LTCC Course: Graph Theory 2019/20
§1: Graph colouring

Peter Allen

13 January 2020

More or less everything we will be discussing in the lectures and these notes can be found in great detail in Chapter 5 of Diestel (in particular Section 5.4). Also Chapter V of Bollobás (and, again, in particular Section V.4) has everything we need.

The best source for anything related to graph colouring is the book *Graph Coloring Problems*, by T.R. Jensen and B. Toft, John-Wiley (1995).

Vertex Colouring

Given a graph $G = (V, E)$ and a finite *colour set* C, a (*proper*) *vertex colouring of* G is a function $\varphi : V \to C$ so that for all edges $uv \in E$ we have $\varphi(u) \neq \varphi(v)$. The *chromatic number* $\chi(G)$ of G is the smallest number of colours a colour set must have so that a vertex colouring exists.

We usually use positive integers for the colours, $C = \{1, \dots, k\}$ for some k. And then we can say that $\chi(G)$ is the minimum k so that there is a vertex colouring $\varphi : V \to \{1, \dots, k\}$. If G admits a vertex colouring with k colours, we say G is *k-colourable*. So saying that G is k-colourable is equivalent to stating that $\chi(H) \leq k$.

Many of the bounds on the chromatic number are in terms of vertex degrees. For example, the obvious greedy algorithm gives a proof that $\chi(G) \leq \Delta(G) + 1$. We define the *degeneracy* $\deg(G)$ *of* G to be the maximum of $\delta(H)$ over all subgraphs H of G. It is obvious that $\deg(G) \leq \Delta(G)$ holds for all graphs G.

Proposition 1. *For all graphs G we have $\chi(G) \leq \deg(G) + 1$.*

Proof. Suppose $\deg(G) = m$ and G has n vertices. We begin by finding an order $V(G) = \{v_1, \dots, v_n\}$ such that for each $1 \leq i \leq n$, the vertex v_i has at most m neighbours in $\{v_1, \dots, v_{i-1}\}$. For each $i = n, n-1, \dots, 1$ in that order, choose v_i to be a vertex in $V(G) \setminus \{v_{i+1}, \dots, v_n\}$ with least neighbours in $V(G) \setminus \{v_{i+1}, \dots, v_n\}$. Since $G\big[V(G) \setminus \{v_{i+1}, \dots, v_n\}\big]$ is a subgraph of G, by definition of degeneracy v_i has at most m neighbours in $V(G) \setminus \{v_{i+1}, \dots, v_n\}$.

Now we colour $V(G)$ greedily in the order we found: for each $i = 1, \dots, n$ in that order, colour v_i with the smallest number colour not used on any neighbour of v_i in $\{v_1, \dots, v_{i-1}\}$. There are at most m colours used on these vertices, so the smallest number colour not used is at most $m + 1$, as desired. □

This is a proof by algorithm—it finds the desired colouring in polynomial time (polynomial in n), and the algorithm actually also finds the degeneracy of G (check you see why this is true). Graphs such as the complete bipartite graphs $K_{n,n}$, consisting of $2n$ vertices divided into two equal parts, with the edge set consisting of all crossing pairs, show that the degeneracy can be much larger than the chromatic number, though: $K_{n,n}$ has chromatic number two (colour one part with colour one and the other with colour two) and degeneracy n.

In the other direction, a clique K_t (obviously) needs t colours in a proper colouring, and it has maximum degree and degeneracy equal to $t-1$, so Proposition 1 can be sharp. With some effort, you can say a little more.

Theorem 2 (Brooks, 1941)**.** *A connected graph G that is not a complete graph or an odd cycle satisfies $\chi(G) \le \Delta(G)$.*

Proof. We'll only give a sketch of the proof to highlight the main idea. First, the case $\Delta(G) = 2$ is easy: such a connected graph G can only be a path or a cycle; paths and even cycles are 2-colourable, which leaves only the odd cycles.

