

Elementary Graph Theory

Robin Truax

March 2020

Contents

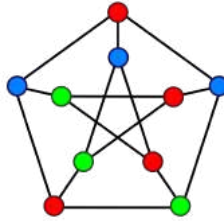
1	Basic Definitions	2
1.1	Specific Types of Graphs	2
1.2	Paths and Cycles	3
1.3	Trees and Forests	3
1.4	Directed Graphs and Route Planning	4
2	Finding Cycles and Trails	5
2.1	Eulerian Circuits and Eulerian Trails	5
2.2	Hamiltonian Cycles	6
3	Planar Graphs	6
3.1	Polyhedra and Projections	7
3.2	Platonic Solids	8
4	Ramsey Theory	9
4.1	Ramsey's Theorem	9
4.2	An Application of Ramsey's Theorem	10
4.3	Schur's Theorem and a Corollary	10
5	The Lindström-Gessel-Viennot Lemma	11
5.1	Proving the Lindström-Gessel-Viennot Lemma	11
5.2	An Application in Linear Algebra	12
5.3	An Application in Tiling	13
6	Coloring Graphs	13
6.1	The Five-Color Theorem	14
7	Extra Topics	15
7.1	Matchings	15

1 Basic Definitions

Definition 1 (Graphs). A *graph* G is a pair (V, E) where V is the set of *vertices* and E is a list of “edges” (undirected line segments) between pairs of (not necessarily distinct) vertices.

Definition 2 (Simple Graphs). A graph G is called a *simple graph* if there is at most one edge between any two vertices and if no edge starts and ends at the same vertex.

Below is an example of a very famous graph, called the Petersen graph, which happens to be simple:



Right now, our definitions have a key flaw: two graphs that have exactly the same setup, except one vertex is a quarter-inch to the left, are considered completely different. To solve this issue, we need an idea of “equivalence of graphs”.

Definition 3 (Graph Isomorphism). An *isomorphism of graphs* $G = (V_G, E_G)$ and $H = (V_H, E_H)$ is a bijection $\phi : V_G \rightarrow V_H$ such that (v, w) is in E_G if and only if $(\phi(v), \phi(w))$ is in E_H . If there is an isomorphism $\phi : G \rightarrow H$, we say G and H are *isomorphic graphs*, denoted $G \cong H$.

Essentially, an isomorphism of graphs is a structure-preserving map from the set of vertices of G to the set of vertices of H which is a one-to-one correspondence. Of course, where there is an isomorphism, there is a homomorphism.

Definition 4 (Graph Homomorphism and Subgraphs). A *homomorphism of graphs* $G = (V_G, E_G)$ and $H = (V_H, E_H)$ is a function $\phi : V_G \rightarrow V_H$ such that $(v, w) \in E_G \Rightarrow (\phi(v), \phi(w)) \in E_H$. In particular, if this is an injective function $G \rightarrow H$, then we say that H is a *subgraph* of G (since there is a structure-preserving embedding of H into G).

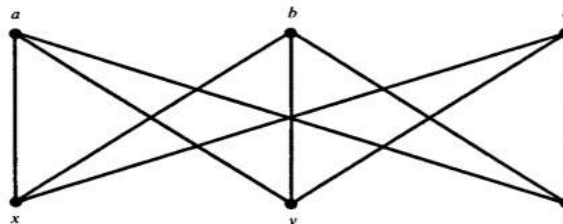
1.1 Specific Types of Graphs

Definition 5 (Null Graph). The graph G with v vertices and no edges is called the *null graph* on v vertices and denoted N_v . For example, a diagram of the graph N_0 is displayed to the right:

Definition 6 (Complete Graph). The graph G on v vertices with one edge between each distinct pair of vertices is called the *complete graph* on v vertices and denoted K_v .

Definition 7 (Bipartite Graph). A graph G is called a *bipartite graph* if one can partition the set of vertices V_G of G into two nonempty parts A and B such that every edge in G is between a vertex in A and a vertex in B .

For example, the “utility graph” (which we will explore more later) is a bipartite graph:



Definition 8 (Complete Bipartite Graph). The complete bipartite graph on $K_{v,w}$ is a graph consisting of:

1. A set A of v vertices,
2. a set B of w vertices,
3. and every possible edge starting from a vertex in A going to a vertex in B .

For example, the utility graph is $K_{3,3}$.

Definition 9 (Degree of a Vertex). The number of edges connected to a vertex a in a graph G is denoted by $\deg a$ and called the *degree of a* . If $\deg a$ is even, we call a an *even vertex*; otherwise we call a an *odd vertex*.

Proposition 1 (Handshaking Lemma). For any graph G with vertex set V_G and edge set E_G ,

$$\sum_{a \in V_G} \deg a = 2|E_G|.$$

In other words, the sum of the degrees of the vertices in G is twice the number of edges in G . By corollary, the sum of the degrees of the vertices in G is even and G has an even number of odd vertices.

Definition 10 (Regular Graphs). A graph is called *regular* if all its vertices have the same degree.

N_v , K_v , $K_{v,v}$ are all classic examples of regular graphs. We will explore some of the interesting regular-type criteria for graphs, such as complete regularity, in later sections.

1.2 Paths and Cycles

Definition 11 (Path). A “trip” in a graph G between vertices along edges is called a *path* if no vertex is visited more than once.

Definition 12 (Cycles). A trip in a graph G is called a *cycle* if its ending and starting vertex are the same, and no other vertex is used more than once.

