MATH 5707 Midterm 2

Sasha Pevzner (edited by Darij Grinberg)

April 13, 2017

1. Let $G = (V, E)$ be a simple graph such that $|E| \geq |V|$. Then there exists an injective map $f : V \to E$ such that each $v \in V$ satisfies $v \notin f(v)$.

   Proof. Let $n = |V|$, and assume without loss of generality that $|E| = n$ (else we could remove edges from G until $|E| = n$ and apply the argument). Also assume WLOG that $V \cap E = \varnothing$ (else, we can rename the vertices). Note that we must have $n \geq 3$, since a simple graph with 2 or fewer vertices can have at most 1 edge. Now we seek a bijection $f : V \to E$ with each $v \in V$ satisfying $v \notin f(v)$. Construct a bipartite graph $(\widetilde{G}; V, E)$, where $v \in V$ is adjacent to $e \in E$ if and only if $v \notin e$. We seek a perfect matching, or simply an E-complete matching of $\widetilde{G}$. First we note that since each $e \in E$ has exactly 2 endpoints in G, it has exactly two elements of V which are *not* neighbors of e in $\widetilde{G}$. That is, for each $e \in E$, we have $\deg_{\widetilde{G}}(e) = n-2$. In other words, $|N(\{e\})| = n-2 \geq 1$, since $n \geq 3$. Thus, we have the Hall condition satisfied for subsets of E of size 1. Furthermore, we have that any nonempty subset P of E satisfies $|N(P)| \geq n-2$ and thus we have verified the Hall condition for all subsets of E with size $\leq n-2$. It remains to verify the Hall condition for subsets P of E with $|P| = n-1$ and $|P| = n$. If $|P| = n$, then $P = E$, and we claim that $N(P) = V$. Indeed, assume the contrary. Then there exists $v \in V$ such that v is isolated in $\widetilde{G}$. By definition of $\widetilde{G}$, this means that v is an endpoint of every edge in G. Since G is simple, each edge in G must have a unique second endpoint. We have n edges, but only at most $n-1$ of these unique endpoints and by the Pigeonhole Principle we've reached a contradiction. Thus, $N(E) = V$, and we have $|N(E)| = |V| = |E|$. Now let $e \in E$, and consider $P = E \setminus \{e\}$. Then $|P| = n-1$. Suppose that $|N(P)| \leq n-2$. This means that we have two vertices incident with all edges in P. But $|P| = n-1 \geq 3-1 = 2$, and thus we have two vertices that must be simultaneously incident with two distinct edges. This can only be achieved if we allow parallel edges, but G is simple, so this cannot be the case. Thus we can conclude $|N(P)| > n-2$, so that $|N(P)| \geq n-1 = |P|$. $\square$

2. Let $G = (V, E)$ be a connected simple graph such that $|E| \geq |V|$. Then there exists an injective map $f : V \to E$ such that each $v \in V$ satisfies $v \in f(v)$.

   Proof. We shall describe a method to construct such a map. Since G is connected, consider a spanning tree $T = (V, F)$ of G, where $F \subset E$. First, for each leaf ℓ of T, define $f(\ell)$ to be its unique incident edge in T. Now, consider T with all leaves (and incident edges) removed. This new graph is T', and we can define $f(\ell)$ for each leaf ℓ of T' to be its unique incident edge in T'. Note that if ℓ is a leaf in T, and ℓ' is a leaf in T', we cannot have $f(\ell) = f(\ell')$ since all edges incident to leaves of T were removed in the construction of T'. We continue this process of pruning leaves and assigning edges until we are left with either one vertex (if T has one center) or two vertices connected by an edge (if T has two centers). If we have two vertices connected by an edge, choose one of them arbitrarily and assign to it the remaining edge of T. Now in either case we have one vertex remaining, call it c, and all other vertices v have been assigned a unique edge $f(v)$ satisfying the property that $v \in f(v)$. Now we must assign an edge to c. If c is incident with an edge in $E \setminus F$, then we can assign c to this edge in the mapping f and we are done. Assume that there does not exist $e \in E \setminus F$ such that c is incident with e. Then, there exists $v \in V \setminus \{c\}$ such that v is incident with an edge in $E \setminus F$. Denote this edge by e. Indeed, since T is a spanning tree, it has $|V| - 1$ edges, but by assumption G satisfies $|E| \geq |V|$ and consequently $E \setminus F \neq \varnothing$. We can now change the mapping f such that $f(v) = e$. If c is adjacent to v, then we can assign the edge $\{c, v\}$ to c in the mapping f and we are done. If c is not adjacent to v, then consider the neighbor w of v that is closest to c in T. Since v is further from c, it was pruned earlier than w and thus the edge $\{w, v\}$ was assigned to v. Now, since $f(v) = e \neq \{w, v\}$, we can assign $\{w, v\}$ to w. Then, we can look at the neighbor u of w that is closest to c in T. Again we see that $\{u, w\}$ was originally assigned to w, but now it is free and can be assigned to u. We continue this process until $u = c$, and we have assigned to c an edge incident with it. □

