Topics in Graph Theory – 1

January 7 and 9, 2014

A *proper (vertex) colouring* of a graph G is an assignment of one colour to each vertex of G, so that adjacent vertices receive different colours. Given a colouring, the set of vertices receiving one particular colour is called a *colour class*. A graph that has a proper colouring with k colours is called a k*-colourable* graph.

The *chromatic number* of a graph G (notation: $\chi(G)$ is the minimum number of colours required for a proper colouring of the graph. To show that any particular graph G has a particular chromatic number $\chi(G) = k$ two aspects must be shown: (1) show that it is possible to colour G with k colours (this can be done by giving an explicit colouring, or by giving a proven colouring method), and (2) show that it is not possible to colour G with $k-1$ colours (this can be done by finding an induced subgraph that cannot be coloured by $k-1$ colours). A graph with chromatic number k is called a k*-chromatic graph.*

Example: the *circulant graph* $C(n;k)$ is defined of the graph with vertex set $\{0, 1, \ldots, n-1\}$, where vertex i is adjacent to vertices $i+1, \ldots i+k$, where addition is taken modulo n. Since vertices $0, 1, \ldots, k$ are all pairwise adjacent, we need at least $k+1$ colours. Assume each vertex i is assigned colour $i \mod (k+1)$. Is this a proper colouring? Assume vertex i and j are adjacent and receive the same colour. This means that $|j-i| \leq k \bmod n$, while $|j-i| = 0 \bmod k+1$. This can only happen if $k+1$ does not divide n. Thus, when $k+1$ divides n, this colouring is a proper colouring, and $\chi(n;k) = k+1$.

An *independent set* in a graph G is a set of vertices of G so that no two of them are adjacent. The *independence number* of G (notation $\alpha(G)$) is the size of the largest independent set in G.

A *clique* in a graph G is a set of vertices of G so that every two of them are adjacent. The *clique number* of G (notation $\omega(G)$) is the size of the largest clique in G.

We use $n(G)$ to denote the number of vertices of G, and $m(G)$ to denote the number of edges of G. We also use n and m if it is clear from the context which graph is referred to.

Theorem 1. *For every graph G, $\chi(G) \geq \omega(G)$ and $\chi(G) \geq n(G)/\alpha(G)$.*

Proof. In a proper colouring, every vertex in a clique must receive a different colour. This proves the lower bound on $\chi(G)$. For the upper bound, note that every colour class of a proper colouring must be an independent set, and thus has size at most $\alpha(G)$. Since every vertex must receive a colour, the result follows. $\square$

Open problem (Erdös-Faber-Lovász conjecture) *Suppose a graph is the union of k complete graphs of size k, where any two complete graphs have at least one vertex in common. Does this graph have chromatic number k?*

One of the most well-known colouring algorithm is greedy colouring. It is a *heuristic*, which means that it does not always lead to an optimal colouring. In fact, the gap between the number of colours used by the greedy algorithm and the chromatic number may be arbitrarily large.

The Greedy Colouring Heuristic
INPUT: a graph G
OUTPUT: a proper (vertex) colouring of G

1. Arrange the vertices of G in linear order $v_1, v_2, \ldots, v_n$.

2. Colour the vertices one by one in this order, assigning to v_i the smallest positive integer (colour) not assigned to one of its already-colouried neighbours.

An example where the greedy colouring can do very badly is the following. Let G be a bipartite graph with bipartition $X = \{x_1, \ldots, x_n\}$ and $Y = \{y_1, y_2, \ldots, y_n\}$, where for each i, x_i is adjacent to all vertices in Y *except* y_i. Now choose the ordering so that $v_{2i+1} = x_i$ and $v_{2i+2} = y_i$ for $i = 1, \ldots, n$. Then the greedy algorithm will assign both x_1 and y_1 colour 1, after which colour 1 cannot be used on any other vertex. Similarly, x_i and y_i will be assigned colour i. Thus, the greedy algorithm uses n colours, while only 2 colours are needed, since G is bipartite.

As the example above shows, the performance of the greedy algorithm depends heavily on the linear order chosen. In fact, there always exists at least one ordering that will give the optimal colouring. We use $[n]$ to denote the set $\{i : 1 \leq i \leq n\}$. Note that $[0] = \emptyset$.

Theorem 2. *For every graph G, there exists a linear order of the vertices such that the greedy algorithm uses $\chi(G)$ colours when following this order.*

Proof. Fix graph G, and let $k = \chi(G)$. Let $c : V(G) \to [k]$ be a proper colouring of G. Now order the vertices of G, $v_1, \ldots, v_n$ so that $c(v_i) \leq c(v_j)$ whenever $i < j$, and let $c^* : V(G) \to [k]$ be the greedy colouring obtained from using this order. We claim that for each i, $c^*(v_i) \leq c(v_i)$. To prove the claim, we use induction on i. Clearly, $c^*(v_1) = 1 \leq c(v_1)$. Now assume the claim holds for all values less than i, and assume that $c^*(v_i) = \ell$. By our choice of ordering, this means that the already-coloured neighbours of v_i have colours $1, \ldots, \ell - 1$. In particular, there must be a vertex v_j with $j < i$ which is adjacent to v_i and has $c^*(v_j) = \ell - 1$. By the induction hypothesis, $c(v_j) \geq c^*(v_j)$, and by our choice of colouring, $c(v_i) \geq c(v_j)$. Thus, $c(v_i) \geq \ell - 1$, and since v_i and v_j are adjacent and c is a proper colouring, this implies that $c(v_i) \geq \ell = c^*(v_i)$. This completes the proof of the claim. Using the claim, we see that the highest colour assigned by c^* must be at most k, and therefore c^* uses at most k colours. Since G is k-chromatic, c^* uses exactly k colours. □

Let $\Delta(G)$ be the maximum degree of G. Clearly the greedy colouring never uses more than $\Delta(G) + 1$ colours (since any vertex can be adjacent to at most $\Delta(G)$ colours). This leads to our first upper bound on the chromatic number.