For $\Delta(G) = \Delta \ge 3$, the only exceptional graph is supposed to be $K_{\Delta+1}$. First, suppose G is not regular. Then there is a vertex v_n with degree at most $\Delta - 1$. We order $V(G)$ by choosing, for each $i = n-1, \ldots, 1$, a vertex v_i with at least one neighbour in $\{v_{i+1}, \ldots, v_n\}$; this is possible because G is connected. Now colour $V(G)$ greedily (as in Proposition 1) in this order, from v_1 to v_n. For each i, the vertex v_i has at most $\Delta - 1$ neighbours preceding it in the order: for $i < n$ because v_i has at most $\Delta(G)$ neighbours in total, at least one of which comes later in the order, and for $i = n$ by choice of v_n. So the greedy colouring uses at most $\Delta - 1 + 1 = \Delta$ colours.

Next, we'd like to make something like this trick work when G is regular. Pick a vertex v_n of $V(G)$ arbitrarily. If all pairs of its neighbours are adjacent, then $G = K_{\Delta+1}$, so we can assume some pair v_1, v_2 are not adjacent. We try to repeat more or less the proof above: for each $i = n-1, \ldots, 3$ in that order, let v_i be a vertex with at least one neighbour in $\{v_{i+1}, \ldots, v_n\}$, and colour $V(G)$ using the greedy strategy in this order. If this is possible (meaning: if it's possible to create this vertex order) then both v_1 and v_2 get colour one (because they're not adjacent), we can colour each vertex $v_3, \ldots, v_{n-1}$ with a colour at most Δ for the same reason as above, and we can colour v_n with a colour at most Δ because although it has Δ neighbours, two have the same colour (namely v_1 and v_2) and so at most $\Delta - 1$ colours are used on its neighbourhood. In other words, we're done.

The only problem is: what if the order we asked for doesn't exist? This happens if and only if $G - \{v_1, v_2\}$ is a disconnected graph. It's fairly easy to prove that if for some v the graph $G - \{v\}$ is disconnected, then we can colour it with Δ colours. Let C be a component of $G - \{v\}$; let $G_1 = G[C \cup \{v\}]$ and $G_2 = G - C$. By induction we can colour both G_1 and G_2 with $\Delta(G)$ colours, and permuting the colours on one if necessary we can ask that the colourings agree on v; and then they give a colouring of G.

What's left is the case that $G - \{v\}$ is connected for every v, but $G - \{v_1, v_2\}$ is disconnected. It is a good exercise (Exercise 1) to show that this case can also be handled by induction. If your proof is a simple generalisation of the proof above dealing with the case $G - \{v\}$ is disconnected, then you missed something, though. $\square$

These results (more or less obviously) come with algorithms that actually give colourings. But in general colouring a graph G optimally is a hard algorithmic problem. It is an easy exercise to give an algorithm which, for input G, either returns a 2-colouring of G or a witness that this is not possible in the form of a cycle of odd length in G. But if 2 is replaced by 3 (or any larger number), then there is no known algorithm which performs much better than the brute-force approach of simply testing all maps $c : V(G) \to \{1, 2, 3\}$ to see if one is a proper colouring. This is one of the (many) incarnations of the famous P versus NP problem (lecture 3); most (but not all) mathematicians believe that there is in fact no polynomial-time (polynomial in $v(G)$) algorithm for this problem. This is something like a reason why problems in the area of graph colouring tend to be difficult: you want to prove a non-trivial result relating to perhaps $\chi(G)$, but if your approach would imply an algorithm which gives a $\chi(G)$-colouring of G, it is unlikely to work.

Much research on graph colouring has focused on finding the best upper bound on the chromatic number of planar graphs. A *planar graph* is a graph that can be drawn in the plane so that the vertices are different points, the edges are simple curves connecting their endvertices, and the interior of each edge contains no vertex and no point from any other edge. We will see a lot more about planar graphs (and more general related concepts) in the next lecture.

The Four Colour Conjecture can be stated as "every planar graph has chromatic number at most 4". It was originally formulated as a conjecture on the number of colours needed to colour a map of contiguous regions in the plane, by Francis Guthrie in 1852. Francis asked the question to his brother Frederick, who was at that time a student of De Morgan at University College London. De Morgan started mentioning the problem to mathematicians he communicated with.

From Euler's Formula (see the next lecture), it is easy to obtain that a planar graph $G = (V, E)$ on at least three vertices satisfies $|E| \leq 3|V| - 6$. Since $\sum_{v \in V} d(v) = 2|E|$, from this we can derive that a planar graph always has a vertex of degree at most five. So we have $\deg(G) \leq 5$, and hence Proposition 1 guarantees that every planar graph G is 6-colourable.