3. Let $D = (V, A)$ be a digraph. Let $k \in \mathbb{N}$. Let u,v, and w be three vertices of D. Assume there exist k arc-disjoint paths from u to v. Assume furthermore that there exist k arc-disjoint paths from v to w. Then, there exist k arc-disjoint paths from u to w.

4. Let $G = (V, E)$ be a simple graph. Define a polynomial χ_G in a single indeterminate x with integer coefficients by

$$\chi_G = \sum_{F \subseteq E} (-1)^{|F|} x^{\mathrm{conn}(V,F)}$$

   Fix $k \in \mathbb{N}$. Then, the number of proper k-colorings of G is $\chi_G(k)$.

Proof. Let $n = |V|$, $m = |E|$. If we wish to count the number of proper k-colorings of G, we can count the total number of k-colorings of G and subtract off the number of "improper" k-colorings of G (i.e. those that fail to be proper). We can characterize an improper k-coloring of G to be one in which there exists $e \in E$ such that both endpoints of e are assigned to the same color. Create a listing $e_1, ..., e_m$ of all of the edges in G. Now, for each $i \in \{1, ..., m\}$, define the set $\mathcal{A}_i$ by

$$\mathcal{A}_i = \{f \in \{1, ..., k\}^V \mid f(u) = f(v), \text{ where } e_i = \{u, v\}\}$$

In other words, $\mathcal{A}_i$ denotes the set of all k-colorings of G where the endpoints of e_i share the same color. Then, we see that the number N of proper k-colorings of G is given by

$$\begin{aligned} N &= \#\{\text{all } k\text{-colorings of } G\} - \#\{\text{improper } k\text{-colorings of } G\} \\ &= |\{1, ..., k\}^V| - \left|\bigcup_{i=1}^{m} \mathcal{A}_i\right| \\ &= k^n - \left|\bigcup_{i=1}^{m} \mathcal{A}_i\right| \end{aligned} \tag{1}$$

By inclusion-exclusion, we have

$$\left|\bigcup_{i=1}^{m} \mathcal{A}_i\right| = \sum_{i=1}^{m} |\mathcal{A}_i| - \sum_{1 \le i < j \le m} (|\mathcal{A}_i \cap \mathcal{A}_j|) + ... + (-1)^{m-1} |\mathcal{A}_1 \cap ... \cap \mathcal{A}_m| \tag{2}$$

Let us consider one such intersection of $\mathcal{A}_i$'s. More precisely, consider $F = \{e_{i_1}, ..., e_{i_\ell}\}$, an ℓ-element subset of E, where $1 \le \ell \le m$. Then,

$$\mathcal{A}_{i_1} \cap ... \cap \mathcal{A}_{i_\ell}$$

represents the set of all k-colorings of G in which the endpoints of each edge of F share the same color. Thus, we can assign a color to each connected component of (V, F) independently, but we'd like all vertices of one connected component to share a color (since they will be connected by edges of F). It follows that

$$|\mathcal{A}_{i_1} \cap ... \cap \mathcal{A}_{i_\ell}| = k^{\text{conn}(V,F)}$$