Theorem 3. *For each graph G, $\chi(G) \leq \Delta(G) + 1$.*

For general graphs, we cannot do much better than that. Brook's Theorem (see any basic graph theory text) gives an upper bound of $\Delta(G)$ for all graphs for all graphs with a few simple exceptions.

Another upper bound can be obtained from the greedy colouring algorithm via the concept of *k-cores*. The k-core of a graph G is the largest induced subgraph of G where all vertices have degree at least k (this subgraph is unique. An algorithm to obtain the k-core is to successively remove vertices of degree less than k until no more such vertices can be found. The graph remaining after the algorithm stops is the k-core.

Theorem 4. *If the k-core of G is empty, then $\chi(G) \leq k$.*

Proof. Suppose the "peeling" described above is applied to obtain the k-core, and let $v_1, \ldots v_n$ be the order in which vertices are removed. Since a vertex is only removed when it has degree less than k in the remaining graph, we see that for each i, v_i has at most $k - 1$ neighbours in the graph induced by $v_{i+1}, \ldots, v_n$. Thus, we can apply the greedy colouring with the ordering

$v_n, \dots v_1$ (thus, the last removed vertex is coloured first), and will need at most k colours. $\square$

There are some classes for which the greedy colouring with a specific ordering always gives an optimal colouring. We mention bipartite graphs, and interval graphs.

Theorem 5. *Let G be a connected bipartite graph, and assume that the vertices are ordered so that, for all $j > 1$, v_j is adjacent to at least one vertex v_i with $i < j$. Then the greedy colouring uses the optimal two colours.*

Proof. Let G be a bipartite graph, and assume the vertices $v_1, \dots, v_n$ are ordered as stated. Assume, by contradiction, that the greedy algorithm is applied, and for some v_i a third colour is needed. Thus, v_i has two coloured neighbours, say v_k and v_ℓ ($k < i$ and $\ell < i$) so that v_k has colour 1 and v_ℓ colour 2. By the restrictions on the ordering, there is a path of coloured vertices from v_k leading to v_1 and back to v_ℓ. Since only two colours are used on this path, and both endpoints are of different colour, this path has odd length. Since v_ℓ and v_k are both adjacent to v_1, this means that G contains an odd cycle, and thus is not bipartite. This completes the proof. $\square$

Note that the theorem above can be applied to each connected component if the bipartite graph is not connected.

An *interval graph* is a graph $G = (V, E)$ where each vertex v corresponds to an interval $[a_v, b_v]$, and two vertices u and v are adjacent precisely when their intervals overlap, i.e. when $a_v \leq a_u \leq b_v$ or $a_u \leq a_v \leq b_u$.

Theorem 6. *Let G be an interval graph. If the vertices are ordered according to the start of their intervals, i.e. so that $v_i = [a_i, b_i]$ and $v_i < v_j$ if $a_i < a_j$, then the greedy colouring achieves an optimal colouring of G with $\omega(G)$ colours*

A graph G is *perfect* if, for each induced subgraph H of G, $\chi(G) = \omega(H)$.

Note that the classes of bipartite graphs and interval graphs are *closed under taking induced subgraphs*, i.e. any induced subgraph is also a member of the class. Thus, from the theorems above we immediately obtain the following result:

Theorem 7. *All interval graphs and all bipartite graphs are perfect.*

How far from perfect can a graph be, i.e. how large can the difference between clique number and chromatic number get? The following construction shows that this gap can be as large as you want.

Mycielski's construction. This construction shows how to construct in an iterative way, graphs G_k ($k \geq 2$) which are triangle-free (so their clique number is 2), but which have chromatic number k.

For $k = 2$, $G_2 = K_2$, the complete graph on two vertices. For $k \geq 2$, we construct G_{k+1} from G_k as follows: Let the vertices of G_k be $v_1, \ldots, v_n$. Add $n+1$ new vertices $u_1, \ldots, u_n$ and w. For $i = 1, \ldots n$, join u_i to every neighbour of v_i in G_k, and to w. This is the construction; we now must show that G_{k+1} so constructed is triangle-free and has chromatic number $k+1$.

First we show that G_{k+1} has no triangles. Note that $u_1, \ldots, u_n$ form an independent set, and w is only adjacent to vertices u_i. Also, by assumption, G_k is triangle-free. Thus, any triangle in G_{k+1} must consist of two vertices v_i and v_j, and a vertex u_k. However, u_k is only adjacent to the neighbours of v_k, so v_i and v_j must both be neighbours of v_k, which contradicts the fact that G_k contains no triangles.

Next we show that $\chi(G_{k+1}) = k+1$. The easy part is the upper bound, which we can find by defining a $k+1$ colouring: first, colour $v_1, \ldots, v_n$ with k colours (this can be done since, by assumption, $\chi(G_k) = k$). Then, give each u_i the same colour as v_i, and give w colour $k+1$. Since u_i is only adjacent to w and to all neighbours of v_i and has the same colour as v_i, this is a proper colouring.

To show that G_{k+1} cannot be coloured with k colours, we argue by contradiction. Suppose there is a colouring with k colours. This colouring is also a k-colouring of G_k. As seen earlier in problem set 1, for every colour j there must exist a vertex v_j of colour j whose neighbours have every colour except j. Now u_j has the same neighbours as v_j, so u_j must also have colour j. So the set $u_1, \ldots u_n$ contains vertices of every colour. But w is adjacent to every vertex u_i, which leads to a contradiction.