Definition 13 (Trail). A trip in a graph G is called a *trail* if no edge is used more than once. A trail which returns to its starting point is called a *circular trail* or a *circuit*.

Problem 1. Formalize the definition of a “trip” in a graph G .

One of the most important criteria for graphs, which we will assume in numerous sections of these notes, is “connectedness”.

Definition 14 (Connectedness). A vertex a is *connected* to another vertex b if there exists a path connecting a and b . If every vertex in a graph G is connected to every other vertex in G , then we call the entire graph G *connected*. Finally, the *connected component of a* is the maximal connected subgraph of G containing a .

1.3 Trees and Forests

Definition 15. A *tree* is a connected graph with no cycles.

In particular, a tree T must be a simple graph and there must be a unique path connecting any distinct pair of vertices in T . Furthermore, any tree with n vertices must have $n - 1$ edges.

Definition 16 (Roots and Layers). Given a tree T , one may label a specific vertex a the *root* of the tree. In this case, a is said to be on *layer 0*. Vertices adjacent to a are said to be on layer 1, new vertices adjacent to those are said to be on layer 2, and so on.

Definition 17 (Terminal Edges and Terminal Vertices). After fixing a root, every branch in the graph must have a final *terminal edge* to a *terminal vertex* (or *leaf*).

Theorem 2. In a connected graph G with n vertices and m edges, the smallest number of edges that must be removed so that no cycles remain is exactly $m - n + 1$.

Proof. First show that any graph contains a tree containing all its vertices – then notice that this tree must have $n - 1$ edges, so we must have removed $m - (n - 1) = m - n + 1$. This number is denoted γ and called the *cycle rank* of G . \square

Problem 2. Prove that if there are n cities, with a function $c(ab)$ giving the cost of constructing a road between city a and city b , then the cheapest way to create a connected graph containing all n cities is given by the following “greedy algorithm”:

1. Connect the two cities with the cheapest connecting link.
2. Add the cheapest possible edge producing a tree with the edges already created.
3. Repeat step 2 until all cities are connected.

Problem 3. Prove that a connected graph G is a tree or consists of a single cycle with trees growing at its vertices if and only if there is an injective map from each edge in G to one of its endpoints.

Problem 4. Prove that a connected graph G which does not admit an injective map from vertices into edges is a tree.

Definition 18 (Forests). A *forest* is a set of disconnected trees. In particular, a forest with k trees and n total vertices has $n - k$ edges.

1.4 Directed Graphs and Route Planning

Definition 19 (Directed Graph). A *directed graph* (or *digraph*) is a graph wherein all edges are replaced by *arrows* (i.e. directed line segments).

Definition 20 (Mixed Graph). A *mixed graph* is a graph where some (but not necessarily all) edges are replaced by arrows.

In particular, by replacing each undirected edge with two arrows (one in each direction), we can reduce any mixed graph to a digraph. An example problem to do with mixed graphs is the “one-way traffic problem”: when can the edges of a graph G be given directions in such a way that there is a directed path from any vertex to any other?

An edge e between a and b is called a *bridge* if there is no way to get from a to b (or vice versa) except through e . Any “terminal edges” are bridges, since their tip (the outer vertex) is only reachable through said edge. Similarly, an edge which is not a bridge is called a *cycle edge*.

Problem 5. Using the above definitions, prove that if G is a connected graph, then we can always direct the cycle edges of G (leaving the bridges undirected) so that there is a directed path from any given vertex to any other.

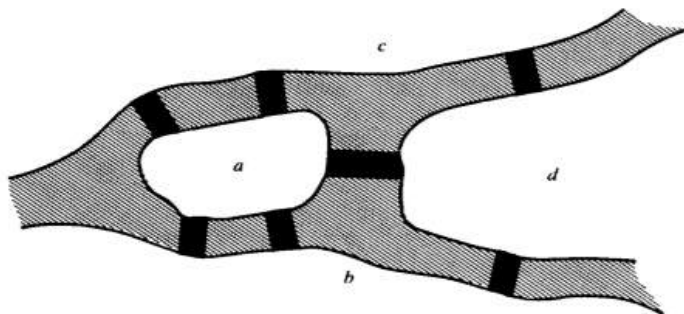
Problem 6. You are given a set V of vertices, a set E of weighted directed edges, a start vertex s and an end vertex t . Implement an algorithm to find the least difficult path from s to t (along with its difficulty), using the following method:

1. Give the start vertex the value 0.
2. For every vertex v adjacent to the startpoint, give v value equal to the minimal weight required to reach v . Keep track of the edge to v giving v its minimal value.

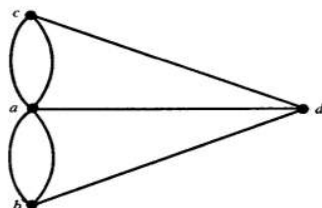
3. Iterate by assigning values to each vertex adjacent to a vertex covering in the last step. Whenever multiple values can be assigned to a vertex, choose the minimal one. Again, keep track of the edges to each vertex giving them their minimal value.
4. Eventually, one gives the end vertex t a value, and can trace back the path using the previous edges recorded.

2 Finding Cycles and Trails

Consider the following problem: is it possible to start on one landmass and use each of the Seven Bridges of Königsberg (as seen in the diagram below) exactly once, ending on the landmass where you started?



Dualizing the above map into a graph, this becomes equivalent to the problem of finding a circular trail passing through every edge in the following graph:



But clearly, to find such a cycle, every vertex must have even degree (as each time one enters the vertex, one must leave it, and vice versa). Since b , c , and d all have degree 3, no such cycle exists and the problem is unsolvable.

