. 7: Every graph has an even number of odd degree vertices.

Proof: Reduce all terms mod 2. Then we have

$$\sum_{\substack{v\in V\\ \deg(v)\equiv 1}} 1 \equiv 0.$$

This says that the number of odd-degree vertices is an even number.

Theorem 1.6, p. 9: Isomorphic graphs have same number of vertices, same number of edges, and same vertex degree multiplicities.

Proof: Let $G = (V, E)$ and $H = (W, F)$ and let $\phi : V \to W$ be a graph isomorphism between G and H. Since ϕ is bijective, $|V| = |W|$. Since ϕ induces a bijective map $\phi' : E \to F$, $|E| = |F|$. For each $k \in \omega$ define

$$V_k = \{v \in V : deg(v) = k\}$$

and

$$W_k = \{w \in W : deg(w) = k\}.$$

We claim that $\phi(V_k) = W_k$ for each k. It suffices to show that $deg(\phi(v)) = deg(v)$ for each $v \in V$. For $v \in V$ we have

$$deg(\phi(v)) = \sum_{f \in F} \chi(\phi(v) \in f) = \sum_{e \in E} \chi(\phi(v) \in \phi'(e)),$$

so it suffices to show that $\phi(v) \in \phi'(e)$ if and only if $v \in e$. Since both ϕ and ϕ^{-1} are isomorphisms, it suffices to prove $v \in e$ implies $\phi(v) \in \phi'(e)$. But this is true by definition: $e = \{v, v'\}$ implies $\phi'(e) = \{\phi(v), \phi(v')\}$.

Theorem 1.7, p. 12: For integers r and n, there exists an r-regular graph of order n if and only if $0 \leq r \leq n-1$ and $\{r, n\}$ contains an even number.

Proof: If such a graph exists then it has an even number of odd degree vertices. So if n is odd then it has an odd number of degree r vertices, forcing r to be even. An upper bound on vertex degree is $n-1$, so $0 \leq r \leq n-1$.

Conversely, let $0 \leq r \leq n-1$ be given. First assume that $r = 2s$. Let $V = \{[0], \ldots, [n-1]\}$ modulo n. If $r = 0$ we set $E = \emptyset$. For $r \geq 2$ we set

$$N([k]) = \{[k+1], [k+2], \ldots, [k+s], [k-1], [k-2], \ldots, [k-s]\}.$$

We must verify that $|N([k])| = r$. So suppose $[k+i] = [k+j]$ where $i \neq j$ and $i, j \in \{1, \ldots, s, -1, \ldots, -s\}$. Then $n|(i-j)$. However, $|i-j| \leq 2s = r < n$, forcing $i = j$.

Next assume $r = 2s+1$. We set

$$N([k]) = \{[k+1], [k+2], \ldots, [k+s+1], [k-1], [k-2], \ldots, [k-s]\}.$$

Now we have $|i-j| \leq 2s+1 = r < n$, so again $|N([k])| = r$.

Theorem 1.8, p. 14: Let G be a bipartite graph. Then $e \leq \frac{v^2}{4}$.

Proof: If the vertex partition has sizes a and $v-a$ then

$$e \leq a(v-a) = \frac{v^2}{4} - (\frac{v}{2} - a)^2 \leq \frac{v^2}{4}.$$

Theorem 1.9, p.15: If G satisfies $v \geq 3$ and $e > \frac{v^2}{4}$ then G has a subgraph isomorphic to K_3.

Proof: If $v = 3$ and $e > \frac{9}{4}$ then $G \equiv K_3$. If $v = 4$ and $e > \frac{16}{4}$ then $G \equiv K_4$ or $G \equiv K_4 - 12$. Each of these has a subgraph isomorphic to K_3. If there is a counterexample to the theorem then let G be one with the fewest number of vertices. It must satisfy $v \geq 5$. Let $\{x, y\}$ be an edge. Then $N(x) \cap N(y) = \emptyset$, which forces $|N(x) \cup N(y)| = |N(x)| + |N(y)| \leq |V| = v$. Now consider $H = G - x - y$. Note $v_H \geq 3$. We have

$$e_H \geq e - v \geq \frac{v^2}{4} - v = \frac{v_H^2}{4} + 1 > \frac{v_H^2}{4},$$

which implies that H has a subgraph isomorphic to K_3, which implies that G does also. Contradiction. So there is no counterexample.

Section 1.3. Connected Graphs and Distance.

Walk: Sequence of vertices, consecutive pairs define edges.

Length of walk: The walk $(v_0, v_1, \ldots, v_k)$ has length k.

Trail: Walk, distinct edges.

Path: Walk, distinct vertices.

Theorem 1.16, p.32: Every uv walk contains subsequence forming a uv path.

Proof: Find a uv walk of minimal length among subsequences of the walk. This cannot have a repeated vertex a, because if $a, b, \ldots, z, a$ appears in the walk, we obtain a shorter one by excising $a, b, \ldots, z$.

Theorem 1.17, p. 33: Let G be a graph with vertex set $\{1, \ldots, n\}$. Let A be its adjacency matrix. The number of ij walks of length k is the ij entry of A^k.

Proof: By induction on k. Since the only walks of length 1 have same equal endpoints, A^0 counts walks of length 0. Now assume that $A^k = (b_{ij})$ counts ij of length k. The number of ij walks of length $k+1$ that pass through p before arriving at j is $b_{ip}a_{pj}$, for a total of $b_{i1}a_{1j} + \cdots + b_{in}a_{nj}$. This is the ij entry in A^{k+1}.

Information from A^k: The number of walks of the form (i, j, i) is the degree of i. Hence $deg(i) = A^2_{ii}$. The number of walks of the form (i, j, k, i) is twice the number of K_3 subgraphs that i belongs to, and is also A^3_{ii}.

Circuit: Closed trail.

Cycle: Closed trail which can be described as a path followed by an edge to the first vertex.

Girth: Smallest cycle length $g(G)$ in a graph G.

Connected Graph: Every pair of vertices has a walk between them.

Connected Component of a Graph: A connected subgraph which is not the proper subgraph of a connected subgraph. No vertex belong to two of these. No edge can belong to two of these. Since every vertex and edge belongs to a component, any graph is a disjoint union of components. The number of these is $k(G)$. Every connected subgraph is wholly contained in a unique component: if H is connected and contains x in C_1 and y in C_2, then every vertex in C_1 has a walk to every vertex in C_2 through H via x and y, which implies $C_1 \cup C_2 \cup H$ is connected.

Distance between two vertices in a connected graph: $d(u, v) =$ minimum length of uv paths. The distance function satisfies triangle inequality.

Theorem 1.18, p. 37: A non-trivial graph ($v \geq 2$) is bipartite if and only if it has no odd cycles.

Proof: Bipartite implies all even cycles: any path has to bounce back and forth, then return. Conversely, suppose there are no odd cycles. It suffices to show that each connected component is bipartite. If there are no edges, we're done. Otherwise, let $\{x, y\}$ be an edge, set $c(x) = 0$ and $c(y) = 1$, and let W be a subset of vertices of maximal size containing x and y such that $G[W]$ is connected and has a proper 2-coloring which includes $c(x) = 0$ and $c(y) = 1$. Then $G[W]$ has to be contained entirely in one connected component C of G We claim that $W = V(C)$, which implies that C is bipartite. Given $z \in V(C)\backslash W$, find a walk in C from x to z. Let z_0 be the first vertex along this walk that belongs to $V(C)\backslash W$. Let $N_W(z_0)$ be the set of neighbors of z_0 in W. By construction, $N_W(z_0) \neq \emptyset$. Write $N_W(z_0) = N_0 \cup N_1$, where $c(x) = 0$ for all $x \in N_0$ and $c(x) = 1$ for all $x \in N_1$. If one of these sets is empty then $G[W \cup \{z_0\}]$ is connected and properly 2-colorable, contradicting the maximality of W. So there must exist $x_0 \in N_0$ and $x_1 \in N_1$. Since $G[W]$ is connected and properly 2-colored, there is a path of odd length from x_0 to x_1. Appending z_0, x_0 to this path produces an odd cycle in G, contrary to hypothesis. So in fact $V(C)\backslash W = \emptyset$ and $W = V(C)$.

Eccentrity of a vertex in a connected graph: $e(v) = \max\{d(v, w) : w \in V(G)\}$.

Diameter and Radius of a connected graph: $diam(G) = \max\{e(v) : v \in V(G)\}$, $rad(G) = \min\{e(v) : v \in V(G)\}$.

Peripheral, antipodal, central vertices: Peripheral vertices satisfy $e(v) = diam(G)$, antipodal vertices satisfy $e(u) = e(v) = d(u,v)$, central vertices satisfy $e(v) = rad(G)$.

Theorem 1.19, p. 38: For a connected non-trivial graph G, $rad(G) \leq diam(G) \leq 2rad(G)$.

Proof: $rad(G) \leq diam(G)$ is immediate. Now choose u, v such that $d(u,v) = diam(G)$. Let w be a central vertex. Then

$$d(u,v) \leq d(u,w) + d(w,v) \leq e(w) + e(w) = 2rad(G).$$

Center and Periphery of a graph: $Cen(G) = G[\text{central vertices}]$ and $Per(G) = G[\text{peripheral vertices}]$.

Theorem 1.20, p. 39: Every $G = Cen(H)$ for some H.

Proof: See the construction on p. 40.

Theorem 1.22, p. 40: A non-trivial graph G is the periphery of another graph if and only if all or none of its vertices has eccentricity equal to 1.

Proof: If all vertices have eccentricity 1 then every vertex is peripheral, so $per(G) = G$. If no vertex has eccentricity 1, form $H = G \vee K_1$. Then $G = per(H)$. Now assume that G has a vertex x with $e(x) = 1$ and a vertex y with $e(y) > 1$. If $G = per(H)$, then there is a vertex $z \in V(G)$ such that $d_H(x,z) = diam(H) > 1$. This implies that z is not adjacent to x in G, which contradicts $e(x) = 1$. So $G \neq per(H)$ for any H.

Section 1.4: Multigraphs and Digraphs

Theorem 1.23: If D is a digraph of size m then

$$\sum_{v \in V(G)} od(v) = \sum_{v \in V(G)} id(v) = m.$$

Proof:

$$m = \sum_{(a,b) \in V \times V} \chi((a,b) \in E) = \sum_{a \in V} \sum_{b \in V} \chi((a,b) \in E) = \sum_{a \in V} od(a).$$

The other formula is proved the same way.

Section 2.1: Nonseparable Graphs

Cut vertex: $k(G-x) > k(G)$.

Theorem 2.1, p. 55: Every non-trivial connected graph G contains at least 2 non-cut vertices.

Proof: Let x be a peripheral vertex. Suppose $G-x$ is not connected. Write $G-x = G_1 + G_2 + \cdots + G_k$. Then x has an edge to at least one vertex x_i in each G_i. Let y be the eccentric vertex of x. Without loss of generality $y \in G_1$. Let P be an x_1, y_1 path in G_1. Then x_2, x_0, P is path in G of length $\geq diam(G)$: contradiction. Hence no peripheral vertex is a cut vertex. Hence y is also not a cut vertex.

Theorem 2.2, p. 56: A vertex x is a cut vertex of G if and only if there exist two other vertices y and z in the same component such that all y, z paths include x.

Proof: Let x be a cut vertex. Then it must be a cut vertex of some connected component G. Write $G = G_1 + \cdots + G_k$ with x adjacent to $x_i \in G_i$ for each i. There is no x_1, x_2 path in $G - x$, so every x_1, x_2 path in G includes x. Conversely, let y and z be vertices in the same component G such that every y, z path includes x. Then $G - x$ has no y, z paths and so must have at least two components, making x a cut vertex.

Nonseparable graph: At least 2 vertices and no cut vertices.

Theorem 2.3, p.57: In a nonseparable graph G of order ≥ 3, every pair of vertices belongs to a cycle.

Proof: Let x, y be given. We will produce a cycle including them by induction on $d(x, y)$. When $d(x, y) = 1$, $e = \{x, y\}$ is an edge in G. If $G - e$ is not connected then $G - e = G_1 + G_2$ is the component structure, and one of these contains a third vertex z. Without loss of generality $z \in G_1$. This makes x a cut vertex: contradiction. Therefore $G - e$ is connected and there is an x, y path P that excludes e. Hence x and y belong to the cycle $C = P + e$. For $d(x, y) = k \geq 2$, let $x_0 = x, x_1, \ldots, x_k = y$ be a path between them. Then $d(x, x_{k-1}) = k - 1$, so there is a cycle C containing x and x_{k-1} made up of two internally disjoint paths P and Q. Let $R = (r_0, r_1, \ldots, r_j)$ be an x, y path in $G - x_{k-1}$. We have $r_0 = x \in P \cup Q$. Let i be maximal such that $r_i \in P \cup Q$, and without loss of generality say that $r_i \in Q$. Then $Q[x, r_i] + R[r_i, y]$ and $P + \{x_{k-1}, y\}$ are internally disjoint, therefore x and y lie in the cycle formed by their union.

Block: Maximal nonseparable subgraph of a graph.