A next, non-trivial, improvement was obtained by Heawood, using ideas from Kempe's incorrect proof of the Four Colour Conjecture.

Theorem 3 (Heawood, 1890; Kempe, 1879)**.** *All planar graphs are* 5*-colourable.*

Proof. We prove this by contradiction. Suppose there is a planar graph which is not 5-colourable, and let G be such a graph with fewest vertices. Now G has a vertex v of degree at most 5, since $e(G) \leq 3v(G) - 6$. Fix a plane drawing of G. Let c be a 5-colouring of $G - v$, which exists by minimality of G. Now the only reason why we cannot extend c to a 5-colouring of G is that v has five neighbours on which c uses all five colours; without loss of generality we can assume the neighbours $u_1, \ldots, u_5$ of v, listed in clockwise order around v, get colours 1 to 5 in that order.

For any given $1 \leq i < j \leq 5$, consider the subgraph G_{ij} of $G - v$ induced by the vertices getting colours i and j. If u_i and u_j are in different connected components of this graph, then we can swap the colours i and j on all vertices of G_{ij} containing u_i, in order to obtain a 5-colouring c' of $G - v$ in which both u_i and u_j are coloured j; this colouring extends to a 5-colouring of G.

But if u_1 and u_3 are in the same component of G_{13}, then there is a path in G_{13} from u_1 to u_3 (a *Kempe chain*). Adding the vertex v to this path, we obtain a cycle C in G. This cycle separates the plane into two regions, one containing u_2 and the other containing u_4 and u_5. Because we have a plane drawing of G, it follows that in particular u_2 and u_4 are in different connected components of G_{24}, so G is 5-colourable. $\square$

It's worth commenting on why this approach doesn't prove the Four Colour Conjecture. Briefly, if we try to repeat this proof with four colours rather than five, we see that when we colour $G - v$, all four colours have to appear in the neighbourhood of v. If v only has four neighbours, then the same Kempe chain method allows us to rearrange colours and we can get a 4-colouring. If v has five neighbours, though, some colour (say 1) appears twice in the neighbourhood of v. We can assume (see Exercise 4) that all faces in the plane drawing of G are triangles, because otherwise we can add edges to G; that can only make it harder to colour. Then the neighbours of v form a cycle, and the colours round that cycle are without loss of generality $1, 2, 1, 3, 4$ in that order. The difficulty now is that in both G_{13} and G_{14} the obvious Kempe chains don't separate the neighbours of v with colours 2 and 4 or 2 and 3 (respectively). By looking a bit more at the structure, you can find more complicated chains which allow you to deal with most such cases; but it's easy to miss a case that isn't covered (which is what Kempe did). See `http://web.stonehill.edu/compsci/lc/four-color/four-color.htm` for the details.

The search for a proof of the Four Colour Conjecture that all planar graphs are 4-colourable eventually led to a proof.

Theorem 4 (Appel & Haken, 1977)**.** *All planar graphs are* 4*-colourable.*

Appel and Haken's proof was a bit controversial, since it relied on computer assistance in checking a very large number of cases, and at the time it wasn't realistically possible for other mathematicians to independently verify the computer part of the proof; the program they used was optimised (by Koch) for specific hardware which could have had a bug (if this sounds far-fetched, a bug in the original Pentium chip did exist and was found by Nicely when an inconsistency appeared in computer searches for enumerations of twin, triplet and quadruplet primes) and rewriting or using other hardware was prohibitively expensive (supercomputer time isn't free). A more recent proof of the Four Colour Theorem by Robertson, Sanders, Seymour and Thomas reduces the need for computer checking and uses a simpler (checkable) program which can be run fairly fast on any modern computer; today the result is not really debated, though no human-checkable proof exists.