2.1 Eulerian Circuits and Eulerian Trails

Definition 21 (Eulerian Circuit). A circular trail passing through every edge once is called an *Eulerian circuit*. A graph which has an Eulerian circuit is called an *Eulerian graph*.

Theorem 3 (Eulerian Circuits). *All connected graphs with vertices of only even degree are Eulerian.*

Proof. Choose an arbitrary vertex a and create the longest possible trail T at a , always leaving a vertex from an edge which we have not used before. Since there are an even number of edges at every vertex, T must come to a halt at a . If T passes through all edges, we have found our Eulerian trail.

Otherwise, there is a vertex b lying on T with (an even number of) edges not traversed by T . We now begin a trail U at b , using only edges not in T . Again, this trail can only stop again at b . But then the trail following T to b , following U , and then following the rest of T back to a is strictly longer – a contradiction. Thus T uses all edges, so it is an Eulerian circuit, as desired. \square

Corollary 3.1. *A connected graph has a trail T_{ab} covering all edges just once (which we call an Eulerian trail) if and only if a and b are the only odd vertices.*

Proof. Connect a and b with an edge, create a Eulerian cycle, and then drop the edge ab . \square

Theorem 4. *A connected graph with $2k$ odd vertices contains a family of k distinct trails which, together, traverse all edges of the graph exactly once.*

Proof. Let the odd vertices in the graph be denoted by a_1, \dots, a_k and b_1, \dots, b_k . When we add the edges a_1b_1, \dots, a_kb_k to the graph, all the vertices become even and we get an Eulerian cycle T . When the edges are dropped out again, T breaks up into k separate trails, still covering all edges. \square

2.2 Hamiltonian Cycles

Definition 22 (Hamiltonian Cycles). A *Hamiltonian cycle* in a graph is a cycle that passes through each of the vertices exactly once. A graph with a Hamiltonian cycle is called a *Hamiltonian graph*.

Though it is extremely easy to determine if a graph is Eulerian, finding a Hamiltonian cycle or even showing that one exists is extremely difficult (we have found no efficient general criterion). Indeed, the problem of finding a Hamiltonian cycle is known to be *NP*-complete (that is, if it is possible to do it in polynomial time, then $P = NP$).

3 Planar Graphs

Definition 23 (Planar Graph). If a graph can be drawn on the plane in such a way that the edges have no intersections or common points except at the vertices, then we call it *planar*.

Problem 7. Prove that K_5 is not planar using the Jordan Curve Theorem.

Problem 8. Prove that $K_{3,3}$ is not planar using the Jordan Curve Theorem.

Theorem 5 (Euler's Formula). *For a connected planar graph G with V vertices, E edges, and F faces (including the infinite face), $V - E + F = 2$.*

Proof. Induct on E : start with 0 edges (which implies one vertex) – for which the theorem clearly holds – and add edges. Each time, you either add a new vertex with your edge or a new face with your edge, keeping the sum $V - E + F$ constant. \square

Theorem 6 (Fáry's Theorem). *Any simple planar graph can be drawn without crossings in such a way that all edges are straight line segments.*

Theorem 7 (Kuratowski-Wagner Theorem). *A subdivision of a graph is formed by subdividing its edges into paths of one or more edges. Then a finite graph is planar if and only if it contains no subgraph which is not a subdivision of K_5 or $K_{3,3}$.*

I encourage you to learn more about the proofs of the above theorems, but we will not discuss the proofs here – they are quite long for the minimal reward (at least with respect to what we are going to work on next).

Problem 9. Use the Kuratowski-Wagner Theorem to prove that the Petersen graph is not planar.

However, this is not an easy test to perform in general. Instead, usually one uses invariants of planar graphs to prove that a graph is *not* planar. For example, if G is a simple connected planar graph with v vertices and e edges, then $e \leq 3v - 6$.

Problem 10. Prove that $e \leq 3v - 6$ by adding edges to any planar graph until all faces are triangles, and then using Euler's Formula.

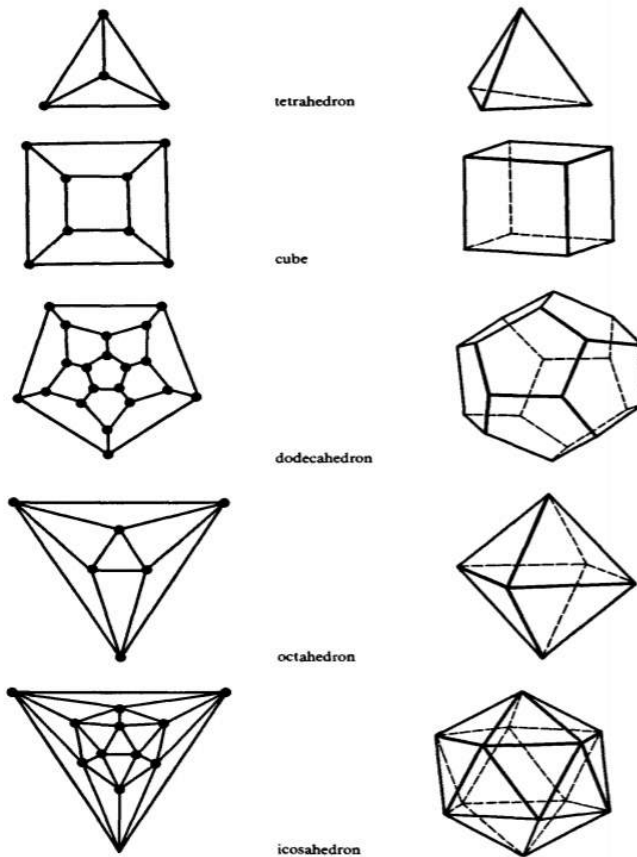
Also look at this link which shows an amazing and beautiful way to draw planar graphs in planar form using algebraic graph theory.

