17th Lecture: 2/3

Definition 17.1. Let $G = (V, E)$ be a graph and k a positive integer. A *(vertex) k-coloring* of G is a function $f : V \to \{1, 2, \ldots, k\}$ such that $f(v_1) \neq f(v_2)$ whenever $\{v_1, v_2\} \in E$.

We say that $f(v)$ is the *color* assigned to vertex v. So, in words, a k-coloring of G is an assignment of colors to vertices so that at most k colors are used in all and vertices which are *neighbors*, i.e.: joined by an edge, always get different colors.

Definition 17.2. The *chromatic number* of a graph G, denoted $\chi(G)$, is the smallest k for which there exists a k-coloring of G.

Example 17.3. $\chi(K_n) = n$ since every pair of vertices are neighbors in K_n and so any coloring must use a different color at each vertex. This observation is generalised in Proposition 17.7 below.

Example 17.4. Let C_n denote a cycle of length $n \geq 3$. Then

$$\chi(C_n) = \begin{cases} 2, & \text{if } n \text{ is even,} \\ 3, & \text{if } n \text{ is odd.} \end{cases}$$

To see this, observe that if one tries to 2-color a cycle one vertex at a time then one will have no choice but to alternate back and forth between the two colors. This will work if n is even, but if n is odd, then we'll need to use a third color at the last vertex.

Remark 17.5. **(i)** To design a general algorithm for computing the chromatic number of a graph G, and/or to find an explicit $\chi(G)$-coloring of it, is another classic example of an NP-complete problem. So it's a a hard problem.
(ii) This problem is often popularised as the *Scheduling Problem*. Suppose you want to make a schedule of exams for courses at Chalmers in Lp3. The basic constraint is that you can't schedule two exams in the same time slot if there is at least one student taking both the courses involved. One would like to make a schedule involving as few slots as possible, simply because each slot costs money to run (rooms must be booked, guards employed etc.). This can be modeled as a graph coloring problem. One lets $G = (V, E)$, where V is the set of all courses which will have an exam this period, and an edge is placed between courses v_1 and v_2 if and only if there is at least one student taking both of them. Then to determine $\chi(G)$ is to determine the minimum number of slots needed in an exam schedule and to determine a specific such schedule is to determine an explicit $\chi(G)$-coloring of the graph. As noted in **(i)**, this problem is extremely hard for general graphs. However, for smallish graphs one can do a kind of simple trial-and-error and "hope for the best", since there is a very simple algorithm for finding *some* coloring, which if you're lucky may not use too many colors. This *greedy coloring algorithm* will be described in Theorem 17.19 below. In practice, one is usually satisfied with any coloring which uses "sufficiently few" colors, even if one doesn't know if it's optimal.

Lower bounds for the chromatic number. The most obvious lower bound is gotten from the following concept:

Definition 17.6. Let $G = (V, E)$ be a graph. The *clique number* of G, denoted $\omega(G)$, is the maximum number of vertices in a complete subgraph of G.

Proposition 17.7. *For any graph G one has*

$$\chi(G) \geq \omega(G). \tag{17.1}$$

PROOF: Obvious, since the vertices of a clique are all neighbors and must therefore all be assigned different colors in any coloring.

It is easy to find examples of graphs in which the inequality in (17.1) is strict. The simplest set of examples is odd cycles C_n, for which $\chi(C_n) = 3$ and $\omega(C_n) = 2$ (see Example 17.4 above). A similar family of examples are the *wheel graphs* W_n $(n \geq 3)$, which consist of an n-cycle all of whose nodes are also joined to a central $(n+1)$:st node (so the union of a C_n and a $K_{1,n}$). It is obvious that $\omega(W_n) = 3$. However, $\chi(W_n) = 4$ when n is odd since the outer n-cycle will then require three colors and a fourth color must be used at the central vertex.

The gap between $\omega(G)$ and $\chi(G)$ can be arbitrarily large. Before exploring this issue, we note that one can at least give a somewhat nice characterisation of graphs with chromatic number 2.

Theorem 17.8. *For a graph G with at least one edge, the following are equivalent:*
(i) G is bipartite
(ii) G has no cycles of odd length
(iii) $\chi(G) = 2$.

PROOF: (i) $\Rightarrow$ (ii) is obvious since, in a bipartite graph, any path must cross back and forth between the two sides and hence, after an odd number of steps we'd be at the opposite side from where we started and could not have completed a cycle. It is also pretty obvious that (i) $\Leftrightarrow$ (iii). For in a bipartite graph there are no edges between vertices on the same side and hence all such vertices can be given the same color. Conversely, if G can be 2-colored, then the vertices in each color form the two sides of a bipartition.

