Perfect Graph Theory	Instructor: Padraic Bartlett

Lecture 1: Motivation: Why $\chi(G)$ Isn't a Local Property

Week 1 — *Mathcamp 2011*

1 The Backstory

In last year's introduction to graph theory course (indeed, in pretty much every introductory graph theory course) one of the first theorems we proved was the following:

Theorem 1 *A graph*[1] *G is bipartite*[2] *if and only if it doesn't contain any odd-length cycles*[3].

It's been a while for most of you since you were last in a graph theory course, so as a warm-up we'll reprove this result, to get you all used to the terminology again:

Proof. First, let's take any graph G that contains a cycle C_{2k+1} of odd length. Consider any coloring of $V(G)$'s vertices by the colors red and blue; this coloring will also color the vertices of C_k. We claim that this coloring will force C_{2k+1} to contain a monochromatic edge.

Indeed, if it didn't, then (if you suppose v_{k+1} was red without any loss of generality,) you'd have

- v_{k+1} being red force
- v_k, v_{k+2} to be blue, which will force
- v_{k-1}, v_{k+3} to be red, which wil force
- $\ldots$
- which will force v_1, v_{2k+1} to both be the same color; a contradiction.

So it suffices to prove the other direction. Without loss of generality we can assume G is connected[4], by simply applying our proof to each of G's connected components.

To do this, take any vertex $y \in V(H)$, and construct the following sets:

- $N_0 = \{w : d(v, y) = 0\}$
- $N_1 = \{w : d(v, y) = 1\}$
- $N_2 = \{w : d(v, y) = 2\}$

[1]A **graph** is a collection of vertices V and unordered pairs of distinct vertices E.

[2]A graph is **bipartite** iff there's a way to color all of its vertices either red or blue, so that none of its edges have both of their endpoints the same color.

[3]A **cycle** C_n is a graph correpsonding to a n-gon.

[4]A graph G is **connected** if there is a path between any two vertices in G; a **path** in G is just a sequence of vertices and edges in G

- ...
- $N_n = \{w : d(v, y) = n\}$

Notice that every vertex shows up in precisely one of these sets, as $d(x, y)$, the length of the shortest path from x to y, is a well-defined object. As well, notice that for any $x \in N_k$ and any path P given by $y = v_0 e_{01} v_1 e_{12} \ldots e_{k-1,k} v_k = x$, each of the vertices v_j lies in N_j. This is because each of these has a path of length j from y to v_j (just take our path and cut it off at v_j), and has no shorter path (by definition.)

Now, color all of the vertices in the even N-sets red, and all of the vertices in the odd N-sets blue. We claim that there are no monochromatic edges.

To see this, take any edge $\{v_1, v_2\}$ in our graph H. Let $d(y, v_1) = k$ and $d(y, v_2) = l$, P_1 and P_2 a pair of paths from v_1, v_2 to y with those lengths k, l, x to be the furthest-away vertex from y that's in both of those paths, and P_1', P_2' be the paths we get by starting P_1, P_2 at x instead and proceeding to v_1, v_2.

Then, either x was one of v_1 or v_2 (in which case the distances of v_1 and v_2 from y have different parities, and therefore v_1 and v_2 are different colors), or the combination of P_1', P_2', and the edge $\{v_1, v_2\}$ forms a cycle, which cannot have odd length; this forces the lengths of P_1' and P_2' to again have different parities, and therefore (again) forces v_1 and v_2 to have different colors. There are therefore no monochromatic edges, and we have created a proper 2-coloring of our graph G, which means it's bipartite.

Later on in the course, we took this notion of "bipartite graphs" – graphs that can have their vertices colored red and blue, so that no edge has both of its endpoints the same color – and extended it to the idea of a k-coloring. We repeat this definition here:

Definition. A graph G is **k-colorable** if we can assign the colors $\{1, \ldots k\}$ to the vertices in $V(G)$, in such a way that every vertex gets exactly one color and no edge in $E(G)$ has both of its endpoints colored the same color. For a fixed graph G, if k is the smallest number such that G admits a k-coloring, we say that the chromatic number of G is k, or that G is k-chromatic or k-partite, and write $\chi(G) = k$.

A natural question to ask, given our first theorem, is whether we can come up with a similar characterization for (say) all of the 3-chromatic graphs– in other words, whether we can classify all of the graphs that admit a 3-coloring by a criterion as simple as "doesn't have any odd cycles." Surprisingly, there doesn't seem to be any such simpler classification; while graph theorists have certainly uncovered tons of 3-chromatic graphs, there doesn't seem to be any unifying theme or property tying them together beyond needing 3 colors to properly color them.

A (perhaps more promising) question to ask, then, is this: is there a way to just simply classify all of the graphs with **large** chromatic number? In other words, if I tell you that a graph has chromatic number $3 \cdot 10^8$, can you tell me anything about that graph at all?

One of the simplest bounds we came up with, in this class, was that the chromatic number of a graph was bounded below by its **clique number**[5]:

$$\omega(G) \leq \chi(G).$$

[5] The clique number of a graph, $\omega(G)$, is the largest value of k for which there is an induced subgraph in G isomorphic to K_k.

Is this bound **tight**? I.e. if we look at that graph with chromatic number $3 \cdot 10^8$, will we find a large complete subgraph sitting somewhere inside of it? Such a claim certainly sounds promising: certainly, if you try checking some examples by hand or constructing a graph with a high chromatic number, you'll almost certainly find that it has a subgraph isomorphic to a rather large complete subgraph.