Fact: if $B \neq C$ are blocks in a graph then $|V(B) \cap V(C)| \leq 1$. Reason: Suppose $x, y \in B \cup C$. Then $B \cup C$ has a cut vertex z. But it can't be x because $(B - x) \cap C$ is connected through y, and similarly it can't be any vertex in $B \cap C$. If $z \in B \backslash C$ then $(B - z) \cup C$ is connected, and if $z \in C \backslash B$ then $B \cup (C - z)$ is connected. Hence there are no cut vertices in $B \cup C$: contradiction.

Fact: If $B \neq C$ are blocks and $V(B) \cap V(C) = \{x\}$, then x is a cut vertex of G. Reason: Let $b \in V(B) \backslash \{x\}$ and let $c \in V(C) \backslash \{x\}$. If $G - x$ is connected then there is a b, c path P that excludes x. However, this makes $B \cup C \cup P$ nonseparable: contradiction.

Fact: If x is a cut vertex of a connected G then x belongs to at least two blocks. Reason: Write $G - x = G_1 + G_2 + \cdots + G_k$. For each $i \leq k$ let $x_i \in V(G_i)$ be such that $\{x, x_i\} \in E(G)$. Let B_i be the unique block containing $\{x, x_i\}$. Since there is no path from x_i to x_j in $G - x$, $(B_i \cup B_j) - x$ is not connected, hence $B_i \neq B_j$. Hence x belongs to at least k blocks.

Taken together, these facts say that the cut vertices are the vertices shared by blocks and that no two blocks share more than one vertex.

Fact: if u and v belong to the same block B and if $\{u, v\}$ is an edge in G, then this belongs to B. Reason: If the edge falls in some block C different from B, then the two blocks would share the vertices u and v: contradiction. Hence for any subset $W \subseteq V(B)$, $G[W]$ is a subgraph of B.

End block: A block containing exactly one cut vertex, i.e. intersecting exactly one other block in a vertex.

Block-cut graph: Let G be a connected graph with cut vertices $x_1, \ldots, x_a$ and blocks $B_1, \ldots, B_b$. Form the bipartite graph with vertices $x_1, \ldots, x_a, B_1, \ldots, B_b$ and an edge of the form $\{x_i, B_j\}$ if and only if $x_i \in B_j$.

Lemma: The block-cut graph is connected and acyclic.

Proof: By construction, every block has an edge to a cut vertex. So it suffices to show that every pair of cut vertices has a path between them in the block-cut graph. Let x_i and x_j be cut vertices. Let $u_0, u_1, \ldots, u_k$ be a path between them in G. Then each edge $\{u_i, u_j\}$ belongs to a unique block. Record the sequence of blocks corresponding to these edges. Every time there is a change in block, the corresponding vertex in the sequence belongs to two blocks, so must be a cut vertex. Hence the blocks and cut vertices along the way form a path in the block-cut graph.

Next, assume there is a cycle in the block-cut graph. Each transition x, B, x' corresponds to at least one path $P_{x,x'} \subseteq B$. Choosing one path per transition x, B, X' in the cycle, there is a closed walk beginning and ending at a cut vertex x in G. Follow this walk starting at x and continuing as far along it as possible through distinct vertices until we arrive at a vertex y via the edge $\{y', y\}$. If the next step in the walk is in the same path that gave rise to it, then the next vertex cannot be y'. But if the next step in the walk is not in the same path that gave rise to it, then again the next vertex cannot by y' because no edge can belong to two blocks. Hence the next step returns to a vertex previously encountered on the path from x to y. This gives rise to a cycle. Hence the cycle edges belong to one block B_i. Therefore two different paths contribute to this cycle. This implies that two different blocks $B_j \neq B_k$ contribute to this cycle. So there is an edge in $B_i \cap B_j$ and an edge in $B_i \cap B_k$. But either $i \neq k$ or $j \neq k$, so there is a cycle edge that lives in two blocks. Contradiction: an edge lives in only one block. Therefore there are no cycles in the block-cut graph.

Fact: If x and y are vertices and P and Q are distinct xy paths, then $P + Q$ contains a cycle. Reason: Find the longest initial segment shared by P and Q. The terminal vertex t cannot be y because the paths are distinct. The paths converge at some point. Let c be the first vertex after y at which the two paths converge. Then t and c have two internally disjoint paths between them.

Fact: If x and y belong to different blocks B and B' of a connected graph and v is a cut vertex on any xy-path, then x and y belong to different connected components of $G - v$. Reason: xy path implies unique BB' path containing v. So every xy path passes through v.

Theorem 2.7, p. 58: Every connected graph containing cut vertices has at least two end blocks.

Proof: Let P be a path of maximal length in the block-cut graph. Let u be an endpoint of P. Then u must have degree 1, otherwise it has an edge not belonging to P to a vertex along P, forming a cycle: contradiction. Since every cut vertex has degree ≥ 2 in the block-cut graph, u must be a block and it must contain exactly one cut vertex, i.e. u is an end block. The other endpoint of P must be another end block.

Theorem 2.8, p. 59: Let G be a connected graph with cut vertices. Then G contains a cut-vertex such that all (or all but one) of the blocks it lives in is an end block.

Proof: Peripheral vertices of a tree are leaves. In the block-cut graph, leaves are blocks, not cut-vertices. Let B be a peripheral vertex of the block-cut graph, and let x be the vertex adjacent to it on the geodesic from B to an eccentric vertex C of B. Since the block-cut graph is acyclic, x is adjacent to exactly one block along the path. If x is adjacent to another block B' then it cannot be along the path. Since there is a path of length $diam(G)$ from B' to C, B' is a peripheral vertex and has degree 1, which makes it an end block. So all but possibly one block that x is adjacent to is an end block.

Theorem 2.9: The center of a connected graph is a subgraph of a single block.

Proof: It suffices to show that the central vertices belong to a single block. Suppose not. Let c_1 and c_2 be central vertices in different blocks. Let P be a path between them. Since the edges of P are wholly contained in blocks, there has to be a consecutive pair of edges that lie in different blocks. The vertex shared by these edges is neither c_1 nor c_2 and must be a cut vertex. Call it v. By facts we have established above, there is no c_1c_2 path in $G-v$, so c_1 and c_2 are in two different connected components of $G-v$. Now find $u \in G$ so that $d(v,u) = e(v)$. Then u lives in a connected component which must be different from, say, c_2. Since every c_2u path passes through v, $d(c_2,u) = d(c_2,v) + d(v,u) = d(c_2,v) + e(v) > e(v)$. This contradicts the centrality of c_2. Therefore all central vertices belong to the same block.

Section 2.2: Trees

Bridge: edge e in graph G such that $k(G-e) > k(G)$.

Theorem 2.10, p. 62: Let G be connected. e is a bridge edge iff e does not belong to a cycle.

Proof: If e belongs to a cycle then it's removal does not disconnect G so it is not a bridge. Conversely, if e is not a bridge then $G-e$ is connected, so there is a path between the endpoints of e in $G-e$, forming a cycle containing e.

Tree: Connected graph, no cycles (all edges bridge edges).

Theorem 2.11, p. 64: A graph is a tree iff every pair of vertices is joined by a unique path.

Proof: Tree implies no cycles implies unique paths. Unique paths implies no cycles implies tree.

Corollary 2.12, p. 65: Every nontrivial tree contains at least two leaves.

Proof: Peripheral vertices in a tree are leaves.

Theorem 2.13, p. 65: In a tree, $m = n - 1$.

Proof: By induction on n. True when $n = 1$. More generally, delete a leaf vertex to drop down to induction hypothesis.

Theorem 2.14, p. 65: Skip.

Theorem 2.15, p. 66: Skip.

Corollary 2.16, p. 66: A forest of k trees satisfies $m = n - k$.

Proof: Adding $k - 1$ edges produces a tree with $m = n - 1$.

Theorem 2.17, p. 67: No cycles and $m = n - 1$ implies tree.

Proof: By Corollary 2.16, $k = n - m = 1$.

Theorem 2.18, p. 67: Connected and $m = n - 1$ implies tree.

Proof: Delete cycle edges until we have a tree. Then $m = n-1$, so no edges were deleted.

Fact: Any connected collection of n edges encompassing $n + 1$ vertices is a tree.

Proof: By induction on n. When $n = 1$, we have one edge and 2 vertices, so a tree. Now assume n edges encompassing $n + 1$ vertices is a tree. Consider $n + 1$ edges encompassing $n + 2$ vertices. The vertex degree sum implies that the sum of the $n + 2$ vertex degrees is $2n + 2$. This implies that at least one vertex degree is 1. Deleting this vertex from the graph on these edges, we are left with n edges and $n + 1$ vertices. These form a tree by the induction hypothesis, so the entire collection is a tree since it is the addition of a non-cycle edge.

Theorem 2.19, p. 67: Equivalent conditions for a tree.

Theorem 2.20, p. 67: If G is a graph with $\delta(G) \geq k - 1$, then G contains an isomorphic copy of every tree on k vertices.

Proof: By induction on k. True for $k = 1$ and $k = 2$. More generally, let $k \geq 3$ with $\delta(G) \geq k - 1$. Given T with k vertices and $k - 1$ edges, delete a leaf u and its edge $\{u, v\}$, producing T'. There is a copy of T' in G. The image of v in G has degree $\geq k-1$, yet at most $k-2$ of the edges out of v belong to T', so we can use one more to reproduce T.

Section 2.3: Spanning Trees

Theorem 2.23, p. 74: There are n^{n-2} trees of order n.

Proof: Given a tree on n vertices, do the following: find the smallest leaf x, record the vertex y it is adjacent to, then delete x from the tree, producing a smaller tree. Keep on going until there is a single edge left. Since there are $n-1$ edges, we produce $n-2$ numbers. We will prove that the mapping from trees to lists is bijective by induction on $n \geq 3$. Hence the number of trees is the number of lists, which proves the formula.

Base case: $n = 3$. Let the vertices be x_1, x_2, x_3. Each tree on these is determined by the central vertex. The tree with central vertex x_i gets mapped to x_i, so the mapping is bijective for $n = 3$.

Now assume the mapping is bijective for a given $n \geq 3$. Let S and T be trees on vertex set $\{x_1, \ldots, x_{n+1}\}$, and suppose $L(S) = L(T) = (y_1, \ldots, y_{n-1})$. Let x_s be the smallest leaf in S and let x_t be the smallest leaf in T. Let $S' = S - x_s$ and let $T' = T - x_t$. Since $L(S') = (y_2, \ldots, y_{n-1}) = L(T')$, the induction hypothesis yields $S' = T'$. This accounts for all the vertices except x_s in S and x_t in T, so $x_s = x_t$ and $S = T$. Hence the map is injective.

We will now show that the map is surjective. Let $(y_1, \ldots, y_{n-1})$ be a list in the characters $1, \ldots, n+1$. Let x_a be the smallest character not appearing on this list. Then $(y_2, \ldots, y_{n-1})$ is the image of a tree T on the vertices $\{x_1, \ldots, \widehat{x_a}, \ldots, x_{n+1}\}$ by the induction hypothesis. One of these vertices is y_1. Form the tree $\widehat{T} = T + \{y_1, x_a\}$. We claim that $\widehat{T}$ produces $(y_1, \ldots, y_{n-1})$. It suffices to prove that x_a is the smallest leaf in $\widehat{T}$. The leaves of $\widehat{T}$ are x_a and $L \backslash \{y_1\}$ where L denotes the set of leaves of T. But none of the leaves of T appear in the list $(y_2, \ldots, y_{n-1})$ by the way the list is generated. Since x_a is the smallest vertex not appearing in the list, x_a is the smallest leaf.

The proof of this theorem suggests how to produce a tree giving rise to a list $(y_1, \ldots, y_{n-2})$ with elements from $\{x_1, \ldots, x_n\}$: find the smallest x_a not in $\{y_1, \ldots, y_{n-2}\}$ and join to y_1. Then find the smallest x_b not in $(x_a, y_2, \ldots, y_{n-2})$ and join to y_2. Then find the smallest x_c not in $(x_a, x_b, y_3, \ldots, y_{n-2})$ and join to y_3. Keep on going. After we have found $n-2$ edges, there are two vertices not accounted for. Join these by an edge.

Spanning tree of a connected graph: a subgraph containing all the vertices of the graph which is a tree.

Ways to form a spanning tree: Find $n-1$ edges that form an acyclic subgraph. Or: chop out cycle edges one by one. Example of the first method: find a partition $V_0 \cup V_1 \cup \cdots \cup V_k$ of the vertex set such that $|V_0| = 1$ and for all $i \geq 1$ and $x \in V_i$ there exists $y(x) \in V_{i-1}$ such that $\{x, y(x)\} \in E$. The collection of edges $\{\{x, y(x)\} : x \notin V_0\}$ forms a spanning tree. Taking $V_0 = \{u\}$ for an arbitrary u and $V_i = \{x \in V : d(x, u) = i\}$ is one way to produce this partition. Another example: pick an arbitrary edge e_1. Having chosen e_1 through e_k, if there are any vertices outside all these edges, pick any edge e_{k+1} that contains a new vertex. (If the graph is connected, there will have to be one.) Keep on going until $n-1$ edges have been chosen.