3.1 Polyhedra and Projections

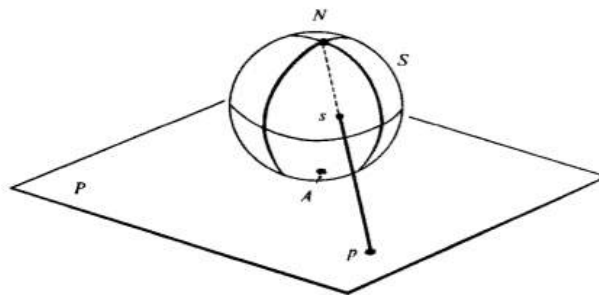
Definition 24 (Schlegel Diagrams). Suppose one has a polyhedron in 3-space above the xy -plane. Then choose a point above of the polyhedron as the projection point, and draw rays down through the polyhedron and onto the xy -plane. Vertices map to vertices, edges map to edges, and faces map to faces.

Corollary 7.1. For all convex polyhedra with V vertices, E edges, and F faces, $V - E + F = 2$.

Following are the Schlegel diagrams of the Platonic solids.



Projections are a very useful tool when analyzing planar graphs. One extremely interesting type of projection from the plane to the sphere is called “stereographic projection”, a common technique from projective geometry and Euclidean geometry (for example, it is the basis for the classic proof on finding all Pythagorean triples). Below is a diagram showing the process of stereographic projection:



Theorem 8. *Let v be a vertex of a connected planar graph G . Then G can be embedded in the plane in such a way that v is on the exterior face of the embedding.*

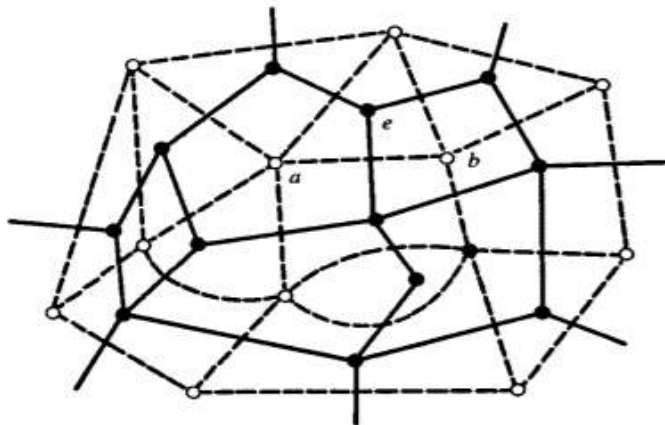
Proof. Let \overline{G} be a planar embedding of G . Stereographically project \overline{G} onto the sphere to get an embedding into the sphere called \overline{X} . Let f be a face that has v as one of its vertices: choose a point x in the interior of f and fix it. Rotate the sphere until x is the north pole, and project \overline{X} back down onto the plane. x becomes the point at infinity, so f is the exterior face, so v is on the exterior face of the new planar embedding. \square

3.2 Platonic Solids

Definition 25 (Dual Graph). For any polygonal graph G , we can construct a new polygonal graph G^* called its *dual graph* by the following method:

1. Within each face, including the infinite face, we define a single point.
2. Two such points are connected by an edge e if they correspond to neighboring faces (i.e. faces sharing an edge f). e is drawn such that it crosses f but no other edges of the graph.
3. If there are several boundary edges common to the two faces, one new edge is drawn for each boundary edge.

For example:



Notice that the two graphs have the same number of edges, but the number of vertices and faces is swapped. Also refer back to the figure of the Platonic solids' Schlegel diagrams and notice that the octahedron graph is dual to the cube graph, the icosahedron graph is dual to the dodecahedron graph, and the tetrahedron graph is self-dual.

Definition 26 (Completely Regular). A graph G is completely regular if both G and G^* are regular.

Theorem 9. *There are exactly five completely regular graphs.*

Proof. Suppose that G is regular of degree k and G^* is regular of degree k^* . Then $2E = k \cdot V$ and $2E = k^* \cdot F$. In particular, $E = \frac{kV}{2}$ and $F = \frac{kV}{k^*}$. From Euler's Formula, we see that:

$$n\left(1 + \frac{k}{k^*} - \frac{1}{2}k\right) = 2 \Rightarrow n(2k + 2k^* - kk^*) = 4k^*.$$

But since n and $4k^*$ are positive integers, then the expression in parentheses must also be positive. Thus $2k + 2k^* - kk^* > 0$ and $(k - 2)(k^* - 2) < 4$. The only pairs which work here are $(3, 3)$, $(3, 4)$, $(3, 5)$, $(4, 3)$, and $(5, 3)$. From here, it is not difficult to show that there is only one graph for each pair, as seen in the following table:

k	k^*	V	E	F
3	3	4	6	4
3	4	8	12	6
3	5	20	30	12
4	3	6	12	8
5	3	12	30	20

□

Notice that each completely regular planar graph corresponds to a unique Platonic solid: think about why this is the case.