Then, (2) becomes

$$\begin{aligned}
\left|\bigcup_{i=1}^{m} \mathcal{A}_i\right| &= \sum_{\substack{F\subseteq E,\\ |F|=1}} k^{\mathrm{conn}(V,F)} - \sum_{\substack{F\subseteq E,\\ |F|=2}} k^{\mathrm{conn}(V,F)} + ... + (-1)^{m-1}k^{\mathrm{conn}(V,E)} \\
&= (-1)^{1-1}\sum_{\substack{F\subseteq E,\\ |F|=1}} x^{\mathrm{conn}(V,F)} + (-1)^{2-1}\sum_{\substack{F\subseteq E,\\ |F|=2}} k^{\mathrm{conn}(V,F)} + ... + (-1)^{m-1}k^{\mathrm{conn}(V,E)} \\
&= \sum_{\substack{F\subseteq E\\ F\neq\varnothing}} (-1)^{|F|-1}k^{\mathrm{conn}(V,F)}
\end{aligned}$$

We also note that $k^n = k^{\mathrm{conn}(V,\varnothing)}$, since the graph $(V,\varnothing)$ consists of n isolated vertices. Then, (1) becomes

$$\begin{aligned}
N &= k^{\mathrm{conn}(V,\varnothing)} - \sum_{\substack{F\subseteq E\\ F\neq\varnothing}} (-1)^{|F|-1}k^{\mathrm{conn}(V,F)} \\
&= (-1)^{|\varnothing|}k^{\mathrm{conn}(V,\varnothing)} + \sum_{\substack{F\subseteq E\\ F\neq\varnothing}} (-1)^{|F|}k^{\mathrm{conn}(V,F)} \\
&= \sum_{F\subseteq E} (-1)^{|F|}k^{\mathrm{conn}}(V,F) \\
&= \chi_G(k)
\end{aligned}$$

□

5. (a) For each $n \in \mathbb{N}$, the complete graph K_n has chromatic polynomial

$$\chi_{K_n} = x(x-1)...(x-n+1).$$

Proof. Let $n \in \mathbb{N}$. By 4, we have that the number of proper k-colorings of any graph G is $\chi_G(k)$, and thus we shall show that for each k, we have that the number of proper k-colorings of K_n is given by $k(k-1)....(k-n+1)$. Let $k \in \mathbb{N}$. If we wish to color K_n properly, each vertex must have a different color, since every pair of vertices in K_n is adjacent. As a result we must have $k \geq n$, so assume this is the case. We can color vertices sequentially, so we create a listing $v_1, ..., v_n$ of the vertices of K_n and color them in this order. Then we have k choices with which to color v_1. Since v_2 is adjacent to v_1, we have $k-1$ choices with which to color v_2 since it cannot share its color with v_1. Each time we color a new vertex, we have one less choice for which color we use than for the previous vertex. Once properly colored, there will be $k-n$ unused colors, and so the number of choices with which to color v_n will be $k-n+1$. Thus, after coloring all n vertices, we have that the number of ways to properly k-color K_n is

$$k(k-1)...(k-n+1).$$

So we have shown that $\chi_{K_n}(k) = k(k-1)...(k-n+1)$ for all integers $k \geq n$. Therefore, the two polynomials χ_{K_n} and $x(x-1)...(x-n+1)$ have the same value whenever x is set to be an integer $\geq k$. In other words, the polynomial $\chi_{K_n} - x(x-1)...(x-n+1)$ vanishes whenever x is set to be an integer $\geq k$. Thus, this polynomial has infinitely many zeroes (viz., all integers $\geq k$). Since the only polynomial with infinitely many zeroes is the zero polynomial, we therefore obtain that $\chi_{K_n} - x(x-1)...(x-n+1)$ is the zero polynomial. This solves part (a). □

(b) Let T be a tree. Let $n = |V(T)|$. Then

$$\chi_T = x(x-1)^{n-1}.$$

Proof. Again we will show that T has $k(k-1)^{n-1}$ proper k-colorings for each $k \in \mathbb{N}$. Fix $k \in \mathbb{N}$, and we will induct on n. For the base case, consider $n = 1$. Then the tree consisting of a single vertex has $k = k(k-1)^{1-1}$ proper k-colorings, since we can choose any of the k colors to assign to the unique vertex. Now assume that all trees on $n-1$ vertices (where $n \geq 2$) have $k(k-1)^{n-2}$ proper k-colorings. Consider T on n vertices. Choose a leaf $\ell \in V(T)$ (assuming $n \geq 2$), and remove it along with its incident edge to create a new tree T' on $n-1$ vertices. By our inductive hypothesis, T' can have $k(k-1)^{n-2}$ proper k-colorings. Now, if we reconstruct T from T' and properly k-color the vertices of T', we have $k-1$ ways in which to color ℓ, since it only has one neighbor. Thus, the number of ways to properly k-color T on n vertices is

$$k(k-1)^{n-2}(k-1) = k(k-1)^{n-1}.$$

From here, we proceed as in (a) to prove the polynomial identity $\chi_T = x(x-1)^{n-1}$. □