Theorem 2.24, p. 75: Skip.

Proof of the Cauchy-Binet Theorem and the Matrix Tree Theorem

Cauchy-Binet Theorem: Assume $p \leq q$. Let $A = (a_{ij})$ be an $p \times q$ matrix, let $B = (b_{ij})$ be a $q \times p$ matrix, and write $AB = C = (c_{ij})$. Then

$$\det(AB) = \det(C_1, \ldots, C_p) = \det(\sum_{i=1}^{q} b_{i1} A_i, \ldots, \sum_{i=1}^{q} b_{ip} A_i) =$$

$$\sum_{1 \leq i_1, \ldots, i_p \leq q} b_{i_1 1} \cdots b_{i_p p} \det(A_{i_1}, \ldots, A_{i_p}) =$$

$$\sum_{1 \leq i_1 < i_2 < \cdots < i_p \leq q} \sum_{\sigma \in \mathcal{S}_p} b_{i_{\sigma(1)} 1} \cdots b_{i_{\sigma(p)} p} \det(A_{i_{\sigma(1)}}, \ldots, A_{i_{\sigma(p)}}) =$$

$$\sum_{1 \leq i_1 < i_2 < \cdots < i_p \leq q} \sum_{\sigma \in \mathcal{S}_p} b_{i_{\sigma(1)} 1} \cdots b_{i_{\sigma(p)} p} \operatorname{sgn}(\sigma) \det(A_{i_1}, \ldots, A_{i_p}) =$$

$$\sum_{I \in \binom{[q]}{p}} \det(A_I) \det(B_I)$$

where for a subset I of $[q]$ of size p, A_I is the submatrix of A using the p columns from I and B_I is the submatrix of B using the p rows from I.

Counting spanning trees: Let G be a graph with vertex set $V = \{x_1, \ldots, x_n\}$ and edge set $E = \{e_1, \ldots, e_m\}$ where $m \geq n-1$. Then the number of spanning trees of G is

$$\sum_{H \in \binom{E}{n-1}} \chi(H \text{ is a spanning tree}).$$

Given the resemblance of this formula to the Cauchy-Binet Theorem, it should not be surprising that there is a determinant formula for this expression.

Matrix-Tree Theorem: Let

$$C = ((-1)^{\chi(x_i = \min e_j)} \chi(x_i \in e_j))$$

where $1 \leq i \leq n-1$ and $1 \leq j \leq m$. Then the number of spanning trees is $\det(CC^T)$.

Example:

$G =$ 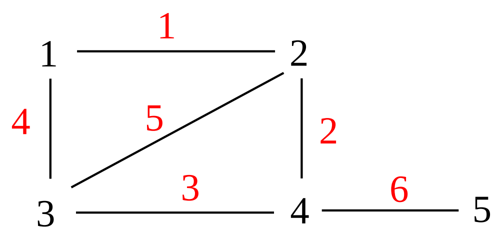

$$C = \begin{bmatrix} -1 & 0 & 0 & -1 & 0 & 0 \\ 1 & -1 & 0 & 0 & -1 & 0 \\ 0 & 0 & -1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & -1 \end{bmatrix}$$

$$\det(CC^T) = 8$$

Spanning trees:

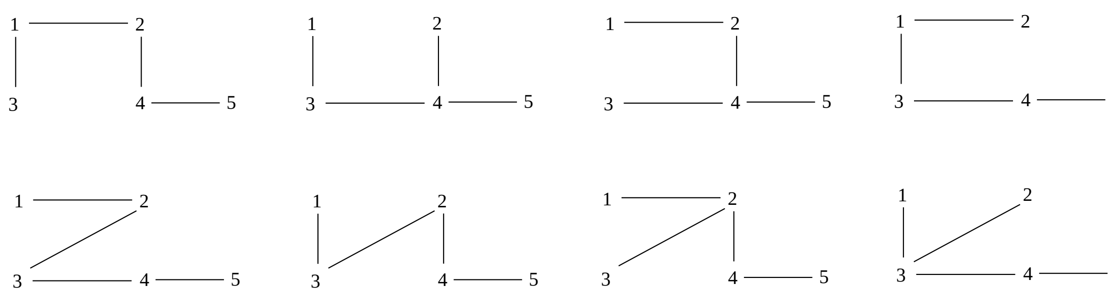

Proof of the Matrix-Tree Theorem: We have

$$\det(CC^T) = \sum_{I \in \binom{[m]}{n-1}} \det(C_I)\det(C_I^T) = \sum_{I \in \binom{[m]}{n-1}} \det(C_I)^2.$$

We will prove that

$$\det(C_I)^2 = \chi(\{e_i : i \in I\} \text{ is a spanning tree})$$

for each $I \in \binom{[m]}{n-1}$.

Let $I \in \binom{[m]}{n-1}$ be given. Name the corresponding edges $f_1, \ldots, f_{n-1}$. Then the ij-entry of C_I is 0 if $x_i \notin f_j$ and is ± 1 if $x_i \in f_j$. These edges form a spanning tree if and only if they are connected and encompass n vertices.

Case 1: $\{f_1, \ldots, f_{n-1}\}$ does not incorporate all n vertices. If x_n is isolated then each column of C_I has a 1 and a -1 in it, so the sum of its columns is the 0 vector, so its columns are linearly dependent and $\det(C_I) = 0$. If some other vertex x_k is isolated then row k in C_I is the 0 vector, which again implies $\det(C_I) = 0$.

Case 2: $\{f_1, \ldots, f_{n-1}\}$ encompasses all n vertices but is not connected. Each component has at least two vertices. The sum of all the rows corresponding to vertices in a component not containing n is 0, hence the rows are not linearly independent and $\det(C_I) = 0$.

Case 3: $\{f_1, \ldots, f_{n-1}\}$ incorporates all n vertices and is connected. The collection of edges forms a spanning tree. Clipping leaf vertices and edges, we can permute the rows and columns of C_I to produce a lower-triangular matrix with ± 1 in each diagonal entry. This implies $\det(C_I) = \pm 1$.

Minimum weight spanning tree: Spanning tree in connected graph whose edge weights have minimum sum.

Kruskal's Algorithm: Let G be a connected graph of order n with edge weights. Choose an edge e_1 of smallest possible weight. Having found the acyclic collection of edges $e_1, \ldots, e_k$, choose e_{k+1} of minimum weight that extends the acyclic collection. Keep on going until no longer possible.

Theorem: Kruskal's Algorithm produces a minimum weight spanning tree.

Proof: First note that when Kruskal's algorithm terminates, all vertices have been incorporated by connectedness. So the result is a spanning tree.

We will prove that $\{e_1, \ldots, e_k\}$ is a subset of a minimum weight spanning tree T_k for each k using an induction argument.

Base Case: Let T be any minimum weight spanning tree. If includes e_1, then set $T_1 = T$. If it doesn't include e_1, the subgraph $T + e_1$ contains a cycle of ≥ 3 edges. Delete f_1, any one of these not equal to e_1. By minimality of $weight(e_1)$, $weight(T + e_1 - f_1) \leq weight(T)$. Since $T + e_1 - f_1$ consists of $n - 1$ edges and encompasses all vertices, it is a tree. So in fact $T + e_1 - f_1$ is a minimum weight spanning tree. We set $T_1 = T - e_1 + f_1$.

Induction hypothesis: There exists a minimum weight spanning tree T_k that contains the edges $e_1, \ldots, e_k$.

We must now construct a minimum weight spanning tree T_{k+1} that contains $e_1, \ldots, e_{k+1}$. If $e_{k+1} \in T_k$ then we set $T_{k+1} = T_k$. But if $e_{k+1} \notin T_k$ then $T_k + e_{k+1}$ contains a cycle. One of the edges in this cycle is e_{k+1}. Since the collection $\{e_1, \ldots, e_{k+1}\}$ is acyclic, one of the edges in the cycle cannot be in this set. Call it f_{k+1}. Then $T_k + e_{k+1} - f_{k+1}$ is a spanning tree. Since $\{e_1, \ldots, e_k, f_{k+1}\} \subseteq T_k$, the collection is acyclic. By the way e_{k+1} was chosen by Kruskal's Algorithm, $weight(e_{k+1}) \leq weight(f_{k+1})$. Therefore $weight(T_k + e_{k+1} - f_{k+1}) \leq weight(T_k)$. Hence $T_k + e_{k+1} - f_{k+1}$ is a minimum weight spanning tree. We set $T_{k+1} = T_k + e_{k+1} - f_{k+1}$.

Prim's Algorithm: Let G be a connected graph of order n with edge weights. Choose an edge e_1 of smallest possible weight. Having found the acyclic collection of edges $e_1, \ldots, e_k$, choose e_{k+1} of minimum weight among all edges that introduce exactly one new vertex. Keep on going until no longer possible.

Theorem: Prim's Algorithm produces a minimum weight spanning tree.

Proof: Prim's algorithm produces a spanning tree by the same argument that Kruskal's algorithm does. We will prove that $\{e_1, \ldots, e_k\}$ is a subset of a minimum weight spanning tree T_k for each k using an induction argument.

Base Case: Same argument as in Kruskal's Algorithm.

Induction hypothesis: There exists a minimum weight spanning tree T_k that contains the edges $e_1, \ldots, e_k$.

We must now construct a minimum weight spanning tree T_{k+1} that contains $e_1, \ldots, e_{k+1}$. If $e_{k+1} \in T_k$ then we set $T_{k+1} = T_k$. But if $e_{k+1} \notin T_k$ then $T_k + e_{k+1}$ contains a cycle. One of the edges in this cycle is e_{k+1}. As in the proof of Kruskal's algorithm, we must chop out some f_{k+1} from this cycle not in $\{e_1, \ldots, e_{k+1}\}$. But we must take care that we can compare the weights of e_{k+1} and f_{k+1}. This requires that f_{k+1} have exactly one vertex in common with $\{e_1, \ldots, e_k\}$. Since the edge e_{k+1} includes a vertex x in common with $\{e_1, \ldots, e_k\}$ and a vertex y not in common with $\{e_1, \ldots, e_k\}$, the cycle it belongs to does also. Walking along the edges in the cycle in the direction of x to y, we eventually encounter another edge with a vertex y' not in $\{e_1, \ldots, e_k\}$ and a vertex x' in $\{e_1, \ldots, e_k\}$. This is the edge f_{k+1} that we choose.

Section 2.4: Connectivity and Edge Connectivity

Vertex Connectivity: $\kappa(G) =$ minimum number of vertices in a set S such that $G - S$ is disconnected or trivial.

Edge Connectivity: $\lambda(G) =$ minimum number of edges in a set S such that $G - S$ is disconnected or trivial.

Bounds on connectivity: $0 \leq \kappa(G), \lambda(G) \leq n - 1$.

Fact: $\kappa(K_n) = n - 1$. Reason: the removal of any vertex leaves a complete graph. Hence one has to remove all but one vertex to achieve disconnected or trivial.

Fact: Let F be a minimal edge cut of a connected graph G. Then $G - F = G_1 + G_2$ is the component structure and F consists of all edges between G_1 and G_2. Reason: Write $G - F = G_1 + G_2 + \cdots + G_k$. There has to be an edge $f \in F$ joining two connected components. Without loss of generality, assume these components are G_1 and G_2. Then $G_1 + G_2 + \cdots + G_k + f$ is connected, therefore $k = 2$ and f must join a vertex in G_1 to a vertex in G_2. Since there can be no surviving edges between G_1 and G_2, F contains every edge that exists between them in G.

Fact: Let G be connected and let U be a minimum vertex cut. Let $G - U = G_1 + \cdots + G_k$ be the component decomposition. Then for each G_i and u_j there is an edge of the form $\{g_i, u_i\}$ for some $g_i \in G_i$. Reason: This must be true if $|U| = 1$. Now consider $|U| \geq 2$. Let $u \in U$ be given. Then $G - (U \backslash \{u\})$ is connected. Let $x_i \in G_i$ be given. There must be a path in $G - (U \backslash \{u\})$ to u_i. The last edge out of G_i cannot be to any other G_j nor to a vertex in $U \backslash \{u\}$, hence must be to u. This edge is of the form $\{g_i, u\}$.

Fact: If $a + b = c$ with $a \geq 1$, $b \geq 1$ then $ab \geq c - 1$. Reason: $ab - c + 1 = ab - a - b + 1 = (a-1)(b-1) \geq 0$. So if $k \geq 1$ and $n - k \geq 1$ then $k(n-k) \geq n - 1$.

Theorem 2.28, p. 91: $\lambda(K_n) = n - 1$.

Proof: This is true when $n = 1$. Now consider $n \geq 2$. Suppose S is a minimal edge cut. Write $G - E = G_1 + G_2$. Then $S = K_{V(G_1),V(G_2)}$, therefore $n - 1 \geq |S| = k(n-k) \geq n - 1$ where $|G_1| = k$. Therefore $|S| = n - 1$.

Theorem 2.29, p. 92: For every graph G, $\kappa(G) \leq \lambda(G) \leq \delta(G)$.