One can generalise all these problems to graphs which can be drawn on surfaces other than the plane without crossings: for example graphs which can be drawn on the torus, and so on. One can always use an argument similar to Heawood's (and in fact Heawood himself gave it) to get an upper bound for the chromatic number of surfaces of any given Euler characteristic (see next lecture for the definition); this bound grows as the surface gets more complicated. You might think, given that Heawood's Theorem 3 is off by one from the best possible answer, that this would be true for more complicated surfaces too: but in fact it usually isn't. Ringel and Youngs showed that for all surfaces other than the sphere (which is basically the same as the plane; more accurately, the sphere is the plane together with a single-point compactification at infinity, and reversing this we can always remove a point not used in the drawing of a given graph on the sphere) and the Klein bottle, Heawood's bound is best possible. Graphs on the Klein bottle can have chromatic number 6, but 7 (which is Heawood's bound for this surface) turns out not to be possible. These results are much easier than Theorem 4. More or less, this happens because n-vertex graphs on any given surface have at most $3n + C$ edges for some C depending on the surface (this is again from Euler's formula) and therefore most vertices have degree at most 6. It's easy to colour them if you have 7 colours to choose from (the case for all surfaces except the sphere and Klein bottle), so the high chromatic number can only come from the (at most constant) number of vertices of degree more than 6; we might as well delete them from the graph and assume every vertex has degree at least 7. It turns out that the best one can do is simply to make the high degree vertices form a clique (this needs a proof!), and what Ringel and Youngs did is to show that the obvious bound from Euler's formula on the maximum size of a clique in a graph on any given surface (except the sphere and Klein bottle) is actually attainable.

Vertex List Colouring

Given a finite set of colours C, a *list assignment* is an assignment $L : V \to \mathcal{P}(C)$ of subsets of the colours to the vertices. So each vertex v gets attached to it a list $L(v)$ of colours. Given a list assignment L, we call G L*-colourable* if there exists a function $\varphi : V \to C$ so that $\varphi(v) \in L(v)$ for all vertices v and such that for all edges $uv \in E$ we have $\varphi(u) \neq \varphi(v)$.

We say that G is k*-list-colourable* or k*-choosable* if G is L-colourable for every list assignment L with $|L(v)| = k$ for all vertices v. And the *list chromatic number* or *choice number* $\mathrm{ch}(G)$ is the smallest k so that G is k-choosable.

By giving all vertices the same list $L(v) = \{1, \ldots, k\}$ of colours, it follows directly that $\mathrm{ch}(G) \geq \chi(G)$. And somehow one would expect that the case when all colours have the same list is the "hardest" to colour, that cases where the lists are not identical are "easier". Surprisingly enough, that is not the case:

Proposition 5. *For all $k \geq 2$, the graph $G = K_{k,k^k}$ (which is 2-colourable) has $\mathrm{ch}(G) > k$.*

Proof. We use k^2 colours, the pairs of integers $[k] \times [k]$. Let the small part of G have vertices $x_1, \ldots, x_k$. We let the list of x_i be $\big\{(i,1),(i,2),\ldots,(i,k)\big\}$ for each i. Now there are k^k sequences of integers $a_1, \ldots, a_k$ with $1 \leq a_j \leq k$ for each j; so for each such sequence we can have a vertex in the large part of G with list $\big\{(1,a_1),(2,a_2),\ldots,(k,a_k)\big\}$. However we colour the vertices $x_1, \ldots, x_k$ from their lists, we will use all the colours on (exactly) one list in the large part. $\square$

An easy positive result is that Proposition 1 works if $\chi(G)$ is replaced by $\mathrm{ch}(G)$; the proof is essentially identical.

Proposition 6. *For all graphs G we have* $\mathrm{ch}(G) \leq \deg(G) + 1$. $\square$

As soon as the concept of list colouring was introduced (usually attributed to Vizing, 1976), determining the choice number of planar graphs became a hot topic. Since there are planar graphs that are not 3-colourable, there are also planar graphs that are not 3-choosable. And the argument of Proposition 6 guarantees that for every planar graph G we have $\mathrm{ch}(G) \leq 6$. So for quite a while it was an open question what the exact number should be. This question was settled in the 1990s with surprisingly simple proofs.

Theorem 7 (Voigt, 1993)**.** *There exist planar graphs that are not 4-choosable.*

Voigt gave a 238-vertex planar graphs with lists of size four from which it cannot be coloured properly; the construction was simplified by Mirzakhani (the late Fields medallist, in one of her first papers while still at school) in 1996 to a 63-vertex graph. We'll describe a variant with 69 vertices.

Proof. We begin by giving a 17-vertex planar graph H and some lists of colours (taken from the set of colours 1 to 5) with the following property: any proper colouring from the lists uses colour 5 somewhere on the outer cycle.