4 Ramsey Theory

Proposition 10. *If there are six people at a party, each pair of which are either mutual acquaintances or mutual strangers, then there is a group of 3 people which are all mutual acquaintances or mutual strangers.*

Proof. Consider person P_1 and their five relationships. By the Pigeonhole Principle, at least three of these relationships, say P_2, P_3, P_4 are the same (without loss of generality we may assume they are friends). Then if P_2 and P_3 are friends, if P_3 and P_4 are friends, or if P_2 and P_4 are friends there is a group of 3 mutual friends. But if they are all strangers, this is a group of 3 mutual strangers, meeting our requirements anyways. □

Definition 27 (Ramsey Number, Preliminary). The Ramsey number $R(m, n)$ is the minimal v such that if every edge in K_v is colored red or blue, there is a subgraph K_m which is entirely red or a subgraph K_n which is entirely blue.

Definition 28 (Ramsey Number, General). The Ramsey number $R(n_1, n_2, \dots, n_k)$ is the minimal v such that if every edge in K_v is labelled using the number $\{1, \dots, k\}$, there is a subgraph K_{n_i} where all edges are labelled with j for some $j \in \{1, \dots, k\}$.

Problem 11. Prove that the Ramsey number $R(3, 3)$ is 6 by finding a coloring of K_5 with no monochromatic subgraph K_3 .

Problem 12. The Ramsey number $R(4, 3)$ is 9.

Problem 13. The Ramsey number $R(4, 4)$ is 18.

Problem 14. The Ramsey number $R(3, 3, 3)$ is 17.

4.1 Ramsey's Theorem

Theorem 11 (Ramsey's Theorem). *The Ramsey number $R(n_1, \dots, n_k)$ exists for any given positive integer k and positive integers n_1, \dots, n_k .*

Proof. We will first prove the case of $k = 2$ (two colors). To do this, we will use the following lemma:

Lemma 12. $R(r, s) \leq R(r - 1, s) + R(r, s - 1)$

Proof. Consider the complete graph on $R(r - 1, s) + R(r, s - 1)$ vertices whose edges are colored with red and blue. Choose an arbitrary vertex v . Then let M denote the set of vertices with red links to v and N denote the set of vertices with blue links to v .

Since the graph has $R(r - 1, s) + R(r, s - 1) = |M| + |N| + 1$ vertices, either $|M| \geq R(r - 1, s)$ or $|N| \geq R(r, s - 1)$. In the former case, if M has a red K_s then so does the original graph and we are done – if it does not, $M \cup \{v\}$ has a blue K_r . The other case follows analogously, and the result follows. □

By induction, the result follows for Ramsey numbers with $k = 2$.

Lemma 13. *If $k > 2$, then $R(n_1, \dots, n_k) \leq R(n_1, \dots, n_{k-2}, R(n_{k-1}), R(n_k))$.*

Proof. Consider the complete graph on $R(n_1, \dots, n_k) \leq R(n_1, \dots, n_{k-2}, R(n_{k-1}), R(n_k))$ vertices and label it with k numbers. Now pretend that $k - 1$ and k are the same number, so the graph is $(k - 1)$ -colored. Due to the definition of $R(n_1, \dots, n_2, R(n_{k-1}), R(n_k))$, either the graph contains a K_{n_i} numbered using just one number for $1 \leq i \leq k - 2$ or a $K_{R(n_{k-1}, n_k)}$ colored using $k - 1$ and k . In this case, then clearly there must be a $K_{n_{k-1}}$ or a K_{n_k} colored using just one color by definition, so we are done. \square

By our result on two colors, and induction, $R(n_1, \dots, n_k)$ is finite for any number of colors, so the Ramsey number exists. \square

4.2 An Application of Ramsey's Theorem

Proposition 14. *There is a function $f(m, n)$ with the following property: if x_1, \dots, x_N is any sequence of distinct real numbers with $N > f(m, n)$ then there is either a monotonic increasing sequence of length greater than m or monotonic decreasing sequence of length greater than n .*

Proof. Take $R(m + 1, n + 1) - 1$. Suppose that $N > f(m, n)$ and we are given a sequence of N distinct real numbers. Take $X = \{1, \dots, N\}$; then, for any distinct $i, j \in X$, color the edge x_i, x_j red if $x_i < x_j$ and blue if $x_j > x_i$. By definition of the Ramsey number, there is either a red $(m + 1)$ -set or a blue $(n + 1)$ -set - which indexes either a monotone increasing subsequence of length $m + 1$ or a monotone decreasing subsequence of length $n + 1$, respectively. \square

4.3 Schur's Theorem and a Corollary

Theorem 15 (Schur's Theorem). *Let r be a positive integer. Then there exists a positive integer $S(r)$ such that if $\{1, \dots, S(r)\}$ is colored with r colors, then there exist x, y, z of the same color satisfying $x + y = z$.*

Proof. Let

$$S(r) = R(\underbrace{3, \dots, 3}_{r \text{ 3's}})$$

Color $\{1, \dots, S(r)\}$ with r colors. Then, consider the graph $K_{S(r)}$, and color the edge ab with the color that $|a - b|$ has in the initial coloring of $\{1, \dots, S(r)\}$. This gives an r -coloring of $K_{S(r)}$, which must admit a monochromatic subgraph K_3 , which implies $|a - b|$, $|b - c|$, and $|a - c|$ are the same color. Assuming without loss of generality that $a > b > c$, we see $x = a - b, y = b - c, z = a - c$ are all positive integers colored the same color, and furthermore that $x + y = z$. \square

Corollary 15.1. *Fix $m \in \mathbb{N}$. There exists q such that for primes $p \geq q$, the equation $x^m + y^m = z^m$ has a solution in $\mathbb{Z}/p\mathbb{Z}$.*

Proof. Let $q = S(m) + 1$. Let g be a generator of the cyclic group $G = (\mathbb{Z}/p\mathbb{Z})^\times$. For any $x \in G$, we can write $x = g^a$. Furthermore, by the division algorithm, we can write $x = g^{mj+i}$ for some j and $0 \leq i < m$. We label the elements of G with m numbers using the i in the representation above.