Proof: We have $\kappa(K_n) = \lambda(K_n) = \delta(K_n) = n - 1$. Now assume the graph G is not complete. If it is disconnected then $\kappa(G) = \lambda(G) = 0$ and $\delta(G) \geq 0$. Now assume that G is connected and not complete. Since G is not complete, $\delta(G) \leq n - 2$ and this has to be an upper bound on $\lambda(G)$. Let X be a minimal edge cut. Then $G - X = G_1 + G_2$ is the component decomposition and X consists of all edges between G_1 and G_2. If every vertex in G_1 has an edge with every vertex in G_2 then $|X| = k(n-k) \geq n-1$, which contradicts $\lambda(G) \leq n - 2$. So in fact there is a vertex $x_1 \in G_1$ and a vertex $x_2 \in G_2$ such that $\{x_1, x_2\} \notin G$. We will choose one vertex per edge in X to create a vertex cut.

Let $\{u_1, u_2\} \in X$ be given. Then either $u_1 \neq x_1$ or $u_2 \neq x_2$. Choose the first vertex that satisfies this condition. Deleting it eliminates this edge from the graph but leaves the other vertex in place. When done, we have created a vertex cut of size $\leq |X|$ since it leaves no edge between x_1 and x_2, so $\kappa(G) \leq |X| = \lambda(G)$.

Theorem 2.30, p. 92: $\kappa(G) \leq \frac{2m}{n}$**.**

Proof: Average vertex degree is $\frac{2m}{n}$, so $\kappa(G) \leq \delta(G) \leq \frac{2m}{n}$.

Theorem 2.31, p. 93: For every cubic graph G, $\kappa(G) = \lambda(G)$.

Proof: The formula is true if G is not connected: both parameters are zero. Now assume connected. The formula is true if $\kappa(G) = 3$ since $3 = \kappa(G) \leq \lambda(G) \leq \delta = 3$. Now consider $1 \leq \kappa(G) \leq 2$.

If $\kappa(G) = 1$: we need only find an edge cut of size 1. Let u be a cut vertex. Then $G - u = G_1 + G_2$ or $G_1 + G_2 + G_3$. One of these components must have a single edge to u, giving rise to a bridge edge.

If $\kappa(G) = 2$: Let $\{u, v\}$ be an minimal vertex cut. We must find an edge cut of size 2. Then $G - u - v = G_1 + G_2$ or $G = G_1 + G_2 + G_3$. Considering the possible cases, we can always find an edge cut of size 2.

Theorem 2.32, p. 93: Let G be a graph of order n and let $1 \leq k \leq n-1$. If $\delta(G) \geq \frac{n+k-2}{2}$ then G is k-connected.

Proof: Let U be a vertex cut of size p. Let $G - U = G_1 + \cdots + G_j$ be the component decomposition. The smallest component has $\leq \frac{n-p}{2}$ vertices, which implies that $\frac{n+k-2}{2} \leq \delta(G) \leq \frac{n-p}{2} - 1 + p$. This forces $p \geq k$.

Theorem 2.33, p. 94: A nontrivial graph G satisfies $\lambda(G) \geq k$ if and only if, for every nonempty $W \subseteq V$, there are at least k edges joining W to $V - W$.

Proof: Assume $\lambda(G) \geq k$. Let W be given. If there were fewer than k edges between W and $V - W$ then you could disconnect the graph by removing these edges. But this contradicts $\lambda(G) \geq k$. So there have to be at least k edges between W and $V - W$. Conversely, suppose the W condition holds. Let U be a minimum vertex cut. Then $G - U = G_1 + G_2$ and U consists of all edges between G_1 and G_2. There must be at least k of these, so $\lambda(G) \geq k$.

Theorem 2.34, p. 95: Let G be a connected graph. Then $diam(G) = 2$ implies $\lambda(G) = \delta(G)$.

Proof: Let F be a minimum edge cut. Then $G - F = G_1 + G_2$ and F consists of all edges between G_1 and G_2. We need only show $|F| \geq \delta(G)$. Setting f_u = number of edges in F with endpoint u, we have

$$|F| = \sum_{u \in G_1} f_u.$$

Note that $f_u = \deg(u) -$ (# edges from u into G_1), hence

$$f_u \geq \deg(u) - |G_1| + 1 \geq \delta(G) - |G_1| + 1.$$

This implies

$$|F| \geq |G_1|(\delta(G) - |G_1| + 1).$$

We would like to show

$$|G_1|(\delta(G) - |G_1| + 1) \geq \delta(G).$$

Rearranging, this is equivalent to

$$\delta(G)(|G_1| - 1) \geq |G_1|(|G_1| - 1).$$

This will be true if $|G_1| = 1$, and for $|G_1| > 1$ we require that $\delta(G) \geq |G_1|$. Since we know that $\delta(G) \geq |F|$, we require that $|F| \geq |G_1|$. This will be true if each $f_u \geq 1$. If some $f_u = 0$, we can replace G_1 with G_2. Reason: for each $v \in G_2$ there is a path to u of length ≤ 2. The last edge connects two vertices in G_1 so the first edge connects v to a vertex in G_1. This implies $f_v \neq 0$ for each $v \in G_2$.

Section 2.5: Menger's Theorem

uv **separating set:** Assuming u and v are not adjacent, a uv separating set is a set of vertices S such that $G - S$ leaves u and v in different components. Denoting by s_{uv} the minimum size of a uv-separating set, we have $s_{uv} \leq \deg(u), \deg(v)$.

Fact 1: Let s_{uv} be the minimum number of vertices in a uv-separating set, and let p_{uv} the maximum size of a collection of internally disjoint uv paths. Then $p_{uv} \leq s_{uv}$. Reason: For any collection of k internally disjoint uv paths, every uv separating set must include at least one vertex from each path. Hence $k \leq s_{uv}$. This implies $p_{uv} \leq s_{uv}$.

Fact 2: Let S be a minimal uv separating set. Then every uv path must pass through some vertex of S (since S is separating) and for each $x \in S$ there is a uv path that passes through x (by minimality of $|S|$).

Fact 3: Every vertex in a minimal uv-separating set has degree ≥ 2 since it lies on a uv-path.

Fact 4: Set S be a minimal uv-separating set. Then for each $x \in S$ there is a uv path such that x is the first vertex in S on the path. Reason: Suppose $x_0 \in S$ is never first. Then removing all the vertices in S except x_0 disrupts all paths, which makes $S - \{x_0\}$ a uv-separating set, violating the minimality of S.

Fact 5: Let S be a minimal uv-separating set in G and suppose $x \in S$. Let s'_{uv} be the size of a minimal uv-separating set in $G - x$. Then $s'_{uv} = s_{uv} - 1$. Reason: $S - x$ is uv separating in $G - x$, hence $s'_{uv} \leq s_{uv} - 1$. Now suppose $s'_{uv} < s_{uv} - 1$. Let T be a uv-separating set in $G - x$. Removing T from G eliminates all uv paths avoiding x and removing x eliminates the rest, hence $T + x$ is a uv-separating set in G which is too small. Therefore $s'_{uv} = s_{uv} - 1$.

Fact 6: Let $w \neq u, v$ be a vertex that does not belong to any minimal uv separating set of G. Then $s_{uv} = s'_{uv}$, where s'_{uv}, where s'_{uv} is the size of a minimal uv separating set in $G' = G - w$. Reason: let S be a minimal uv separating set. Then it doesn't contain w and is a uv-separating set in $G - w$. This implies $s_{uv} \geq s'_{uv}$. Suppose that $s_{uv} > s'_{uv}$. Let T be a uv-separating set in $G - w$. If we remove T from G we leave only those uv-paths passing through w, and if we next remove w we disrupt all uv paths in G. This makes $T + w$ a uv separating set in G, and its size implies that it is minimal. Contradiction. Therefore $s_{uv} = s'_{uv}$.

Fact 7: Let S be a minimal uv separating set. If P is a path in $G - S$ with endpoint u and Q is a path in $G - S$ with endpoint v then $V(P) \cap V(Q) = \emptyset$. Reason: if they had a vertex in common then u and v belong to the same component of $G - S$.

Fact 8: Let G be a graph containing vertices $u \neq v$. Let S be a minimal uv-separating set. Then one of the following three scenarios must hold:

Scenario 1. $S \cap N(u) \cap N(v) \neq \emptyset$. (Corresponds to Case 1 of Menger's Theorem.)

Scenario 2. $S \not\subseteq N(u)$ and $S \not\subseteq N(v)$. (Corresponds to Case 3 of Menger's Theorem.)

Note that Scenarios 1 and 2 are not mutually exclusive.

Scenario 3. Neither Scenario 1 nor Scenario 2 hold. Hence $S \subseteq N(u)$ and $S \cap N(v) = \emptyset$ or $S \subseteq N(v)$ and $S \cap N(u) = \emptyset$. (Corresponds to Case 2 of Menger's Theorem.)

Theorem 2.36, p. 98 (Menger's Theorem): $p_{uv} = s_{uv}$.

Proof: By induction on m.

Base Case: $m = 0$: $p_{uv} = 0$ and $s_{uv} = 0$.

Induction hypothesis: for all graphs with $< m$ edges, $p_{uv} = s_{uv}$ for any u and v in the graph.

Now let G be a graph with m edges and let u, v be two non-adjacent vertices in G. There is nothing to prove if $s_{uv} \leq 1$, so we will assume that $s_{uv} \geq 2$. We will exploit the induction hypothesis by deleting at least one edge from the graph. We will organize the argument according to the three scenarios described above.

Case 1: G has a uv-separating set S of size s_{uv} that satisfies Scenario 1, i.e. $S \cap N(u) \cap N(v) \neq \emptyset$. Choose $x \in S$ adjacent to both u and v. Then $S - x$ is a minimal uv separating set in $G' = G - x$ and $s'_{uv} = s_{uv} - 1$. By the induction hypothesis, $p'_{uv} = s'_{uv} - 1$. Let $P_1, \ldots, P_{s'_{uv}-1}$ be internally disjoint uv paths in $G - x$. Then $P_1, \ldots, P_{s'_{uv}-1}, (u, x, v)$ is a collection of s_{uv} internally disjoint uv paths in G, which implies $p_{uv} \geq s_{uv}$, which implies $p_{uv} = s_{uv}$.

Case 2: G has a uv-separating set S of size s_{uv} that satisfies Scenario 2: $S \not\subseteq N(u)$ and $S \not\subseteq N(v)$. Define H_u to be the graph consisting of the union of all paths in $G - v$ originating from u that contain exactly one vertex from S. By Fact 3, $S \subseteq H_u$. We also have $H_u \subseteq G - v$, hence $|E(H_u)| \leq m - \deg(v) \leq m - s_{uv}$.

Claim: $|E(H_u)| < m - s_{uv}$. For suppose $|E(H_u)| = m - s_{uv}$. This implies $\deg(v) = s_{uv}$, which implies $E(H_u) = E(G - v)$. It also implies that $N(v)$ is another minimal uv-separating set. Choose $w \in N(v) \backslash S$. Since $\deg(w) \geq 2$, it belongs to an edge ww' where $w' \neq v$. This places $ww' \in E(H_u)$, which places w on a path from u to S with exactly one vertex in S. Truncating this path at w and appending the edge wv we obtain a uv path that avoids S completely. Contradiction. Therefore $|E(H_u)| < m - s_{uv}$.

Now form the graph G_u by adjoining a new vertex u' to H_u and joining each $x \in S$ to u'. Then $|E(G_u)| < m$.

Claim: $s_{uu'} = s_{uv}$. We certainly have $s_{uu'} \leq s_{uv}$ since S is a uu' separating set in G_u. To obtain $s_{uu'} \geq s_{uv}$ it suffices to show that every uu' separating set in G_u is a uv separating set in G. If T is not a uv-separating set in G then there is some uv path in G that avoids T, giving rise to a us path in H_u avoiding T, giving rise to a uu' path in G_u avoiding T. So the claim is proved.

By the induction hypothesis, there exist s_{uv} internally disjoint uu' paths in G_u. This gives rise to s_{uv} paths in G, endpoints of the form u and some $s \in S$, internally disjoint from S and sharing only the vertex u. We can similarly form H_v and find s_{uv} paths in G, endpoints of the form v and some $s \in S$, internally disjoint from S and sharing only the vertex v. These paths intersect only in S, so they can be glued together to form s_{uv} internally disjoint paths in G. Hence $p_{uv} \geq s_{uv}$, which implies $p_{uv} = s_{uv}$.

Case 3: None of the uv-separating sets of G of size s_{uv} fall in Case 1 or Case 2. For each such S, either $S \subseteq N(u)$ and $S \cap N(v) = \emptyset$ or $S \subseteq N(v)$ and

$S \cap N(u) = \emptyset$. Without loss of generality there exists a minimal uv-separating set S of the first type: $S \subseteq N(u)$ and $S \cap N(v) = \emptyset$. Let $P = (u, x, y, \ldots, v)$ be a uv path of minimal length. Then $y \notin N(u)$ by minimality of the length of P. Let e be the edge $\{x, y\}$. Write $G' = G - e$. Then $s'_{uv} \leq s_{uv}$.