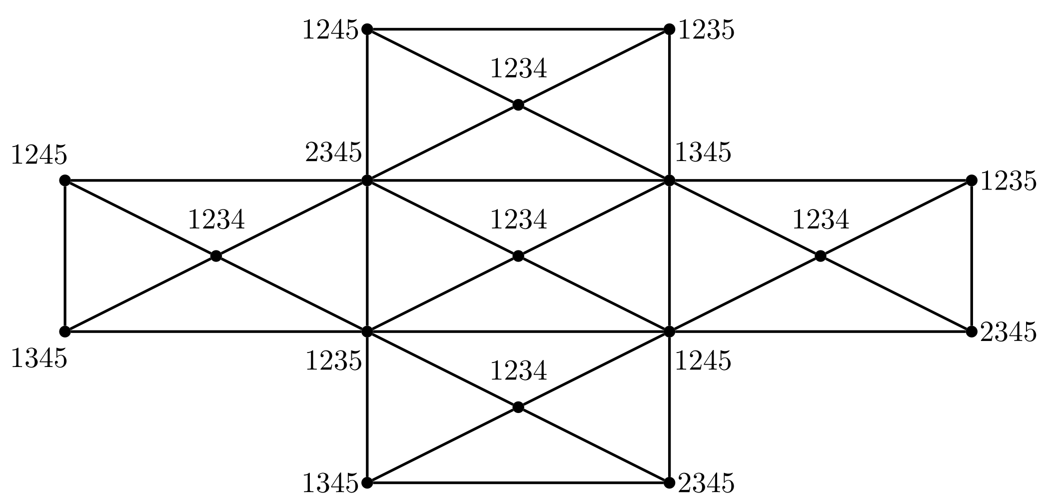

Checking that H indeed cannot be coloured properly from the lists without using 5 on the outer face is an easy case analysis. One way is to consider the colours of the two upper vertices in the central square. If the left uses 4 and the right 3, then in the upper square 5 has to be used. If the right uses 4 and the left 3 then in the middle square the colour 5 has to be used. By symmetry, we cannot use 1 and 2 to colour the bottom two vertices of the central square, nor 2 and 3 for the left two, nor 1 and 4 for the right two.

Now there are two cases left. Suppose we do not use 5 on the outer face. First, suppose we colour the top right vertex of the central square with 1. Then, as we just showed, we have to use 2 on the bottom right vertex, so 3 on the bottom left, so 4 on the top left—and now the central vertex is not colourable. Second, if we don't colour the top right vertex with colour 1, then it gets colour 3 or 4; so the top left vertex has to get colour 2. We can now deduce that the bottom left, bottom right and top right vertices have colours 1, 4, 3 in that order, and again the central vertex is not colourable.

To go from this graph H to the desired counterexample, we take four disjoint copies of this graph, called H_1, H_2, H_3 and H_4. We use the same lists on each graph as in H, *except* that we swap colours 1 and 5 in the lists for H_1, we swap 2 and 5 in the lists for H_2, and so on. By construction, any proper colouring of the resulting graph from its lists has to use colour i on the outer cycle of H_i for each i. We add one more vertex v adjacent to all vertices on all four outer cycles and give it the list $\{1, 2, 3, 4\}$; now the result cannot be properly coloured from its lists, and it's easy to see that it is planar. □

Note that this construction actually has chromatic number 3 (Voigt's construction is not 3-colourable). Mirzakhani was able to save a few vertices by noticing that one can identify a few vertices of some graphs H_i and H_j (on the outer faces with the same lists).

Giving a matching upper bound, Thomassen proved:

Theorem 8 (Thomassen, 1994)**.** *All planar graphs are 5-choosable.*

Thomassen's proof is stunningly simple and subtle at the same time. Before we describe its main idea, we need one more concept. A planar drawing of a planar graph in the plane (called, slightly confusingly, a *plane graph*) divides the plane into a number of connected "regions". Such a region is called a *face* of the plane graph. All faces have a finite size, except the one on the outside, which is usually called the *outer face*.

To prove Theorem 8, we only need to consider planar graphs that are connected and which can be drawn in the plane such that each face is bounded by exactly three edges – we call such a face a *triangle*. Indeed, if a planar graph has a face with more than three edges on its boundary, than we can add an edge so that the graph remains planar (see Exercise 4). Any colouring of this larger graph is certainly a colouring of the original graph.