(c) The chromatic polynomial of the path graph P_3 is

$$\chi_{P_3} = x^3 - 2x^2 + x$$

Proof. Let $V = V(P_3)$, $E = E(P_3)$, and a, b be the two edges of E (that is, $E = \{a, b\}$). Then, by definition of the chromatic polynomial, we have

$$\chi_{P_3} = (-1)^{|\varnothing|}x^{\text{conn}(V,\varnothing)} + (-1)^{|\{a\}|}x^{\text{conn}(V,\{a\})} + (-1)^{|\{b\}|}x^{\text{conn}(V,E\{b\})} + (-1)^{|E|}x^{\text{conn}(V,E)} \quad (1)$$

Since P_3 has 3 vertices, we have that

$$\text{conn}(V, \varnothing) = 3.$$

When we remove 1 edge from E, we are left with an isolated vertex and a path of length 1. In other words,

$$\text{conn}(V, E\{a\}) = \text{conn}(V, E\{b\}) = 2.$$

Thus, (1) becomes

$$\begin{aligned}\chi_{P_3} &= (-1)^0 x^3 + (-1)^1 x^2 + (-1)^1 x^2 + (-1)^2 x^1 \\ &= x^3 - 2x^2 + x\end{aligned}$$

□

6. Let G be a tree. Let x, y, z, w be four vertices of G. Then the two larger ones among the numbers $d(x, y) + d(z, w)$, $d(x, z) + d(y, w)$, and $d(x, w) + d(y, z)$ are equal.

Proof. We will use induction on $|V(G)|$. Since x, y, z, w need not be distinct, we have that our base case is when $|V(G)| = 1$. Then $x = y = z = w$, and $d(x, y) + d(z, w) = d(x, z) + d(y, w) = d(x, w) + d(y, z) = 0$. Now let $n \geq 2$, and assume the statement holds for all trees G' with $|V(G')| = n - 1$. Now consider G with $|V(G)| = n$. Choose a leaf $\ell \in V(G)$ and remove it along with its incident edge to obtain a tree G' on $n-1$ vertices. If none of x, y, z, or w is equal to ℓ, then by our induction hypothesis we are done, so assume that one of x, y, z, w is equal to ℓ. Without loss of generality, assume $w = \ell$. Assume also that w is distinct from x, y, and z. Let w' be the unique neighbor of w in G. Then $w', x, y, z \in V(G')$. Since w is a leaf, we have for any $v \in V(G) \setminus \{w\}$, $d_G(v, w) = d_G(v, w') + 1$. Therefore,

$$\begin{aligned}d_G(x, y) + d_G(z, w) &= d_G(x, y) + d_G(z, w') + 1 \\ d_G(x, z) + d_G(y, w) &= d_G(x, z) + d_G(y, w') + 1 \\ d_G(x, w) + d_G(y, z) &= d_G(x, w') + d_G(y, z) + 1\end{aligned}$$

But then for any $u, v \in V(G')$, we have $d_G(u, v) = d_{G'}(u, v)$. By our induction hypothesis, in G' we have that the largest two among $d_{G'}(x, y) + d_{G'}(z, w')$, $d_{G'}(x, z) + d_{G'}(y, w')$, and $d_{G'}(x, w') + d(y, z)$ are equal. Then the largest 2 among $d_{G'}(x, y) + d_{G'}(z, w') + 1$, $d_{G'}(x, z) + d_{G'}(y, w') + 1$, and $d_{G'}(x, w') + d(y, z) + 1$ are equal. Thus the statement holds when x, y, and z are distinct from w (but not necessarily from each other).

Now consider the case where, without loss of generality, $z = w$. Then $d(z, w) = 0$, and we are comparing the three numbers $d(x, y)$, $d(x, z) + d(y, w)$, and $d(x, w) + d(y, z)$. Note that the latter two are equal, and thus we need to show that we cannot have $d(x, y) > d(x, z) + d(y, w)$. But by the symmetry of distance and the triangle inequality, $d(x, z) + d(y, w) = d(x, w) + d(w, y) \geq d(x, y)$. Thus the statement holds. $\square$

7. Let $G = (V, E, \phi)$ be a multigraph. Let A, B, and C be three subsets of V such that the sub-multigraphs $G[A]$, $G[B]$, and $G[C]$ are connected.