Claim: $s'_{uv} = s_{uv}$. Reason: Suppose $s'_{uv} < s_{uv}$. Then $S' \cup \{x\}$ and $S' \cup \{y\}$ are uv-separating sets which are minimal by virtue of their size. Therefore $S' \cup \{x\} \subseteq N(u)$ and $S' \cup \{y\} \subseteq N(v)$. If $z \in S'$ then z is a member of a minimal uv-separating set which is adjacent to both u and v, which is not permitted in Case 3. Hence $S' = \emptyset$, which forces $s_{uv} = 1$, which contradicts $s_{uv} \geq 2$. Hence the claim is true.

By the induction hypothesis, G' has s_{uv} internally-disjoint uv-paths. Hence G does also.

Theorem 2.37 (Whitney's Theorem), p. 101: Let G be a non-trivial graph. Then $\kappa(G) \geq k$ iff every pair of vertices in G, not necessarily non-adjacent, has at least k internally-disjoint paths between them.

Proof: First assume that every pair of vertices has this property. Choosing two vertices at random, we see that there are at least $k+1$ vertices in the graph, namely the 2 vertices and the internal vertices from $k-1$ of the internally-disjoint paths. Now let X be an arbitrary set of $k-1$ vertices. Then $G - X$ has at least 2 vertices. Let a and b be vertices in G. At least 1 ab path survives, hence $G - X$ is connected. This implies $\kappa(G) \geq k$.

Conversely, suppose $\kappa(G) \geq k$. Let u and v be vertices in G. If non-adjacent, there are k internally disjoint paths joining them. If adjacent, then they are non-adjacent in $G' = G - e$. Hence there are $\kappa(G')$ internally-disjoint uv-paths in G', which yields $\kappa(G') + 1$ internally disjoint paths between u and v in G. We just need to verify that $\kappa(G') + 1 \geq k$. If $\kappa(G') + 1 < k$ then $\kappa(G') \leq k - 2$, hence there is a set of vertices X with $|X| \leq k - 2$ whose removal leaves 1 vertex or a disconnected graph. Both are impossible: the former implies G has $\leq k - 1$ vertices, which is impossible, and the latter implies $X \cup \{x\}$ is a separating set of G of size $\leq k - 1$, which is also impossible.

Corollaries 2.38, 2.39, 2.40, pp. 101 – 102: See exercises 3, 4, 5.

Theorem 2.41, p. 102: If G is k-connected and $k \geq 2$, then every k vertices of G lie in some cycle.

Proof: Choose k arbitrary vertices $v_1, \ldots, v_k$. Since the graph is 2-connected, v_1 and v_2 lie in a common cycle $C(2)$. So we have the base case of an induction argument. Assume that $v_1, \ldots, v_r$ lie on a common cycle $C(r)$ for some $2 \leq r < k$.

Case 1: The cycle contains only the vertices $v_1, \ldots, v_r$. Without loss of generality they appear in this order along the cycle. By Corollary 2.39 there are paths $v_{r+1}v_1, v_{r+1}v_2, \ldots, v_{r+1}v_r$ paths that share only the vertex v_{r+1}. Let $C(r+1)$ be the v_1v_r path along the cycle followed by the v_rv_{r+1} path followed by the $v_{r+1}v_1$ path.

Case 2: The path from v_1 to v_r along the cycle does contain other vertices. Let v_0 be another vertex. Without loss of generality the vertices $v_0, v_1, \ldots, v_r$ appear in this order along the cycle. Use these vertices to partition the cycle into subpaths $[v_0, v_1), [v_1, v_2), \ldots, [v_r, v_0)$. By Corollary 2.13 there are paths $v_{r+1}v_0, v_{r+1}v_1, \ldots, v_{r+1}v_r$ paths that share only the vertex v_{r+1}. Truncate these at the moment they hit the cycle. By the pigeon-hole principle, two of these paths fall in some subpath $[v_i, v_{i+1})$. Add these two paths to the cycle structure and delete the subpath along the cycle where these paths meet the cycle. In this way we don't discard any of the vertices $v_1, \ldots, v_r$ and we create a cycle $C(r+1)$ containing v_{r+1}.

Section 4.1: Strong Digraphs

Definition: A digraph is strongly connected iff for each $u \neq v$ there exists a directed uv path.

Theorem 4.1, p. 149: Every directed uv walk gives rise to a directed uv-path.

Proof: Cut out self-intersections.

Theorem 4.2, p.149: A digraph is strongly connected iff there exists a closed directed walk through all the vertices.

Proof: If such a walk exists the graph is strongly connected. Conversely, given a strongly connected digraph D, concatenate directed paths from each vertex to the next to build the closed directed walk.

Theorem 4.3 (Robbin's Theorem), p. 150: A nontrivial graph G has a strongly connected orientation iff it has no bridge edges ($\lambda(G) \geq 2$).

Proof: If G has a strongly-connected orientation then every edge lives in a cycle (the directed paths between the endpoints are internally disjoint), hence no edge is a bridge edge. Now assume that G has no bridge edges. We will construct the strong orientation.

Pick an edge at random. Its removal leaves a connected graph, so the edge lives in a cycle. Orient the edges around the cycle uniformly. We have found a strongly-orientable subgraph. We will show that for each strongly orientable subgraph $H \neq G$ there is a larger strongly orientable subgraph H'. So we can keep on enlarging until we produce a strong orientation of G.

If H contains all the vertices of G, we can toss in any missing edge and use an arbitrary orientation to produce H'.

If H does not contain all the vertices of G, Let $u \in V(G) \backslash V(H)$ be given. By connectedness there is a path from H to u that meets H in exactly one vertex. Let hv be the first edge. Then $G - hv$ is connected, so it contains a path P from v to H. Orient all the edges in $hv + P$ in the same direction and add them to H to produce a strongly connected H'.

Definition: A connected digraph is Eulerian if there is a closed directed walk through all edges used exactly once. More generally, a digraph is Eulerian if each connected component is Eulerian.

Theorem 4.4, p. 151: Let D be a non-trivial digraph. Then D is Eulerian (one can walk through all the edges of each connected component of the underlying graph using directed edges) iff $id(v) = od(v)$ for all vertices v.

Proof: If D is Eulerian you can start at any vertex v_0 in a connected component and walk around all the edges exactly once in the same direction, ending at v_0. As you go you can take a survey of the edges incident to each vertex, and you will find that for each one in there is a corresponding one out, determined by the closed Euler trail. When done you have accounted for all the edges and proved $id(v) = od(v)$ for all v.

Conversely, suppose the in-degree/out-degree condition holds. Then there must be a directed cycle in the graph: find a directed path of maximal length. The last out-degree must be at least one and so there must be a directed edge from the last vertex to a vertex along the path. Truncating if necessary we obtain a directed cycle. Now we have the basis of an induction argument according to number of edges. If just 3 edges, done. If more than this, start by finding a directed cycle C, then delete the edges to obtain D' satisfying the same in-degree/out-degree condition. Then D' is Eulerian trail by the induction hypothesis, and use C to glue together the Euler bits that had contact with C.

Definition: An open Euler trail in a connected digraph is a walk through all edges once that begins and ends at different vertices.

Theorem 4.5, p. 151: A connected digraph has an open Euler trail between u and v iff $od(u) = id(u) + 1$, $id(v) = od(v) + 1$, and for all other vertices x, $id(x) = od(x)$.

Proof: Exercise 9. (Add or subtract a directed edge between them, form Eulerian digraph, modify.)

Theorem 4.6, p. 151: If a connected digraph D has two vertices u and v such that, for some positive integer k, $od(u) = id(u) + k$, $id(v) = od(v) + k$, and for all other vertices x, $id(x) = od(x)$, then D contains k edge-disjoint directed uv paths.

Proof: Correct the degree deficiencies by adjoining new vertices $w_1, \ldots, w_k$ and directing v into each of these and directing each of these into u. The result is Eulerian. Walk around the directed edges, finding k disjoint uv directed walks which can be truncated to paths.

Definition: A Hamilton Cycle in a graph is a cycle that passes through all n vertices.

Theorem 3.6, p. 119: If $n \geq 3$ and $\deg(u) + \deg(v) \geq n$ for all non-adjacent u, v then there is a cycle in the graph containing all the vertices (Hamilton Cycle).

Proof: First, the graph must be connected: otherwise we could partition the vertices into V_1 and V_2 with no edges between these sets, and the degree of each vertex in V_1 would be $\leq |V_1| - 1$ and the degree of each vertex in V_2 would be $\leq |V_2| - 1$ and we would violate the degree sum condition. Secondly, any maximal path can be used to create a cycle, either by joining the endpoints or using the degree-sum condition to do the same thing as in the proof of Exercise 15, Section 1.3, p. 42. Thirdly, if the cycle created by a maximal path did not incorporate all vertices then connectedness allows us to glue an edge to this cycle and open the thing up to create a longer path: contradiction. So there is a Hamilton Cycle.

Corollary 4.10, p. 153: If D is a digraph in which $od(v) \geq n/2$ and $id(v) \geq n/2$ then there's a directed Hamilton Cycle. (Of course, this implies the graph is strongly connected.)

Proof: For $n = 2$ and $n = 3$ there has to be a directed edge between every pair of vertices. Now consider $n \geq 4$. Then every out-degree is ≥ 2. Let Q be a path of maximal length. Its endpoint has directed edges back to the path, and since there are at least $n/2$ of these, we can find a cycle of length at least $n/2 + 1$ (we did this before in Theorem 1.24, p. 50). Now let C be a cycle of maximal length. Then it has at least $n/2 + 1$ edges. We would like to argue that it has n edges. Say that it has $< n$ edges. We will obtain a contradiction.

There has to be a vertex not on the cycle. Let P be a path of maximal length l edges ($l+1$ vertices) not intersecting C. Assume P begins with u and ends with v. Since u has in-degree $\geq n/2$ but at most $n/2 - 2$ of these edges originate from P, there are $a \geq 2$ vertices on C with an edge to u. Similarly, since v has out-degree $\geq n/2$, there are $b \geq 2$ vertices on C on the other end of edges directed out of v. Whenever there is an edge $c_1 \to u$ where $c_1 \in C$ there must be a gap of at least $l + 1$ vertices along C to an edge of the form $v \to c_2$ where $c_2 \in C$, otherwise we can create a longer cycle by taking a detour along P. So for each $c \to u$ edge, the next vertex c' cannot be part of a $v \to c'$ edge. Scoop up $a - 1$ vertices c that are not part of a $v \to c$ edge, and scoop up $l + 1$ more of these between the last $c \to u$ edge and the first $v \to c'$ edge following it. We have found at least $a + l$ vertices along C not contributing to the out-degree of v.

Now we have some inequalities to manipulate: $id(u) \geq n/2$ with a contribution of a edges into u from C leaves at least $n/2-a$ edges into u from P (there are no other vertices with edges into u by maximality of P). Since P contains $l+1$ vertices, this implies $l \geq n/2-a$. Hence $a+l \geq n/2$. On the other hand, we have identified $a+l$ vertices on C that do not contribute to the outdegree of v, therefore $n/2 \leq od(v) \leq n - 1 - (a + l) \leq n - 1 - (n/2) = n/2 - 1$. Contradiction. Therefore C has n edges in it and is a directed Hamilton Cycle.

Skip the material on Line Digraphs.

Section 4.3: Flows in Networks

Network: directed graph with source vertex having in-degree 0, target vertex having out-degree 0, non-negative integer edge capacities.

Network flow: an edge function which respects capacities and satisfied conservation of flow: $0 \leq f(xy) \leq cap(xy)$ for each edge xy and, for all

$v \in V \backslash \{S, T\}$,

$$\sum_{xv \in E} f(xv) = \sum_{vx \in E} f(vx).$$

Network cut: an ordered pair (P, Q) satisfying $P, Q \subseteq V$, $S \in P$, $T \in Q$, $P \cap Q = \emptyset$.

Cut capacity:

$$cap(P, Q) = \sum_{(p,q) \in (P \times Q) \cap E} cap(pq).$$

Flow across a cut:

$$flow(P, Q) = \sum_{(p,q) \in (P \times Q) \cap E} f(pq).$$

Flow from one subset to another:

$$flow(X, Y) = \sum_{(x,y) \in (X \times Y) \cap E} f(xy).$$

Value of a flow: $val(f) = flow(\{S\}, V)$.

The basic problem: find a flow of maximum value.

Lemma: For any flow f and for any cut (P, Q),

$$val(f) = flow(P, Q) - flow(Q, P).$$

Corollary: for any flow f and for any cut (P, Q), $val(f) \leq cap(P, Q)$.

Theorem: when $val(f) = cap(P, Q)$ occurs, $val(f)$ is maximal among all flows and $cap(P, Q)$ is minimal among all cuts.

Proof: Let f' be given. $val(f') \leq cap(P, Q) = val(f)$. Let (P', Q') be given. $cap(P, Q) = val(f) \leq cap(P', Q')$.

Theorem: Given any network, it is always possible to find a flow and a cut satisfying $val(f) = cap(P, Q)$.