A *near-triangulation* is an embedding of a connected planar graph in which each face, except for the outside face, is a triangle, and the outside face is a simple cycle (no repeated vertices). Thomassen then proves the following for near-triangulations :

Theorem 9 (Thomassen, 1994)**.** *Let G be a planar graph that can be drawn in the plane as a near-triangulation, and suppose that $v_1v_2 \ldots v_kv_1$ is the sequence of vertices encountered when walking along the boundary of the outside face. Suppose that the vertices of G have been assigned lists of colours with the following properties:*

– *v_1 and v_2 have a list of one colour each, with $L(v_1) \neq L(v_2)$;*

– *the other vertices on the outer face have a list of three colours;*

– *all vertices not on the outer face have a list of five colours.*

Then there exists a colouring of G where each vertex receives a colour from its list.

Note that the lists of colours given to each vertex according to Theorem 9 is always at most five. This means that Theorem 8 is a direct consequence of Theorem 9.

The point of the complicated Theorem 9 is that it is easy to prove by induction on $v(G)$; most of the work is to find the statement of Theorem 9.

Proof of Theorem 9. We prove the theorem by induction on $v(G)$. The base case $v(G) \leq 3$ is trivial, so suppose G is a minimal counterexample with at least four vertices.

Let C be the cycle in G bounding the outer face. First suppose that C has a chord; that is, there is an edge of G between two vertices of C which is not an edge of C. Then G is split into two smaller graphs G_1 and G_2 by this chord which are near-triangulations and whose outer faces are the two parts of C together with the chord. At least one of these graphs, say G_1, contains both v_1 and v_2, so by induction we can colour it properly from its lists. This fixes the colours of both vertices of the chord in G_2, but the remaining vertices of the outer face of G_2 are not v_1 nor v_2 so have lists of size at least three; by induction we can complete the colouring.

Second, suppose C has no chord. Let the neighbours of v_k be $v_{k-1}, u_1, u_2, \ldots, u_m, v_1$ in order around v_k. Since the interior faces are all triangles, these vertices are the vertices of a path in G from v_{k-1} to v_1. Since C has no chord, we have $m \geq 1$ and the vertices $u_1, \ldots, u_m$ are not on C.

We now use induction to colour $G - v_k$ and extend the colouring to v_k. Let x and y be two colours in $L(v_k)$ which are not in $L(v_1)$. We change the lists of $u_1, \ldots, u_m$ by removing x and y from these lists; since they originally had lists of size five, the resulting lists are of size three and so the induction statement applies; we can colour $G - v_k$ from the new lists. Now neither x nor y is used on any of $u_1, \ldots, u_m, v_1$ in this colouring, and at most one is used on v_{k-1}; so we can choose the other to colour v_k as desired. □

Edge Colouring

In this section we allow graphs to have *multiple edges* (but still no loops). When we want to exclude multiple edges, we use the term *simple graph*.

Given a graph $G = (V, E)$ and a finite *colour set* C, a (*proper*) *edge colouring of* G is a function $\varphi : E \to C$ so that for any two adjacent edges $e, f \in E$ (i.e., e and f have at least one common endvertex) we have $\varphi(e) \neq \varphi(f)$.

We use terminology similar to vertex colouring. So a graph G can be *k-edge-colourable*, and the minimum k for which G is k-edge-colourable is the *edge chromatic number* or *chromatic index* $\chi'(G)$. For a graph $G = (V, E)$, the *line graph* $L(G) = (V_L, E_L)$ is the graph that has the edges of G as vertices: $V_L = E$; and two edges are adjacent in the line graph if they have a common endvertex in G.

It is easy to see that edge colouring a graph G is the same as vertex colouring the line graph $L(G)$. So in that sense, edge colouring is just a special case of vertex colouring. But by arguing in this way, we ignore the special structural properties of the edge set of a graph that can be used when analysing edge colouring, but are not present in vertex colouring in general.

Another way to look at this is that a vertex-colouring of G is a partition of $V(G)$ into independent sets; an edge-colouring of G is a partition of $E(G)$ into matchings.