To complete the proof of the theorem, it thus suffices to show that (ii) $\Rightarrow$ (iii). So suppose G has no odd cycles. We may assume G is connected, since G can be 2-colored provided each of its connected components can. We will describe an explicit step-by-step procedure for 2-coloring G. Call the two available colors red and blue.

Step 0: Pick a vertex v at random and color it red.
Step 1: Color each neighbor of v blue. The point is that we can do this, since if any two of v's neighbors were themselves neighbors, then together with v they would form a 3-cycle, contradicting the assumption that G has no odd cycles.
...
Step k: At the k:th step we use the color not used at step $k-1$. The vertices colored at

this step are those which are neighbors of the vertices colored at step $k-1$ and which have not yet been colored. Equivalently, the vertices colored at step k are those whose graph distance[1] from the vertex v is exactly k. The point is that no two of these vertices are neighbors. For if two of them were, say v_1 and v_2, then G would possess a cycle of length $2l+1$ for some $1 \leq l \leq k$, got by taking a shortest path from v_1 back to the common ancestor of v_1 and v_2, followed by a shortest path back to v_2 and finally the edge $\{v_2, v_1\}$. This completes the proof.

Remark 17.9. In particular, Theorem 17.8 says that a graph with no cycles at all, and at least one edge, has chromatic number 2. Such graphs are called *forests* and a connected forest is a *tree*. We study such graphs in the next lecture.

The gap between clique number and chromatic number. The point of this subsection is to give some intuition as to why it is a hard (NP-complete) problem to compute $\chi(G)$ in general. There are two stand-out results which we present below, Theorem 17.11 and Corollary 17.16. Their common theme is that they indicate that the chromatic number is a "global" rather than a "local" property of a graph. In other words, even if a graph lacks "local bottlenecks", such as short cycles (Theorem 17.11) or large cliques (Corollary 17.16), vertices which are far apart can still influence one another in curious ways so as to make the chromatic number large.

Definition 17.10. The *girth* of a graph G, denoted $g(G)$, is the minimum length of a cycle in G. If G has no cycles, one sets $g(G) := +\infty$.

The following result is due to Erdős:

Theorem 17.11. *For each positive integer t, there exists a graph G such that*

$$\min\{\chi(G),\, g(G)\} \geq t. \tag{17.2}$$

The proof of this theorem is beyond the scope of the course though it's not "too hard". What is interesting is that Erdős' proof uses a *probabilistic method* and hence, in particular, he only proves the existence of graphs satisfying (17.2) for large t, and does not find any explicit examples. As far as I'm aware, it is still an open problem to construct such graphs explicitly, though Erdős' proof is 70+ years old ! If you're interested, you can find his proof in my 2020 lecture notes for MMG610 (file 16.E) or in the book

N. Alon and J. Spencer, *The Probabilistic Method*, 3rd edition (2008), Wiley.

For our purposes, the point is that any graph with girth at least 4 has clique number

[1] The *distance* between two vertices v and w in a graph G, denoted $d(v, w)$, is the minumum length of (i.e.: number of edges in) a path between them. If there is no path between them (in other words, if they lie in separate connected components of G), then $d(v, w) := +\infty$. Graph distance satisfies the abstract properties of a *distance function*, i.e.:
(i) $d(v, w) = 0$ if and only if $v = w$
(ii) Symmetry: $d(v, w) = d(w, v)$
(iii) Triangle inequality: $d(v, x) \leq d(v, w) + d(w, x)\ \forall\, v, w, x \in V(G)$.

2. So one can have $\omega(G) = 2$ and $\chi(G)$ arbitrarily large.

Definition 17.12. Let $G = (V, E)$ be a graph. A subset W of V is said to be an *independent set* if $\{w_1, w_2\} \notin E$ whenever $w_1, w_2 \in W$. In other words, there are no edges whatsoever between the vertices of W.

The *independence number* of G, denoted $\alpha(G)$, is the maximum number of vertices in an independent set.

Definition 17.13. Let $G = (V, E)$ be a graph. The *(graph) complement* of G, denoted $\overline{G}$, is the graph whose vertices are the same as those of G and such that $E(\overline{G}) = \binom{V}{2} \backslash E(G)$. In other words, the edges of $\overline{G}$ are precisely those which are absent from G.

Proposition 17.14. *Let $G = (V, E)$ be a graph. Then*
(i) $\alpha(G) = \omega(\overline{G})$.
(ii) $\chi(G) \geq \frac{|V|}{\alpha(G)}$.