Proof: First, what would this look like? On the one hand, $val(f) = cap(P,Q)$, and on the other hand $val(f) = flow(P,Q) - flow(Q,P)$. Since $flow(P,Q) \leq cap(P,Q)$ and $flow(Q,P) \geq 0$, we must have $flow(P,Q) = cap(P,Q)$ and $flow(Q,P) = 0$.

Idea of the construction: start with any arbitrary flow f, then pump it up as far as possible.

Method: Inspect any any semi-path from S to T, regardless of edge-directions. Measure slack as the ability to increase flow up to capacity in the forward direction and the ability to decrease flow down to zero in the backward direction. Find the minimum slack value and adjust all edges accordingly. This pumps up the flow value and preserves conservation of flow.

A procedure for finding a path, which also detects when the flow is maximal: label the vertex S with $*$. At present we have a partition of the vertex set into labeled vertices P and unlabeled vertices Q. Given an edge pq with slack $s > 0$, we label q with $(p, +, s)$. Given an edge qp with slack $s > 0$, we label q with $(p, -, s)$. Continue labeling so long as positive slack is detected across edges between P and Q in either direction. There are two possible outcomes.

1. T ends up with a label. We can backtrack through the labels all the way to S and identify a semi-path along which to pump up the flow. Make it so.

2. T does not end up with a label. Consider what this says about (P,Q).

(a) $S \in T$.

(b) $T \in Q$.

(c) $P \cup Q = V$.

(d) $P \cap Q = \emptyset$.

(e) For any pq edge, slack is zero, hence $flow(pq) = cap(pq)$.

(f) For any qp edge, slack is zero, hence $flow(qp) = 0$.

(g) $val(f) = flow(P,Q) - flow(Q,P) = cap(P,Q) - 0$. Hence flow is maximal.

Turning the crank means label and halt with either 1 or 2. Every time we halt with 1, we turn the crank again. We can only turn the crank a finite number of times and end with 1, because the flow can only increase so much. Eventually we end with 2.

Theorem 4.37: Let u and v be distinct vertices in a digraph D. Then the maximum number of edge-disjoint st paths is equal to the minimum number of edges in an uv separating set.

Proof: Set p_{uv} equal to the maximum number of edge-disjoint uv paths and set s_{uv} equal to the minimum number of edges in a uv-separating set. Let P be a collection of p_{uv} edge-disjoint uv paths. Obviously no uv-separating set can contain less than $|P|$ edges, so $s_{uv} \geq p_{uv}$. We need to prove somehow that $s_{uv} \leq p_{uv}$.

Create a network N from D by adding an edge Su and an edge vT. Assign $c(e) = 1$ to all edges in D, and infinite capacity to Su and vT. By conservation of flow, every flow f in this network gives rise to an edge-disjoint collection of $val(f)$ uv paths. Now find a flow f and a cut (P, Q) in the network satisfying $val(f) = cap(P, Q)$. Since there is infinite capacity along Su, $u \in P$. Since there is infinite capacity along vT and $T \notin P$, $v \notin P$. So every uv path starts in P, ends in Q, and incorporates an edge in $(P \times Q) \cap E(D)$. This makes $(P \times Q) \cap E(D)$ a uv-separating set in D. So

$$s_{uv} \leq |(P \times Q) \cap E(D)| = cap(P, Q) = val(f) \leq p_{uv}.$$

Theorem 4.38: Let u and v be distinct vertices in a diagraph D such that $uv \notin E(D)$. Then the maximum number of internally disjoint uv paths in D is equal to the minimum number of vertices in a uv separating set.

Proof: Set p_{uv} equal to the maximum number of internally disjoint uv paths and set s_{uv} equal to the minimum number of edges in a uv-separating set. We have to remove at least one vertex from each path in a collection of internally-disjoint uv paths to separate u from v, so $p_{uv} \leq s_{uv}$. We cannot use the same network idea as before, however, because flows give rise to edge disjoint paths that are allowed to share vertices.

The fix is to create create a new digraph D', inserting a new vertex x' (bottleneck vertex) for every original vertex, adding the new edge $x'x$, and replacing every original edge xy by the edge xy'. Now create the network N by adding Su', vT, using unit capacities everywhere except for Su', uu', $v'v$, and vT. Every flow corresponds to a collection of internally-disjoint uv paths in D', hence in D. So as before we can find a cut (P, Q) such that

$$|(P \times Q) \cap E(D')| = cap(P, Q) = val(f) \leq p_{uv}.$$

Each edge in $(P \times Q) \cap E(D')$ is of the form xy' or $x'y$. Since the edges are internally disjoint, the collection of unprimed vertices they contain is distinct. Removing these vertices breaks all paths between u and v. Hence we obtain a uv-separating set of edges of size $|(P \times Q) \cap E(D')|$, which yields

$$s_{uv} \leq |(P \times Q) \cap E(D')| = cap(P, Q) = val(f) \leq p_{uv}.$$

Skip Chapter 5: running short on time.

Chapter 6: Planar Graphs

Section 6.1: The Euler Identity

Planar graph: Can be drawn in plane with no crossing edges.

Planar graphs define regions. It turns out that all possible planar representations of a connected graph create the same number of regions.

Theorem 6.1, p. 225: For a connected planar graph, $r = e - v + 2$.

Proof: If the graph is a tree, there is one region, which satisfies the formula. If the graph is not a tree, there is a cycle somewhere. Find a cycle of minimal size. There will be no chords. Stuff everything inside this cycle (imagine drawing the image on a rubber ball, then puncturing the ball inside the cycle then tearing and spreading out flat. Things will deform but the number of regions will remain the same). So without loss of generality the planar representation can be described as a big cycle with stuff inside. Delete one of the edges of the outer cycle. In the process, lose one interior region and lose one edge. Keep on going until you obtain a tree. We have $r' = e' - v + 2$ in the tree, hence $r = e - v + 2$.

Theorem 6.3, p. 226: If a connected planar graph contains a cycle then $e \leq 3v - 6$.

Proof: Stuff everything inside the cycle as before. Now place 2 dots about every edge near the middle. The dots are partitioned by region, and each region is bounded by a cycle. So each region contributes at least 3 dots. So the number of dots is $\geq 3r$. On the other hand, the number of dots is $2e$, therefore $2e \geq 3r$. That is, $2e \geq 3e - 3v + 6$. This yields $e \leq 3v - 6$.

Note: This theorem can be refined to incorporated the minimum size k of a cycle. This yields $(k - 2)e \leq kv - 6$. In particular, for bipartite graphs, $2e \leq 4v - 6$ so $e \leq 2v - 4$.

K_{33} and K_5 are non-planar: they have too many edges using the previous theorems.

Corollary 6.6, p. 227: Every planar graph contains a vertex of degree ≤ 5.

Proof: Too many high-degree vertices implies too many edges to be planar.

Maximal Planar Graph: The addition of any edge creates a non-planar graph. These objects must be connected. A tree cannot be maximal planar. Hence a maximal planar graph can be stuffed inside a cycle, and the cycle must have exactly 3 edges to it. Also, the interior regions are bounded by 3-cycles. This implies $2e = 3r$, i.e. $e = 3v - 6$.

Nearly Maximal Planar: Can be stuffed inside a cycle. All interior cycles are 3-cycles. The outer one doesn't have to be a 3-cycle.

Theorem 6.8, p. 228: Maximal planar graphs with $v \geq 4$ satisfy $\delta \geq 3$.

Proof: This follows from the fact that maximal planar graphs are 3-connected (Exercise 6).

Skip pp. 229 – 232.

The Five Regular Polyhedra: See the description on page 221 and the figures on page 223. These can be projected to the surface of a sphere, which creates a connected plane graph with all vertex degrees the same and all regions having the same number of edges. To prove that there are no others, note that if degrees are d and bounding cycle lengths are L then $2e = vd = rL$ by a dot counting argument. Combined with $r - e + v = 2$ we obtain

$$2\frac{e}{L} - e + 2\frac{e}{d} = 2$$

$$\frac{2}{L} - 1 + \frac{2}{d} = \frac{2}{e}$$

$$\frac{2}{L} - 1 + \frac{2}{d} > 0$$

$$\frac{2}{L} + \frac{2}{d} > 1$$

$$2d + 2L > dL$$

$$(d-2)(L-2) = dL - 2d - 2L + 4 < 4.$$

Since $d \geq 3$ and $L \geq 3$ the only solutions to this inequality are

$$(d, L) \in \{(3,3), (3,4), (3,5), (4,3), (5,3)\}.$$

Section 6.2: Planarity versus Nonplanarity

Kuratowski subgraph of a graph: A subgraph which can be described as subdivision of K_5 or $K_{3,3}$ (interrupt edges by degree 2 vertices).

Petersen Graph: Satisfies $e \leq 3v - 6$ but not $(k-2)e \leq kv - 6$ using $k = 5$, hence non-planar. Circle-chord method yields a K_{33} configuration: see illustration.

Branch and subdivision vertices in a Kuratowski subgraph: Branch vertices are the original vertices of K_5 and K_{33}. Subdivision vertices are the inserted vertices of degree 2.

Minimal non-planar graph: A non-planar graph such that every proper subgraph is planar.

X-lobe of a graph G: Let X be subset of vertices of G and let G_i be a connected component of $G - X$. An X-lobe of G is the induced subgraph $G[G_i + X]$.

Edge contraction: Let G be a graph and let $e = \{x, y\}$ be an edge in G. $G \cdot e$ is the graph obtained from G by shrinking the edge xy down to a point z. In the process we lose the vertices x and y, gain the vertex (xy), and any edge xv or yv in G becomes the edge $(xy)v$ in $G \cdot e$.

Example: The Petersen graph can be contracted down to K_5. See illustration. If G is planar then $G \cdot e$ is planar. So if some contraction of G is nonplanar then G is nonplanar. Hence P is nonplanar.

Lemma 1: Let C be a cycle of a planar graph G. Then there is a way to draw G so that the edges of C all border the infinite region.

Proof: Stuff everything inside as before.

Lemma 2: Every minimal non-planar graph G (all proper subgraphs planar) is 2-connected.

Proof: First note that any minimal non-planar graph must be connected. We must show that there are no cut-vertices.

Suppose x is a cut vertex. Let $G - x$ have components $H_1, \dots, H_k$. Each lobe $G[H_i + x]$ is planar by minimality of G. Each non-tree among these can be redrawn so that x is bordering the infinite region. Each tree among them has x bordering the infinite region. These can be glued together to form a planar representation of G. Contradiction. So there are no cut vertices.

Lemma 3: Suppose $G - x - y$ is not connected. If G is non-planar then adding the edge xy to some $\{x, y\}$-lobe of G yields a non-planar graph.

Proof: Let the components of $G - x - y$ be $H_1, \dots, H_k$. Suppose every $G[H_i + x + y] + xy$ is planar. Draw each such configuration so that the edge xy borders the infinite region. There is a way to glue all these things together to create a planar graph, and this includes G as a subgraph. Contradiction. So some $G[H_i + x + y] + xy$ is non-planar.

Lemma 4: If there exists a minimal example G of a non-planar graph with no Kuratowski subgraph, then it is 3-connected.

Proof: G is a minimal non-planar graph. By Lemma 2 is G is 2-connected. If G is not 3-connected then it has a minimal vertex cut $\{x, y\}$. Let the connected components of $G-x-y$ be $H_1, \dots, H_k$. Then wlog $G[H_1+x+y]+xy$ is non-planar. Moreover, $G[H_1 + x + y] + xy$ has fewer elements than G, therefore by minimality of G it contains a Kuratowski subgraph K. Since K cannot be a subgraph of G, $xy \in K$ and $xy \notin G$. Since $\{x, y\}$ is a minimal vertex cut, there is an edge from x to H_2 and from H_2 to y. Hence there is a path P from x to y which is internally disjoint to $G[H_1 + x + y]$ and hence to K. We can replace the edge xy in K by the path P to create a Kuratowski subgraph $K - xy + P$ in G. Contradiction. Therefore G is 3-connected.

Lemma 5: Every 3-connected graph G with at least 5 vertices has an edge e such that $G \cdot e$ is also 3-connected.

Proof: Suppose this edge cannot be found. Let G be 3-connected and let $e = xy$ be an edge in G. Since $G \cdot e$ is not 3-connected, it has a vertex cut $\{u, v\}$. We claim that $(xy) = u$ or $(xy) = v$.

Suppose in fact $(xy) \neq u$ and $(xy) \neq v$. Then there must be some vertex in $(G \cdot e) - u - v$ with no path to (xy). Since G has $v \geq 5$ vertices, there must be some vertex in $G - u - v$ with no path to x and no path to y. Contradiction.

So now we know that $G \cdot e$ has a separating set of the form $\{(xy), z\}$. This creates a separating set $\{x, y, z\}$. Of all ways to choose the edge $e = xy$, choose one which maximizes the vertices in the largest connected component of $G-x-y-z$. Since $\{x,y,z\}$ is a minimal vertex cut of G, x, y and z have edges to each component of $G-x-y-z$. So if H is the largest component of $G-x-y-z$ and H' is another component of $G-x-y-z$, we get the diagram on page 249 of Douglas West's textbook.