By Schur's Theorem, there exist a, b, c all colored the same such that $a + b = c$: but this implies that $g^{mj_a+i} + g^{mj_b+i} = g^{mj_c+i}$ for some j_a, j_b, j_c and one fixed i . By multiplying through by g^{-i} and rearranging,

$$(g^{j_a})^m + (g^{j_b})^m = (g^{j_c})^m,$$

then by assigning $x = g^{j_a}, y = g^{j_b}, z = g^{j_c}$, we see that there exist x, y, z with $x^m + y^m = z^m$, exactly as desired. \square

5 The Lindström-Gessel-Viennot Lemma

Definition 29 (Locally Finite Graph). A *locally finite graph* is a graph with a (possibly infinite) set of vertices such that the set of edges of any single vertex is finite.

Definition 30 (Acyclic Graph). An *acyclic graph* is a graph with no cycles.

Definition 31 (Acyclic Plane). Consider the graph G with vertices \mathbb{Z}^2 with directed edges $(x, y) \rightarrow (x, y + 1)$ and $(x, y) \rightarrow (x + 1, y)$. In particular, G is a locally finite directed acyclic graph and we call it the *acyclic plane*.

Definition 32 (Path System). Given “sources” a_1, \dots, a_n and “sinks” b_1, \dots, b_n (both sets of distinguished vertices) in a locally finite directed graph G , a *path system* \mathcal{P} is an n -tuple of paths, sending each source to a distinct sink. Such a path system has an associated permutation: the permutation of $\{1, \dots, n\}$ corresponding to which source goes to which sink. We denote this permutation $\sigma(\mathcal{P})$.

Definition 33 (Non-Intersecting Path System). A path system is said to be a *non-intersecting path system* if no two distinct paths share a vertex.

Lemma 16 (The Lindström-Gessel-Viennot Lemma). *Consider a locally finite acyclic directed graph G with sources a_1, \dots, a_n and sinks b_1, \dots, b_n . Define a weight on each edge e in G , denoted $\omega(e)$, and define the weight of a path P in G to be the product of the weights of the edges in P , denoted $\omega(P)$. For any two vertices a and b , define $e(a, b)$ to be the sum of the weights of the paths between a and b .*

Also let the weight of a path system \mathcal{P} , denoted by $\omega(\mathcal{P})$, be $\prod_{P \in \mathcal{P}} \omega(P)$. Finally, let M be the matrix whose i, j -entry is $e(a_i, b_j)$. Then, if \mathcal{N} is the set of all non-intersecting path systems in G ,

$$\det M = \sum_{\mathcal{P} \in \mathcal{N}} \text{sign}(\sigma(\mathcal{P})) \omega(\mathcal{P})$$

Corollary 16.1. *If the only possible permutation of a non-intersecting path system \mathcal{P} is the identity permutation, then $\det M = \sum_{\mathcal{P} \in \mathcal{N}} \prod_{P \in \mathcal{P}} \omega(P)$.*

Corollary 16.2. *If the weight function is 1 (that is, the graph is unweighted) and furthermore the only possible permutation of a non-intersecting path system is the identity, then $\det M$ is the number of non-intersecting path systems in G .*

5.1 Proving the Lindström-Gessel-Viennot Lemma

For such a complex theorem to state, the proof is remarkably simple. Intuitively, we start with the Leibniz formula for determinants, and then just need to define an involution on the path systems with fixes the non-intersecting path systems but reverses the sign of the permutation of the intersecting path systems (so they all cancel out except the non-intersecting path systems).

Proof. Now recall the Leibniz formula for determinants: if A is a matrix with i, j entry $a_{i,j}$, then

$$\det(A) = \sum_{\sigma \in S_n} \text{sign}(\sigma) \prod_{i=1}^n a_{i, \sigma(i)}$$

In particular, for our matrix M , we see that:

$$\det(M) = \sum_{\sigma \in S_n} \text{sign}(\sigma) \prod_{i=1}^n e(a_i, b_{\sigma(i)}) = \sum_{\sigma \in S_n} \text{sign}(\sigma) \sum \{\omega(\mathcal{P}) : \mathcal{P} \text{ a path system with } \sigma(\mathcal{P}) = \sigma\}$$

Now, if \mathcal{S} is the set of all path systems in G , then the above sum (and therefore $\det M$) is equal to:

$$\sum_{\mathcal{P} \in \mathcal{S}} \text{sign}(\sigma(\mathcal{P}))\omega(\mathcal{P})$$

Now we seek to prove that the given sum over all $\mathcal{P} \in \mathcal{S} \setminus \mathcal{N}$ is 0. To do this, consider an involution on the intersecting path systems defined so:

1. Let i be the smallest index such that the path from a_i intersects another path and let j be the largest index such that the path from a_j intersects the path from a_i .
2. Let v be the vertex which is the first intersection between the paths from a_i and a_j .
3. Swap the two paths (the path from a_i and the path from a_j) *after* v , thereby negating the sign of the path system.