PROOF: (i) is immediate from the definitions above. For (ii), note that in any coloring of G, the vertices assigned a fixed color must form an independent set. Hence, the number of colors used times the maximum size of an independent set cannot be less than the total number of vertices in the graph (since each vertex gets *some* color). In particular, $\chi(G)\alpha(G) \geq |V|$, v.s.v.

We now wish to prove a result which says that, "typically", $\chi(G)$ is much larger than $\omega(G)$. To do so, we first need to make precise what it means to choose an n-vertex graph uniformly at random.

Let $V = V_n = \{1, 2, \ldots, n\}$. For each of the $\binom{n}{2}$ pairs $\{i, j\}$ toss a fair coin and insert the edge $\{i, j\}$ if and only if it shows heads. This produces what's called an *(Erdős-Renyí) random graph* and is denoted $G(n, 1/2)$. What it is, more precisely, is a uniform distribution over all possible graphs with vertex set V_n. Note that we are distinguishing here between graphs which don't have exactly the same set of edges, even if they are isomorphic.

Theorem 17.15. *Let $\varepsilon > 0$. Let E_n be the event that $G(n, 1/2)$ contains a clique of size at least $(2 + \varepsilon)\log_2 n$. Then $\mathbb{P}(E_n) \to 0$ as $n \to \infty$.*

PROOF: Let $k = \lceil (2 + \varepsilon)\log_2 n \rceil$. Since $\frac{(2+\varepsilon)^2}{2} = 2 + 2\varepsilon + \frac{\varepsilon^2}{2} > 2 + 2\varepsilon$, if n is sufficiently large then $\frac{k(k-1)}{2} > (2 + 2\varepsilon)(\log_2 n)^2$. We have

$$E_n = \bigcup_{i=1}^{\binom{n}{k}} E_{n,\,i},$$

where the $\binom{n}{k}$ subsets of $V(G(n,\,1/2))$ are ordered arbitrarily, and $E_{n,\,i}$ is the event that the subgraph induced by the i:th subset is a clique. By MP,

$$\mathbb{P}(E_{n,\,i}) = \left(\frac{1}{2}\right)^{\binom{k}{2}}$$

and hence, when n is sufficiently large,

$$\mathbb{P}(E_{n,\,i}) < 2^{-(2+2\varepsilon)(\log_2 n)^2}.$$

Thus, for n sufficiently large,

$$\mathbb{P}(E_n) \leq \sum_{i=1}^{\binom{n}{k}} \mathbb{P}(E_{n,\,i}) < \binom{n}{k} 2^{-(2+2\varepsilon)(\log_2 n)^2} <$$

$$< n^k (2^{\log_2 n})^{-(2+2\varepsilon)(\log_2 n)} \leq n^{1+(2+\varepsilon)(\log_2 n)} n^{-(2+2\varepsilon)(\log_2 n)} = n^{1-\varepsilon(\log_2 n)},$$

which goes to zero as $n \to \infty$, for any fixed $\varepsilon > 0$.

Corollary 17.16. *Let $\varepsilon > 0$. Then, as $n \to \infty$,*

$$\mathbb{P}(\omega(G(n,\,1/2)) > (2+\varepsilon)(\log_2 n)) \to 0, \tag{17.3}$$

$$\mathbb{P}\left(\chi(G(n,\,1/2)) < \frac{n}{(2+\varepsilon)(\log_2 n)}\right) \to 0. \tag{17.4}$$

Proof: Eq. (17.3) is an immediate consequence of Theorem 17.15. For (17.4), first observe that Proposition 17.14 implies that

$$\chi(G(n,\,1/2)) < \frac{n}{(2+\varepsilon)(\log_2 n)} \Leftrightarrow \alpha(G(n,\,1/2)) > (2+\varepsilon)(\log_2 n) \Leftrightarrow \omega(\overline{G(n,\,1/2)}) > (2+\varepsilon)(\log_2 n).$$

But the random graph $G(n,\,1/2)$ and its complement have the same distribution (since each edge appears with probability $1/2$ in each, independent of all other edges), which implies that the events $\omega(G(n,\,1/2)) > (2+\varepsilon)(\log_2 n)$ and $\omega(\overline{G(n,\,1/2)}) > (2+\varepsilon)(\log_2 n)$ have the same probability. The probability of the former goes to zero, by (17.3), so we're done.

To summarise, Corollary 17.16 says that a graph on n vertices will typically have clique number no larger than about $2\log_2 n$, whereas its chromatic number will typically be about $\frac{n}{2\log_2 n}$ or larger. When n is large, this is way bigger than $2\log_2 n$ so, typically, the chromatic number is way bigger than the clique number.