Let u be a neighbor of z in H'. Let v be such that G has a separating set $\{z,u,v\}$. To achieve a contradiction we will find a connected component of $G-z-u-v$ that is larger than H.

First note that $G[H+x+y-v]$ is connected: Consider the cases.

Case 1: $v=x$. Then $G[H+y]$ is connected.

Case 2: $v=y$. Then $G[H+x]$ is connected.

Case 3: $v \in H$. We know that $G-z-v$ is connected. Given two vertices in $G[H+x+y-v]$, find a path between them in $G-z-v$ and shrink it to a path in $G[H+x+y-v]$.

Now that we know that $G[H+x+y-v]$ is connected, it has to belong to a connected component H'' of $G-z-u-v$ which has at least as many vertices as $G[H+x+y-v]$ and strictly greater vertices than H. Contradiction. So yes, we can find $e \in G$ such that $G \cdot e$ is 3-connected.

Lemma 6: If $G \cdot e$ has a Kuratowski subgraph then so does G.

Proof: Let K be a Kuratowski subgraph in $G \cdot e$. Write $e = xy$. If $(xy) \notin K$ then K is a Kuratowski subgraph of G. Now suppose $(xy) \in K$. If (xy) is a subdivision vertex of K, let the edges it belongs to be $u(xy)$ and $(xy)v$. By considering the possibilities in G we can see that G has a Kuratowski subgraph. If (xy) is a branch vertex of K and exactly one of the edges $(xy)u_i$ in K corresponds to xu_i in G and the rest correspond to yu_i in G then x is a subdivision vertex of a Kuratowski subgraph in G (or y if the roles of x and y are reversed). The only remaining case is when K is a subdivision of K_5 and the four edges $(xy)u_1$, $(xy)u_2$, $(xy)u_3$,$(xy)u_4$ in K correspond to xu_1, xu_2, yu_3, yu_4 in G. Writing the branch vertices of K as (xy), v_1, v_2, v_3, v_4, there are paths joining each v_i to v_j as well as paths from x to v_1 and v_2 and paths from y to v_3 and v_3, as well as the edge xy. Tossing the v_1v_2 path and the v_3v_4 path, we obtain a subdivision of K_{33} out of the remaining paths, with branch vertices x, v_3, v_4 on the right and branch vertices y, v_1, v_2 on the left.

Lemma 7: If G does not have a Kuratowski subgraph and $G \cdot e$ is 3-connected and planar, then G is planar.

Proof: We know by Lemma 6 that $G \cdot e$ does not have a Kuratowski subgraph. Now draw a planar representation of $G \cdot e$. Removing (xy), the remaining graph is 2-connected. Therefore (xy) and the edges to its neighbors in $G \cdot e$ are bounded by a cycle C. The vertices of C belong to G. The neighbors of x and y in G belong to C. Let the neighbors of x be $x_1, \ldots, x_j$ in cyclic order around C and let the neighbors of y be $y_1 \ldots, y_k$ in cyclic order around C. Note that there could be some overlap among these sets of neighbors. It is clear that $G - x$ and $G - y$ are planar since they are isomorphic to subgraphs of $G \cdot e$, and if x has ≤ 1 neighbors then we can insert y and its edges to $x, y_1, \ldots, y_k$ to create a planar representation of G. Now assume that x has at least 2 neighbors in C. We will consider the ways $y_1, \ldots, y_k$ can be distributed around C.

Case 1: y has at least three neighbors z_1, z_2, z_3 in common with x. Using C we can create a K_5 subdivision in G. Contradiction. So Case 1 cannot happen.

Case 2: y shares at most two neighbors in common with x and the rest of neighbors of y all fall between two consecutive neighbors x_i, x_{i+1} of x. Then we can insert y in the triangle formed by x, x_i, and x_{i+1} and draw all the edges out of y to create a planar representation of G.

Case 3: y shares at most two neighbors in common with x but the rest of the neighbors of y do not fall between two consecutive neighbors x_i, x_{i+1} of x. In other words, y has neighbors z_1 and z_2 that alternate with neighbors x_i and x_{i+1} of x. Using C we can create a $K_{3,3}$ subdivision in G. Contradiction. So Case 3 cannot happen.

Theorem: Every graph that does not have a Kuratowski subgraph is planar.

Proof: If the theorem is false, then there is a minimal counterexample, G. G is non-planar, does not have a Kuratowski subgraph, and by Lemma 4 G is 3-connected. Since K_4 and its subgraphs are planar, G must have at least 5 vertices. By Lemma 5, G has an edge e such that $G \cdot e$ is 3-connected. By Lemma 6, $G \cdot e$ does not have a Kuratowski subgraph. By minimality of G, $G \cdot e$ must be planar. By Lemma 7, G must be planar. Contradiction. So the theorem is true.

Chapter 8: Vertex Colorings

A proper k-coloring of a graph is a function $f : V \to \{c_1, \ldots, c_k\}$ such that the endpoints of each edge receive different colors.

The chromatic number of a graph, $\chi(G)$, is the smallest number of colors needed to properly color the graph.

Some facts about $\chi(G)$: $\chi(K_n) = n$. If $H \subseteq G$ then $\chi(G) \geq \chi(H)$. A non-trivial graph G is bipartite iff $\chi(G) = 2$. $\chi(C_{2k}) = 2$ and $\chi(C_{2k+1}) = 3$. $\chi(G + H) = \max(\chi(G), \chi(H))$. $\chi(G \vee H) = \chi(G) + \chi(H)$. $\chi(G \Box H) = \max(\chi(G), \chi(H))$ using $c(g, h) = c(g) + c(h)$ mod the larger chromatic number.

Theorem 8.9, page 315: Let G be a graph with n vertices and independence number α (maximum number of mutually non-adjacent vertices in G). Then

$$n/\alpha \leq \chi \leq n - \alpha + 1.$$

Proof: Given a proper χ-coloring, there are at most α vertices colored i. So there are at most $\chi\alpha$ vertices. $n \leq \chi\alpha$ implies $n/\alpha \leq \chi$. Now find α independent vertices and color them 1. This leaves $n - \alpha$ vertices to color, and we can assign then $n - \alpha$ distinct colors other than 1. Hence $\chi \leq 1 + n - \alpha$.

Application: let's say committees 1 through n are meeting at the Joint Meetings in 50 minute time slots. We want to rent the fewest number of hotel conference rooms by the hour to accommodate all the committees. Let G be the graph with vertex set $1, 2, \ldots, n$ and an edge between vertices $i \neq j$ if they cannot meet at the same time due to shared membership. Then an assignment of suitable time slots to the committees is a proper coloring, and the fewest number of time slots needed is $\chi(G)$.

Section 8.3: Bounds for the Chromatic Number.

Greedy algorithm for coloring a graph: Set $c(v_1) = 1$. Having colored v_1 through v_i, assign v_{i+1} the smallest positive integer not assigned to any of its neighbors among $v_1, \ldots, v_i$.

Theorem 8.19, p. 330: The greedy coloring algorithm yields a proper coloring, and the colors used are in the range $\{1, 2, \ldots, \Delta + 1\}$.

Proof: We will show that vertices v_1 through v_i have no edge to a vertex of the same color and each has a color $\leq \Delta + 1$ by induction on i.

First, consider v_1. It receives the color 1. Suppose it has an edge to vertex v_i. Then v_i does not receive the color 1. Hence the base case is true.

Now assume that none of the vertices v_1 through v_i have an edge to a vertex of the same color, and assume that all these vertices are colored with a positive integer $\leq \Delta+1$. Consider the vertex v_{i+1}. Then it does not receive the color of any of its neighbors among v_1 through v_i, and since the colors of each of these are in $[1, \Delta+1]$ but there are fewer than $\Delta+1$ of these, it receives a color in this range. If v_{i+1} has an edge to v_j for some $j > i+1$, then by construction v_j does not receive the same color as v_{i+1}. Hence vertices v_1 through v_{i+1} have the desired properties.

Note that we can strengthen this to say that if every vertex has at most p lower-index neighbors then the greedy algorithm yields a proper coloring with at most $p+1$ colors.

Theorem 8.22 (Brook's Theorem), p. 331: If G is connected and is not a complete graph and not an odd cycle then $\chi(G) \leq \Delta$.

Proof: First suppose that G has a vertex of degree $\leq \Delta-1$. Call this vertex v_n. It has at most $\Delta-1$ lower index neighbors. Now number its neighbors arbitrarily using the highest possible indices. Since these have at least one higher index neighbor, they can have at most $\Delta-1$ lower index neighbors. Keep on going. Now use Greedy Algorithm.

Now suppose that G is Δ regular. We can assume $\Delta \geq 3$, because $\Delta = 2$ implies $G = C_{2n}$, which is 2-colorable. If it has a cut-vertex x, each x-lobe has a vertex of degree $\leq \Delta-1$ and can be Δ-colored. Permute the colors in each component so that x is colored 1 in each, then glue together.

Now assume G is Δ regular and has no cut vertices. Suppose we can find a vertex x with two non-adjacent neighbors y and z so that $G-y-z$ is connected. Then it is possible to label the vertices $v_1, \ldots, v_n$ in such a way that each v_i up to v_{n-1} has at most $\Delta-1$ lower-indexed neighbors: Label x with v_n and label all the vertices except y and z using the vertex names v_3 through v_n as in the first paragraph. Now set $y = v_1$, $z = v_2$. Use greedy coloring on the v_1 through v_{n-1}. Then v_1 and v_2 both get colored 1, so v_n, while it has Δ neighbors, is adjacent to at most $\Delta-1$ colors. So v_n can be colored with a color in the range $\{1, 2, \ldots, \Delta\}$.

Finally, we have to show that the vertices x, y, z can be found. We have assumed that $\kappa(G) \geq 2$. Choose an arbitrary vertex a. Then $\kappa(G-a) \geq 1$. We will consider two cases.

Case 1: $\kappa(G-a) \geq 2$. Since G is regular and not complete, a is not adjacent to some vertex. This implies $e(a) \geq 2$. Say $d(a, A) \geq 2$. Let $P = a, b, c, \ldots, A$

be an a, A path of minimal length. Then a and c cannot be adjacent, otherwise there is a shorter a, A path. We can set $x = b$, $y = a$, $z = c$.

Case 2: $\kappa(G - a) = 1$. Write $G' = G - a$. Then G' has a cut vertex, so at least two blocks, so at least two end-blocks. Consider the path between them in the cut-block graph: $B_1, x_{12}, B_2, x_{23}, B_3, \dots, x_{n-1,n}, B_n$. Deleting x_{12} from G' severs all connections between B_1 and the rest of G', so there must be an edge from $b_1 \in B_1 - x_{12}$ to a since G has no cut vertices. Similarly, there must be an edge from $b_n \in B_n - x_{n-1,n}$ to a. The vertices b_1 and b_n are not adjacent, otherwise the edge between them would live entirely in one block but have endpoints in two blocks. We can set $x = a$, $y = b_1$, $z = b_n$. This works because $G - a - b_1 - b_n$ is connected and a has third edge which joins it to $G - a - b_1 - b_n$.

Theorem 8.27 (The Gallai-Roy-Vitaver Theorem), p.337: If D is an orientation of G then $\chi(G) \leq L + 1$ where L is the longest directed path length in D. Equality is possible for some D.

Proof: Given an orientation D, let D' be a maximal acyclic sub-digraph with respect to number of elements. D' contains all vertices of D. For each vertex u let $c(u)$ = length of longest directed path in D' ending in u. Let $\{u, v\}$ be an edge in G. We must show $c(u) \neq c(v)$. If $u \to v$ belongs to D' then $c(u) < c(v)$. If $u \to v$ doesn't belong to D' then it's addition to D' creates a directed cycle, hence a path from v to u in D'. Hence $c(v) < c(u)$. So we have a proper coloring using colors in the range $\{0, 1, \dots, L\}$. This implies $\chi(G) \leq L + 1$.

Given a minimal coloring, orient the edges from smaller to larger color. Along a path of length L edges one encounters $1+L$ different colors, so $\chi(G) \geq 1+L$. Combined with $\chi(G) \leq 1 + L$ we have $\chi(G) = 1 + L$.

Coloring an infinite graph: Let G be a graph with vertex set $\{v_0, v_1, v_2, \dots\}$. Assume that every finite subgraph of G is k-colorable. Then G is k-colorable.

Proof: Note that if you can successfully k-color $G[v_0, v_1, \dots, v_n]$, you will not necessarily be able to k-color $G[v_0, v_1, \dots, v_{n+1}]$ by using the same color choices for $v_1, \dots, v_n$ and choosing a color for v_{n+1}: perhaps v_{n+1} is adjacent to each color. So we have to proceed carefully.