Clearly this is an involution, and since it negates the sign, we must have that the given sum over all $\mathcal{P} \in \mathcal{S} \setminus \mathcal{N}$ is 0, as desired, so we have our desired result:

$$\det M = \sum_{\mathcal{P} \in \mathcal{N}} \text{sign}(\sigma(\mathcal{P}))\omega(\mathcal{P})$$

□

5.2 An Application in Linear Algebra

Theorem 17. *Consider the following matrix:*

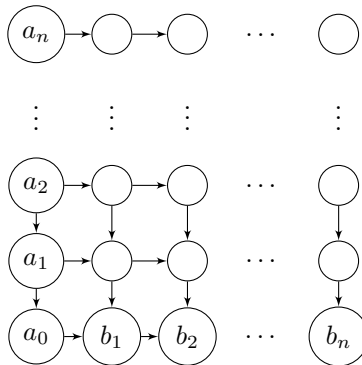
$$M = \begin{bmatrix} \binom{0}{0} & \binom{1}{0} & \binom{2}{0} & \cdots & \binom{n}{0} \\ \binom{1}{0} & \binom{2}{0} & \binom{3}{0} & \cdots & \binom{n+1}{0} \\ \binom{2}{0} & \binom{3}{0} & \binom{4}{0} & \cdots & \binom{n+2}{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \binom{n}{0} & \binom{n+1}{1} & \binom{n+2}{2} & \cdots & \binom{2n}{n} \end{bmatrix}$$

In short, we're inserting Pascal's triangle into a matrix – for example, for $n = 4$, the matrix becomes:

$$M = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 2 & 3 & 4 \\ 1 & 3 & 6 & 10 \\ 1 & 4 & 10 & 20 \end{bmatrix}$$

For any n , the determinant of this matrix is 1.

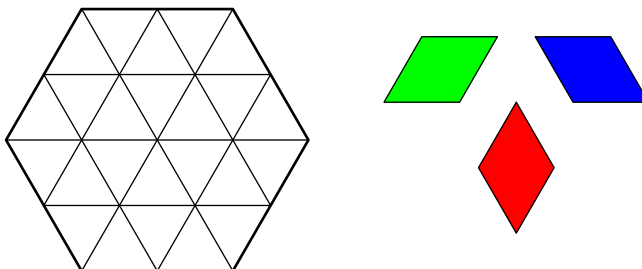
Proof. Consider the $n \times n$ finite acyclic directed graph G where all edges are given weight 1:



Notice that there is no b_0 . However, we define b_0 to be the same vertex as a_0 . It is not difficult to see that the number of paths from a_i to b_j is the i, j -entry of M (indexing from 0 in the rows and columns). But notice that there is only one possible non-intersecting path system, where each a_i travels right until it is directly above b_i , and then travels downward. This path system is associated with the identity permutation, whence by the Lindström-Gessel-Viennot Lemma, $\det M = 1$. \square

5.3 An Application in Tiling

Problem 15. How many ways can you tile the hexagon on the left (along the lines drawn) with the rhombi on the right? You can use as many rhombi as you want.



Hint: The answer follows easily from the the Lindström-Gessel-Viennot Lemma: create a vertex at each midpoint of a horizontal line segment, and create a directed edge down-and-to-the-left and down-and-to-the-right at from each vertex (where possible).

Then each non-intersecting path system generates a unique tiling of the graph (down-to-the-left represents a green rhombus, down-to-the-right represents a blue rhombus, and empty spaces become red rhombi). In fact, a more general result is true:

Problem 16. Prove that the number of ways to tile the regular hexagon with side length n in the manner described above is given by the following formula:

$$\begin{vmatrix} \binom{2n}{n} & \binom{2n}{n-1} & \binom{2n}{n-2} & \cdots & \binom{2n}{1} \\ \binom{2n}{n+1} & \binom{2n}{n} & \binom{2n}{n-1} & \cdots & \binom{2n}{2} \\ \binom{2n}{n+2} & \binom{2n}{n+1} & \binom{2n}{n} & \cdots & \binom{2n}{3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \binom{2n}{2n-1} & \binom{2n}{2n-2} & \binom{2n}{2n-3} & \cdots & \binom{2n}{n} \end{vmatrix}$$

Again, this follows immediately from the Lindström-Gessel-Viennot Lemma.

6 Coloring Graphs

Definition 34 (Vertex Coloring). A *vertex coloring* of a graph G is a coloring of every vertex in G such that no vertices which share an edge share a color.

Definition 35 (Chromatic Number). The *chromatic number* of a graph G , denoted $\chi(G)$, is the smallest number of colors required to vertex-color G .

The problem of coloring graphs originated with the issue of coloring maps. In any good map, every region (say, state) is colored a different color from all adjacent regions. This is not hard if one is given any reasonable amount of colors (5 or 6 is plenty to do it easily). However, we want to find the minimal number of colors necessary to color such a map.

Our first step is to notice that because a map is naturally a planar graph, we can take the dual graph of the map to make our map coloring problem the following *graph coloring problem*.

Definition 36. A *graph coloring* is a coloring of every vertex in a graph such that no vertex which share an edge share a color.

Problem 17. How many colors are required to color a planar graph?

Of course, if we don't restrict ourselves to planar graphs, there is no limit on the number of colors required: the graph K_n requires n colors. Thus, this is the most interesting form of the graph coloring problem.