For a graph with multiple edges, we count the degree of a vertex as the number of *edges* incident with that vertex. So $\Delta(G)$ is the maximum number of edges incident with any vertex of G.

It is obvious that $\chi'(G) \geq \Delta(G)$. Since any edge can be adjacent to at most $2\,(\Delta(G) - 1)$ other edges, this gives the following trivial bounds for the chromatic index.

Proposition 10. *For all graphs G we have $\Delta(G) \leq \chi'(G) \leq 2\,\Delta(G) - 1$.*

But in fact, the relationship between chromatic index and maximum degree is much closer than the relationship between chromatic number and maximum degree, as the following classical results illustrate.

Theorem 11 (Kőnig, 1916)**.** *For a bipartite graph G we have $\chi'(G) = \Delta(G)$.*

Proof. There are two steps to proving this. First, we show that any bipartite graph G is a subgraph of some bipartite H all of whose vertices have degree $\Delta(G)$. Second, a well-known consequence of Hall's theorem is that any regular bipartite graph has a perfect matching. So we can simply remove sequentially perfect matchings from H (which maintains regularity) until no edges remain; this gives a proper $\Delta(G)$-edge-colouring of H, which in particular contains the edge-colouring of G we want.

One way to do the first step (not particularly efficient, but easy) is as follows. Suppose G is not regular. We first add isolated vertices to G if necessary to make it balanced (same number of vertices in each part). Think of the bipartition of G as top and bottom. We take a second identical copy of G with the top and bottom parts swapped. We join any vertex whose degree in G is less than $\Delta(G)$ to its copy. The result is a graph G' which contains G and which has $\delta(G') = \delta(G) + 1$. Now we repeat this until we have a regular graph.

Another much simpler, way is: add vertices to G to make it balanced if necessary, then add edges to make it regular. The difference is that the first way gives a simple graph; this way does not (it might create multiple edges). But for the second part this is not a problem!

For the second part, we need Hall's theorem. If you don't know this, Wikipedia gives one of the several known proofs. The statement is: A bipartite graph F with parts X and Y has a matching covering X if and only if for every set $S \subseteq X$, there are at least $|S|$ vertices in Y with a neighbour in S.

To go from this to showing that for $k \geq 1$ any k-regular bipartite graph has a perfect matching, consider any set $S \subseteq X$. Let T be the set of vertices in Y with at least one neighbour in S. Since all $k|S|$ edges leaving S go to T, on average a vertex in T has degree at least $k|S|/|T|$, and in particular there is a vertex with degree at least $k|S|/|T|$. Since every vertex has degree k, we have $|T| \geq |S|$. It follows that there is a matching covering X. By symmetry, we have $|X| = |Y|$ so the matching is perfect. $\square$

Theorem 12 (Shannon, 1949)**.** *For a graph G we have $\chi'(G) \leq \frac{3}{2}\,\Delta(G)$.*

Note that Theorem 12 is best possible. Form a graph G on three vertices x, y, z by adding m multiple edges between each pair from x, y, z. Then $\Delta(G) = 2\,m$, but all $3\,m$ edges are adjacent, hence $\chi'(G) = 3\,m$.

It is no accident that the graphs that show that Shannon's bound are best possible have multiple edges. Let $\mu(G)$ denote the maximum edge multiplicity of a graph G. Hence G is simple if and only if $\mu(G) \leq 1$.

Theorem 13 (Vizing, 1964)**.** *For a graph G we have $\chi'(G) \leq \Delta(G) + \mu(G)$.*

Proof. We prove this by induction on $e(G)$. The statement is trivial for graphs with no or one edges. Suppose G is a minimal counterexample, and fix an edge e between x and y of G. By assumption, there is an edge-colouring c of $G - e$ with $\Delta(G) + \mu(G)$ colours, which we fix.

Let $M(v)$ denote the set of colours in $[\Delta(G) + \mu(G)]$ which are missing at v, that is, which are not used by c on any edges containing v. Note that for each vertex v we have $|M(v)| \geq \mu(G)$. If $M(x) \cap M(y) \neq \emptyset$, we can extend c to a colouring of G and are done. If not, we try to modify c in order to get to this situation.