Upper bounds for the chromatic number. For a graph $G = (V,\, E)$ denote

$$\Delta(G) := \max_{v \in V} \deg(v).$$

Definition 17.18. Let $k \in \mathbb{Z}_+$. A graph G is said to be *regular of degree* k, or just k*-regular* for brevity, if every $v \in V$ satisfies $\deg(v) = k$.

Theorem 17.19. *(i) For any graph G one has* $\chi(G) \leq \Delta(G) + 1$.
(ii) Moreover, if G is connected and not regular of degree $\Delta(G)$*, then* $\chi(G) \leq \Delta(G)$.

PROOF: The theorem follows from an analysis of the so-called *greedy coloring algorithm*, which we have already hinted at in several places (e.g.: Example 17.4 and the proof of Theorem 17.8).

(i): Order the vertices arbitrarily, say $v_1,\, v_2,\, \ldots,\, v_n$. Denote the available colors as $c_1,\, c_2,\, \ldots,\, c_{\Delta(G)+1}$. We will show how to color G using the available colors. The algorithm is:

Color the vertices in order, at each step using the smallest color not already used on one of the previously colored neighbors of the current vertex.

To see that this works, consider a general step k. At this step we color v_k. Some of its neighbors may have already been colored, but by definition v_k has no more than $\Delta(G)$ neighbors, hence no more than $\Delta(G)$ colors can have already been used on its previously colored neighbors. So we'll be able to color v_k with one of the $\Delta(G) + 1$ available colors.

(ii): The same greedy algorithm will work if we are careful to order the vertices in such a way that

For every vertex, at most $\Delta(G) - 1$ *of its neighbors appear before it in the ordering.*

So it remains to prove that such an ordering exists. We create the list "backwards". For v_n choose any vertex of degree strictly less than $\Delta(G)$ - our assumption in part (ii) is that at least one such vertex exists. If v_n has r neighbors, say, then put these as $v_{n-1},\, \ldots,\, v_{n-r}$, in any order. Proceed in this way: at each step prepend to the list, in any internal order, the neighbors of the vertices prepended at the previous step which are not yet listed.

First of all, since G is connected, every vertex will eventually be listed, thus giving a complete ordering of all n vertices. The procedure ensures that every vertex, apart from v_n, has at least one neighbor appearing after it on the list, hence at most $\Delta(G) - 1$ neighbors appearing before it. But this is also true of v_n, since it has degree at most $\Delta(G) - 1$ in the first place.

Remark 17.20. The chromatic number of a graph may be much less than its maximum degree. For example, a star graph $K_{1,\,n}$ is bipartite and hence has chromatic number 2 for any n, whereas the maximum degree is n, hence arbitrarily large. See also Theorem 17.21 below.

On the other hand, the number of colors utilised by the greedy algorithm obviously depends, in general, on how the vertices are ordered. In particular, there is always *some* ordering of the vertices of a graph such that the greedy algorithm will yield an optimal coloring. For consider any coloring with $\chi(G)$ colors. Order the vertices such that all those with color 1 come first, then all those with color 2 and so on. If we applied the greedy algorithm with this ordering then it would also require only $\chi(G)$ colors. Note that it would not necessarily produce exactly the same coloring as we started with, since it always gives a vertex the first available color and thus it is possible some vertices would receive a lower numbered color than originally.

OK, so the greedy algorithm could in principle be used to compute $\chi(G)$, by simply testing all possible orderings of the vertices. But there are $n!$ possible orderings, so this does not yield an efficient (polynomial-time) procedure.

We finish by noting a very famous result which gives a striking upper bound on the chromatic number of planar graphs:

Theorem 17.21. (Four-Color Theorem) *If G is a planar graph then $\chi(G) \leq 4$.*

This was first proven in the mid-1970s, though the bound $\chi(G) \leq 5$ had been known for a long time and it is easy (see Demo5, Exercise 3) to prove the bound $\chi(G) \leq 6$. One reason for the fame of this theorem is that it is generally considered to be the first proof of a major result which used electronic computing power in a crucial way. The idea of the proof was a (complicated) kind of induction on the number of vertices in the graph, whereby the authors eventually managed to reduce the proof to checking several thousand specific graphs. There were far too many to check by hand, but by the 1970s computers could manage it. Over the years, various simpler proof of the Four Color Theorem have appeared, though as far as I know every one still reduces eventually to checking a few thousand or so specific graphs.