Theorem 18 (The Four-Color Theorem). *If G is a planar graph, then $\chi(G) \leq 4$.*

This theorem, proven in 1976, is still not human-comprehensible: it requires the checking of over 600 cases. Thus, we cannot document it here (and it is honestly not worth the effort to go through the proof). However, we will prove the Five-Color Theorem in the next subsection.

One common treatment of the chromatic number is via the *chromatic polynomial*, which is the function f_G on the natural numbers defined by $f_G(r) := \{\text{the number of colorings of } G \text{ with } r \text{ colors}\}$. The fact that we use the name chromatic "polynomial" should be surprising: it is not at all obvious that such a function is a polynomial. We will prove now that this is the case:

Definition 37 (Deletions). Let $e = ab$ be an edge of a graph $G = (V, E)$. Then the graph $G - e$ is called the *deletion* of e and equals $(V, E \setminus \{e\})$.

Definition 38 (Contractions). The *contraction* of e is denoted G/e and given by replacing x and y by a new vertex z and replacing any edge $\{v, x\}$ or $\{v, y\}$ with an edge $\{v, z\}$ for every v .

Theorem 19. *For any graph G with an edge e , $f_G(r) = f_{G-e}(r) - f_{G/e}(r)$*

Proof. Divide the set of colors of $G - e$ into two disjoint classes:

1. Those for which x and y receive different colors: these are all valid colorings of G .
2. Those for which x and y receive the same color. These colorings correspond exactly to colorings of G/e .

From this the result follows. □

Corollary 19.1. *This implies that the chromatic polynomial of G is a polynomial in $\mathbb{Z}[r]$ of degree n (where n is the number of vertices of G).*

Proof. Consider the graph with one vertex and no edges, which has chromatic polynomial r and thus satisfies the condition. Then induction on the sum $e + v$ (where v is the number of vertices and e is the number of edges) gives the result for all graphs. □

6.1 The Five-Color Theorem

Lemma 20. *Every planar graph G contains a vertex v such that $\deg v \leq 5$.*

Proof. Trivial from the bound $e \leq 3v - 6$. □

Theorem 21. *For any planar graph G , $\chi(G) \leq 5$.*

Proof. We induct on p , the number of vertices in G . If $p = 1$, then the chromatic number is 1. Now suppose that the theorem holds up through graphs of size P (that is, $p \leq P \Rightarrow \chi(G) \leq 5$). Let G be a graph of order $V + 1$. By the lemma, there is a vertex v with $\deg v \leq 5$. By the inductive hypothesis,

$G - v$ has a 5-coloring. If $\deg v \leq 4$, then there is a color which is not used by any of v 's neighbors, so we can color v and the result follows. Thus we may assume that $\deg v = 5$.

Now notice that it is always possible to find 2 vertices adjacent to v that are not adjacent to each other. If this was not the case, there would be a subgraph of G isomorphic to K_5 , a contradiction with the fact that G is planar.

Now let u and w be vertices adjacent to v but not each other. Form a graph H from G by identifying u and w . Now, H is planar, because we can turn it into a plane graph by deforming u and w to v along uv and wv . Thus there is a 5-coloring for this graph, but that gives a 5-coloring on $G - v$ which has 4 or fewer colors adjacent to v – so we can color the entirety of G by coloring v the last color. \square

7 Extra Topics

7.1 Matchings

Recall the definition of a bipartite graph from section 1. A key example of the application of such a graph is the matching of applicants (set A) to positions (set P). An edge is drawn between an applicant and a position if the applicant is suitable for the job. Then, we seek to know when it is possible for us to assign every applicant a to a position $p \in P$.

Definition 39 (Diversity Condition). For each group of k applicants (for all $k = 1, 2, \dots, n$), there must be at least k jobs for which, collectively, they are qualified.

Theorem 22. *If the diversity condition is fulfilled, one can assign a suitable job to each applicant.*

Proof. Clearly the theorem is true for $n = 1$ and $n = 2$. We will thus try to prove the result by induction: assume that the result is true for $n - 1$ applicants and fewer.

Case 1: Every possible group of k applicants ($k = 1, 2, 3, \dots, n - 1$) qualifies for *more* than k jobs. In this case, pick any one of the n applicants and place him in one of the jobs for which he is suited. Then the diversity condition holds for the remaining $n - 1$ applicants, so by induction there is a way to put all of them too into suitable jobs.

Case 2: There is some possible set A_0 of k_0 applicants whose are qualified for exactly k_0 jobs. Because of the diversity condition, no k of these applications can be qualified for fewer than k jobs. By the induction hypothesis, the k_0 applicants in A_0 can be placed can be assigned suitable jobs.

But then the diversity conditions are still satisfied for the $n - k_0$ applicants left (and the remaining jobs), so by induction the rest of the applicants can also be assigned jobs! Thus the result follows. \square

Definition 40 (Graph Matching). A *graph matching from A into P* on a bipartite graph is a set of edges such that each vertex in A is linked to a different vertex in P . Thus, the above theorem becomes: a matching of A into P is possible iff any k vertices in A are connected by edges to at least k vertices in P .

There are numerous applications of this principle. Here's a quite simple example:

Problem 18. Suppose that are n committees C_1, \dots, C_n and a set of p_1, \dots, p_m people belonging to these committees. We seek to choose a secretary from each committee such that no person is a secretary for multiple committees. Prove that this is possible if and only if every group of k committees includes k distinct people.

Problem 19. Prove that if each committee has at least t members, and no individual belongs to more than t committees, then it is always possible to find a separate secretary for each committee.