Given this, then, the following construction of Mycielski is rather surprising:

Example. The **Mycielski** construction, described here, is a method for turning a triangle-free graph with chromatic number k into a larger triangle-free graph with chromatic number $k+1$. It runs as follows:

- As input, take a triangle-free graph G with $\chi(G) = k$ and vertex set $\{v_1, \dots v_n\}$.
- Form the graph G' as follows: let $V(G') = \{v_1, \dots v_n\} \cup \{u_1, \dots u_n\} \cup \{w\}$.
- Start with $E(G') = E(G)$.
- For every u_i, add edges from u_i to all of v_i's neighbors.
- Finally, attach an edge from w to every vertex $\{u_1, \dots u_n\}$.

Starting from the triangle-free 2-chromatic graph K_2, here are two consecutive applications of the above process:

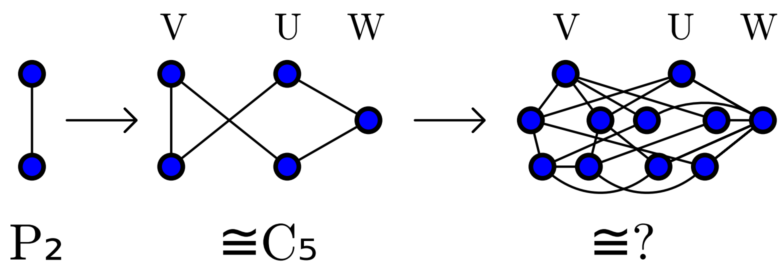

Proposition 2 *The above process does what it claims: i.e. given a triangle-free graph with chromatic number k, it returns a larger triangle-free graph with chromatic number $k+1$.*

Proof. Let G, G' be as described above. For convenience, let's refer to $\{v_1, \dots v_n\}$ as V and $\{u_1, \dots u_n\}$ as U. First, notice that there are no edges between any of the elements in U in G'; therefore, any triangle could not involve two elements from U. Because G was triangle-free, it also could not consist of three elements from V; finally, because w is not connected to any elements in V, no triangle can involve w. So, if a triangle exists, it must consist of two elements v_i, v_j in V and an element u_l in U; however, we know that u_l's only neighbors in V are the neighbors of v_l. Therefore, if (v_i, v_j, u_l) was a triangle, (v_i, v_j, v_l) would also be a triangle; but this would mean that G contained a triangle, which contradicts our choice of G.

Therefore, G' is triangle-free; it suffices to show that G' has chromatic number $k+1$.

To create a proper $k+1$-coloring of G': take a proper coloring $f : V(G) \to \{1, \dots k\}$ and create a new coloring map $f' : V(G') \to \{1, \dots k+1\}$ by setting

- $f'(v_i) = f(v_i)$,

- $f'(u_i) = f(v_i)$, and
- $f(w) = k + 1$.

Because each u_i is connected to all of v_i's neighbors, none of which are colored $f(v_i)$, we know that no conflicts come up there; as well, because $f(w) = k + 1$, no conflicts can arise there. So this is a proper coloring.

Now, take any k-coloring g of G': we seek to show that this coloring must be improper, which would prove that G' is $k+1$-chromatic. To do this: first, assume without any loss of generality that $f(w) = k$ (it has to be colored something, so it might as well be k.)

Then, because w is connected to all of U, the elements of U must be colored with the elements $\{1, \ldots k - 1\}$. Let $A = \{v_i \in V : g(v_i) = k\}$. We will now use U to recolor these vertices with the colors $\{1, \ldots k - 1\}$: if we can do this properly, then we will have created a $k - 1$ proper coloring of G, a k-chromatic graph (and thus arrived at a contradiction.)

To do this recoloring: change g on the elements of A so that $g(v_i)$'s new color is $g(u_i)$. We now claim that g is a proper $k-1$ coloring of G itself. To see this: take any edge $\{v_i, v_j\}$ in G. If both of $v_i, v_j \notin A$, then we didn't change the coloring of v_i and v_j; so this edge is still not monochromatic, because g was a proper coloring of G'. If $v_i \in A$ and $v_j \notin A$, then v_j is a neighbor of v_i and thus (by construction) u_i has an edge to v_j. But this means that $g(u_i) \neq g(v_j)$, because g was a proper coloring of G': therefore, this edge is also not monochromatic!

Because there are no edges between elements of A (as they were all originally colored k under g, and therefore there weren't any edges between them,) this covers all of the cases: so we've turned g into a $k-1$ coloring of a k-chromatic graph. As this is impossible, we can conclude that g cannot exist – i.e. G' cannot be k-colored! So $\chi(G') = k + 1$, as claimed.

As the above construction shows, our bound of $\omega(G) \leq \chi(G)$ can be remarkably useless: using it, we can make graphs where $\omega(G) = 2$ and $\chi(G) \to \infty$!

Surprisingly, this result can be extended to avoiding all short cycles, not just triangles! Erdős in the 60's, as one of the first applications of probabilistic methods to graph theory, showed that almost every graph G on n vertices has high chromatic number and relatively few short cycles, and therefore (by deleting a few edges) there are graphs with chromatic number and girth[6] as large as we like! In other words, there are tons of graphs out there that locally look like trees (i.e. locally look 2-chromatic) but in fact require a ton of colors to properly color.

In 1976, Nešetřil and Rödl created an explicit construction for such graphs. We state it at the end of the HW/recommend it only for rather interested students, as it's kind of complicated (and may scare off some people who would otherwise be fine in this course!)

So: the moral of this lecture is that having a large chromatic number can be a purely global phenomenon, and have very little to do with the local behavior of your graph. As graph theorists, what can we do about this? In our next lecture, we'll start with some ideas. . .

[6] A graph G has **girth** k if the length of the shortest cycle in G is k.