Fix a color c_0 for v_0. Let S_n be the set of finite sequences of the form $(c_0, c_1, \dots, c_n)$, where $G[v_0, v_1, \dots, v_n]$ can be properly k-colored with $c(v_i) =$

c_i for $0 \le i \le n$. Then $S_n \ne \emptyset$ for all n. For each $c \in S_n$ set

$$S(c) = \{d \in \bigcup_{m \ge n} S_m : (d_0, d_1, \ldots, d_n) = c\}.$$

As we have argued above, it is possible that $S(c) = \emptyset$ for a given c. Since

$$S(c_0) = \bigcup_{n \ge 0} S_n,$$

$S(c_0)$ is an infinite set. Since

$$S(c_0) = \{(c_0)\} \cup \bigcup_{i=1}^{k} S(c_0, i),$$

there exists $c_1 \le k$ such that $S(c_0, c_1)$ is an infinite set. Since

$$S(c_0, c_1) = \{(c_0, c_1)\} \cup \bigcup_{i=1}^{k} S(c_0, c_1, i),$$

there exists $c_2 \le k$ such that $S(c_0, c_1, c_2)$ is an infinite set. Keep on going. Color G so that $c(v_i) = c_i$ for each $i \ge 0$. This is a proper coloring: Consider the edge $\{v_m, v_n\}$ where $m < n$. This edge belongs to $G[v_0, v_1, \ldots, v_n]$, which is properly k-colored by $c(v_i) = c_i$ for $0 \le i \le n$, hence $c(v_m) \ne c(v_n)$.

Section 12.3: Ramsey Theory

2-coloring of K_n: an arbitrary assignment of the colors 0 and 1 to the edges of K_n. The color of an edge e is denoted $c(e)$.

Monochromatic clique of K_n: Let $\omega = (E, V)$ be a clique of K_n. We say that ω is i-chromatic if

$$|\{e \in E : c(e) = i| \le 1.$$

In this definition we allow 1-vertex cliques to be monochromatic since they have no edges.

Ramsey numbers: Let $s \ge 1$ and $t \ge 1$ be integers. If there exists a positive integer n such that every 2-coloring of K_n contains a 0-chromatic s-clique or a 1-chromatic t-clique then we say $R(s,t) \le n$ and define $R(s,t)$ to be the minimum value of n such that $R(s,t) \le n$.

Example: We have $R(1, t) = 1$ for all $t \geq 1$.

Example: We also have $R(2, t) = t$. Reason: In any 2-coloring of K_t, either we can find an edge colored 0 or all edges have been colored 1. Hence $R(2, t) \leq t$. Also, the uniform 1-coloring of K_{t-1} contains neither a K_2 colored 0 nor a K_t colored 1. Hence $R(2, t) = t$.

Example: If $R(s, t) \leq n$ if and only if $R(t, s) \leq n$. Reason: Suppose every 2-coloring of K_n contains a 0-chromatic K_s or a 1-chromatic K_t. Let c be an arbitrary 2-coloring of K_n. Then $\overline{c}$ is a 2-coloring of K_n, and a 0-chromatic K_s using $\overline{c}$ yields a 1-chromatic K_s using c and a 1-chromatic K_t using $\overline{c}$ yields a 0-chromatic K_t using c.

Theorem 12.17, page 512: For every $s, t \geq 1$, $R(s, t)$ is computable and, for $s, t \geq 2$,

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

Proof: Suppose that $R(s-1, t)$ and $R(s, t-1)$ are computable. Consider an arbitrary 2-coloring of K_n where $n = R(s-1, t) + R(s, t-1)$. We will show that there is a 0-chromatic K_s or a 1-chromatic K_t in K_n. First, choose a vertex v and define

$$V_0 = \{w \in V \backslash \{v\} : c(wv) = 0\}$$

and

$$V_1 = \{w \in V \backslash \{v\} : c(wv) = 1\}.$$

Then $|V_0| \geq R(s-1, t)$ or $|V_1| \geq R(s, t-1)$, because if both inequalities are violated then the degree of v is no greater than $n-2$ in K_n. Consider the possibilities:

Case 1: $|V_0| \geq R(s-1, t)$. Then there is a 0-chromatic K_{s-1} in $K_n[V_0]$, which yields a 0-chromatic K_s in K_n, or there is a 1-chromatic K_t in $K_n[V_0]$, which yields a 1-chromatic K_t in K_n.

Case 2: $|V_1| \geq R(s, t-1)$. Then there is a 0-chromatic K_s in $K_n[V_0]$, which yields a 0-chromatic K_s in K_n, or there is a 1-chromatic K_{t-1} in $K_n[V_0]$, which yields a 1-chromatic K_t in K_n.

Having established this, we start with the fact that $R(s, t)$ is computable when $s \leq 2$ or $t \leq 2$ and ratchet ourselves up to larger values.

Corollary 12.18, page 513: For every $s, t \geq 1$,

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

Proof: This is true when at least one parameter is ≤ 2. Moreover, assuming $R(s-1,t) \leq \binom{s+t-3}{s-2}$ and $R(s,t-1) \leq \binom{s+t-3}{t-2}$ we obtain

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

A counting problem: Let $n \geq t \geq 1$. How many 2-colorings of K_n contain a monochromatic K_t? Let the subsets of size t from $\{1, 2, \ldots, n\}$ be labeled $V_1, V_2, \ldots, V_k$ where $k = \binom{n}{t}$. Let C_i denote the set of all 2-colorings in which $K[V_i]$ is monochromatic. Then the total number of colorings is $|C_1 \cup C_2 \cup \cdots \cup C_k|$. The exact size of this union requires an inclusion-exclusion argument, but an upper bound is $|C_1| + |C_2| + \cdots + |C_k|$. To produce all the colorings in C_i, color all the edges in $K[V_i]$ the same color and assign arbitrary colors to the other $\binom{n}{2} - \binom{t}{2}$ edges. This yields

$$|C_i| = 2^{\binom{n}{2} - \binom{t}{2} + 1}.$$

Therefore

$$|C_1| + |C_2| + \cdots + |C_k| \geq \binom{n}{t} 2^{\binom{n}{2} - \binom{t}{2} + 1}.$$

Now the total number of 2-colorings of K_n is $2^{\binom{n}{2}}$, so if this number is larger than $\binom{n}{t} 2^{\binom{n}{2} - \binom{t}{2} + 1}$ then there exist 2-colorings of K_n which contain neither a 0-chromatic K_t nor a 1-chromatic K_t. We have proved that

$$2^{\binom{n}{2}} > \binom{n}{t} 2^{\binom{n}{2} - \binom{t}{2} + 1} \implies R(t,t) > n.$$

In other words,

$$2^{\binom{t}{2}} > 2\binom{n}{t} \implies R(t,t) > n.$$

To obtain $2^{\binom{t}{2}} > 2\binom{n}{t}$ it suffices to require $2^{\binom{t}{2}} > 2\frac{n^t}{t!}$. Setting $n = \lfloor c^t \rfloor$, we examine the growth of $f(t) = 2^{\binom{t}{2}}$ versus the growth of $g(t) = 2\frac{c^{t^2}}{t!}$. We have $f(t+1) = 2^t f(t)$, $g(t+1) = \frac{c^{2t+1}}{t+1} g(t)$. Setting $c = 2^{1/2}$ we obtain $\frac{c^{2t+1}}{t+1} = 2^t \frac{2^{\frac{1}{2}}}{t+1} < 2^t$ for $t \geq 1$. Hence $f(t)$ grows faster than $g(t)$ for $t \geq 1$, and one can check that $f(3) > \lfloor g(3) \rfloor$. We have proved

$$R(t,t) > \lfloor 2^{t/2} \rfloor \ (t \geq 3),$$

which is Theorem 12.20, page 515.

Graph Ramsey number: $R(G_0, G_1)$ is smallest n so that every 2-coloring of the edges of K_n contains an i-monochromatic G_i for some i.

Theorem: Let T_m be a tree with m vertices. Then $R(T_m, K_n) = (m-1)(n-1)+1$.

Proof: We first show $R(T_m, K_n) \leq (m-1)(n-1)+1$ by induction on n. The base case, $R(T_m, K_1) \leq 1$, is trivially true, since K_1 has no edges. Now assume $R(T_m, K_n) \leq (m-1)(n-1)+1$. Consider an arbitrary 2-coloring of $K_{(m-1)n+1}$. For each vertex v and define

$$V_0(v) = \{w \in V \backslash \{v\} : c(wv) = 0\}$$

and

$$V_1(v) = \{w \in V \backslash \{v\} : c(wv) = 1\}.$$

Consider the following two cases:

Case 1: $|V_1(v)| \geq (m-1)(n-1)+1$ for some v. By the induction hypothesis, within $G[V_1(v)]$ there is a T_m colored 0 or a K_{n-1} colored 1. The latter produces a K_n colored 1 in $K_{(m-1)n+1}$.

Case 2: $|V_1(v)| \leq (m-1)(n-1)$ for all v. This implies $|V_0(v)| \geq m-1$ for all v. In other words, in the subgraph of $K_{(m-1)n+1}$ consisting of all the vertices and just the edges colored 0, every vertex has degree $\geq m-1$. This must contain a 0-chromatic copy of T_m by a Theorem 2.20, page 67.

We now show that $R(T_m, K_n) > (m-1)(n-1)$. Color the edges of $n-1$ disjoint copies of K_{m-1} with the color 0. This yields no copies of T_m colored 0. Now join all of these vertices not already connected by an edge by an edge and color these 1. The largest clique colored 1 in this construction is K_{n-1}. Hence we have produced a 2-coloring of $K_{(m-1)(n-1)}$ with no 0-colored T_m and no 1-colored K_n.

A generalization of $R(s,t)$: coloring edges of K_n is really a matter of coloring 2-element subsets of $\{1, 2, \ldots, n\}$. To say that such a coloring produces a 0-colored K_s or a 1-colored K_t is to say that there one can find either $S \subseteq \{1, 2, \ldots, n\}$ with $|S| = s$ such that all 2-element subsets of S are colored 0 or we can find $T \subseteq \{1, 2, \ldots, n\}$ with $|T| = t$ such that all 2-element subsets of T are colored 1. We now say that $R(s, t; k) \leq n$ whenever every arbitrary coloring of the k-subsets of $\{1, 2, \ldots, n\}$ yields an s-set S with all k-subsets colored 0 or a t-set T with all k-subsets colored 1.

Theorem: $R(s, t; k)$ is computable.

Proof: Assuming that all $R(i, j; k-1)$ are computable and that $p = R(s-1, t; k)$ and $q = R(s, t-1; k)$ are computable, set

$$N = R(p, q; k-1).$$

We claim that $R(s, t; k) \leq N + 1$. To see this, let c be an arbitrary coloring of the k-subsets of $\{1, 2, \ldots, N+1\}$. Use c to define a coloring of the $(k-1)$-subsets of $\{1, 2, \ldots, N\}$ via

$$c'(S) = c(S \cup \{N+1\}).$$

There are two cases to consider.

Case 1: There is a p-set $P \subseteq \{1, 2, \ldots, N\}$ such that all $(k-1)$-sets of P are colored 0 by c'. So all k-sets of $P \cup \{N+1\}$ that contain $N+1$ are colored 0 by c. Given the size of P, there is a set $S \subseteq \{1, 2, \ldots, N\}$ of size $s-1$ such that all k-subsets of S are colored 0 by c or there is a set $T \subseteq \{1, 2, \ldots, N\}$ of size t such that all k-subsets of T are colored 1 by c. If S exists then $S \cup \{N+1\}$ is a set of size s such that all k-subsets of $S \cup \{N+1\}$ are colored 0.

Case 2: There is a q-set $Q \subseteq \{1, 2, \ldots, N\}$ such that all $(k-1)$-sets of Q are colored 1 by c'. So all k-subsets of $Q \cup \{N+1\}$ that contain $N+1$ are colored 1 by c. Given the size of Q, there is a set $S \subseteq \{1, 2, \ldots, N\}$ of size s such that all k-subsets of S are colored 0 by c or there is a set $T \subseteq \{1, 2, \ldots, N\}$ of size $t-1$ such that all k-subsets of T are colored 1 by c. If T exists then $T \cup \{N+1\}$ is a set of size t such that all k-subsets of $T \cup \{N+1\}$ are colored 1.

Given that all $R(i, j; 2)$ are computable by Theorem 12.17, page 512, each of the $R(s, t; k)$ are computable by a ratcheting-up argument.

Theorem (Erdos-Szekeres 1935): Given a sufficiently large number of non-collinear points in the plane, one can find the vertices of a convex polygon of m sides.

Proof: Let there be $R(m, 5; 4)$ points. Color every subset of 4 of them 0 if they form the vertices of a convex polygon and 1 if they do not. Then either there is a set S of size m such that every 4-set among them are colored 0 (i.e. form a convex polygon) or there is a set T of size 5 such that every 4-set among them are colored 1 (i.e. they don't form a convex polygon). Now T cannot exist, because given any 5 non-collinear points, 4 of them are bound to form a convex polygon. (Proof: take any 3 and form a triangle. If one of the other two are outside, use one to form a convex 4-gon. But if both are inside, use 2 inside and 2 along a side of the triangle to form a convex 4-gon.) Hence the set S exists. Now form the convex hull of these m points. If one of the points is not a vertex of the convex hull then it is an interior point. Triangulate the polygon by drawing chords. The interior points falls in one of these triangles. The three vertices of this triangle plus the interior point do not form a convex polygon, contrary to the defining property of S. Contradiction. So every point of S is a vertex of the convex hull of S, hence the points of S form the vertices of a convex polygon.