   (a) If the sets $B \cap C$, $C \cap A$, $A \cap B$ are nonempty, but $A \cap B \cap C$ is empty, then G has an eclectic cycle.

   Proof. Assume $B \cap C$, $C \cap A$, $A \cap B$ are nonempty, and assume $A \cap B \cap C$ is empty. Let $v \in A \cap B$, $u \in B \cap C$, $w \in C \cap A$. Then, since $G[A]$ is connected, we have that there exists a path $w \to v$ in $G[A]$. Since $G[B]$ is connected, there exists a path $u \to v$ in $G[B]$. Since $G[C]$ is connected, there exists a path $w \to u$ in $G[C]$. Concatenating the latter two paths and removing cycles, we obtain a path $w \to v$ in $G[B \cup C]$ which contains at least one vertex $q \in B \cap C$, and moreover enters this vertex q along an edge of $G[B]$ and exits it along an edge of $G[C]$. This path therefore cannot be identical with the old path $w \to v$ in $G[A]$ (because if it was, then $q \in B \cap C$ would also belong to A, contradicting the assumption that $A \cap B \cap C$ be empty). Hence, we have found two distinct paths $w \to v$ in G. Consequently, there must be a cycle consisting of a nonempty segment of one path and a nonempty segment of the other[1]. With a bit more work, we can see that there exists an **eclectic** cycle.[2] $\square$

   (b) If the subgraphs $G[B \cap C]$, $G[C \cap A]$, and $G[A \cap B]$ are connected, but the subgraph $G[A \cap B \cap C]$ it not connected, then G has an eclectic cycle.

   Proof. Assume the subgraphs $G[B \cap C]$, $G[C \cap A]$, and $G[A \cap B]$ are connected and that the subgraph $G[A \cap B \cap C]$ is not connected. Then, since $G[B \cap C]$, $G[C \cap A]$, and $G[A \cap B]$ are connected, they each have at least one vertex and thus $B \cap C$, $C \cap A$, and $A \cap B$ are all non-empty. If $A \cap B \cap C = \varnothing$, then we can apply (a). Thus assume $A \cap B \cap C \neq \varnothing$. Then there exist distinct $x, v \in A \cap B \cap C$ such that x

[1] Here we are using the following general observation about multigraphs: If s and t are two vertices of a multigraph, and if there exist two distinct paths $s \to t$, then there must be a cycle consisting of a nonempty segment of one path and a nonempty segment of the other.

[2] To obtain such a cycle, concatenate the former path $w \to v$ with the reversal of the latter, obtaining a circuit; then, rotate this circuit to make sure that it starts and ends at q; then, keep removing cycles until only a single cycle remains. The result will be a cycle that enters q along an edge of $G[B]$, exits q along an edge of $G[C]$, and somewhere inbetween also uses an edge of $G[A]$ – hence, an eclectic cycle.

and v lie in two different connected components of $G[A \cap B \cap C]$. Since $G[A \cap B]$ is connected, and x and v belong to $A \cap B$, there exists a path $\mathbf{p}$ in $G[A \cap B]$ from x to v. Similarly, there exists a path $\mathbf{q}$ from x to v in $G[B \cap C]$. We claim that $\mathbf{p}$ and $\mathbf{q}$ are distinct and thus constitute a cycle in G. Indeed, if they were not distinct, then $\mathbf{p}$ would also be a path in $G[B \cap C]$, and thus a path in $G[C]$. But then $\mathbf{p}$ would be a path from x to v in $G[A \cap B \cap C]$, contradicting our assumption that no such path exists. Thus, $\mathbf{p}$ and $\mathbf{q}$ are distinct and we have a cycle in G. This cycle is eclectic, since it uses at least one edge of $G[A \cap B]$ from the path $\mathbf{p}$ (and thus an edge of $G[A]$ and an edge of $G[B]$), and since it uses at least one edge of $G[B \cap C]$ from the path $\mathbf{q}$, this edge must lie in $G[C]$. □