We start by looking at the edges containing y. If we can somehow recolour just these edges, maintaining a proper edge-colouring of $G - e$, in order to free up a colour in $M(x)$, then we are done. We'll consider only a special type of recolouring encoded by a directed multigraph. Let H be a directed multigraph on $N_G(y)$, where we put an arc from u to v for each edge f between v and y in G whose colour is in $M(u)$. The out-neighbours of x are the vertices v such that some edge f between y and v has a colour in $M(x)$; in other words, if we can recolour f we will free up $c(f)$ for the edge e we are eventually trying to colour. Two-arc paths from x correspond to doing two-step recolourings, and so on. Let X be the set of vertices in H we can reach from x. We first consider recolouring only edges adjacent to y.

Claim 1. *For each colour $\alpha \in M(y)$ and $v \in X$ we have $\alpha \notin M(v)$.*

Proof. Suppose not; let α and v be a counterexample. Let $(x, u_1, \ldots, u_t, v)$ be a shortest path in H from x to v. We first give e a colour used on some edge e_1 between y and u_1 which is in $M(x)$. Now we recolour e_1 with a colour used on some edge e_2 between y and u_2 which is in $M(u_1)$, then e_2 with a colour used on some e_3 between y and u_3 which is in $M(u_2)$, and so on until we recolour e_t with a colour used on edge e_{t+1} between y and v which is in $M(u_t)$. These colours and edges exist by the definition of H and because we chose a shortest path (and therefore each new edge we encounter still has the colour given by c). Finally, we recolour e_{t+1} with colour α to obtain a proper edge-colouring of G. $\square$

Another way we can try recolouring is to make the following observation: The subgraph G_{ij} of G consisting of all edges with colour i or j has maximum degree 2, and therefore all its components are either paths or (even) cycles. Suppose that α is some colour in $M(x)$, and β is in $M(y)$. There has to be an edge of colour α at y (otherwise α is in $M(x) \cap M(y)$) and since by assumption there is no edge of colour β, it follows that y is the start of a path component of $G_{\alpha\beta}$. We can swap the colours α and β for all edges on this path, which maintains a proper edge-colouring of $G - e$. Now we can colour e with colour α, unless this swapping destroyed the property $\alpha \in M(x)$; in other words, unless the path in $G_{\alpha\beta}$ has ends y and x. Putting this together with the idea of Claim 1 we get the following.

Claim 2. *For each colour $\alpha \in M(y)$, each $v \in X$ and each $\beta \in M(v)$, the graph $G_{\alpha\beta}$ has a component which is a path from v to y.*

Proof. Suppose not, and let α, β be the colours in a counterexample. Let v be a vertex of X which is at minimum distance from x forming a counterexample, and let P be a shortest path from x to v in H.

We first recolour edges 'following P' to v as in the proof of Claim 1. Let f denote the final edge from y to v whose colour we use in the recolouring; then the only reason we do not obtain a proper edge-colouring is that y is incident to two edges of colour $c(f)$. Since P is a shortest path, we still have $\alpha \in M(y)$ and $\beta \in M(v)$, so by Claim 1 we do not have $\alpha \in M(v)$. Thus there is a component of $G_{\alpha\beta}$ which is a path Q starting at v. If Q does not either end at y or a vertex of P, then we can swap colours α and β on Q, and recolour f with colour α to obtain a proper colouring.

It remains only to argue that Q cannot end at some vertex of P. Suppose to the contrary that its end-vertex z is on P. Because $M(z)$ does not contain α by Claim 1, it follows that $M(z)$ does contain β, and therefore α, β, z is a counterexample to the claim with smaller distance from x to z than from x to v, a contradiction. $\square$

Finally, we can conclude that for any distinct $v, w \in X$ we have $M(v) \cap M(w) = \emptyset$, for if not let β be in both sets and by Claim 2 for any $\alpha \in M(y)$ we see that in $G_{\alpha\beta}$ there are components which are paths starting at v and at w whose other ends are both y; this contradicts their being components.

Now we can complete the proof. On the one hand, from Claim 1, for each $v \in X$ and $\beta \in M(v)$ we have $\beta \notin M(y)$ and therefore there is an edge from y to some w (not v by assumption) of colour β. This gives an arc in H from v to w. It follows that there are at least $|M(v)| \geq \mu(G)$ arcs leaving each $v \in X$ in H, and at least $|M(x)| \geq \mu(G)+1$ arcs leaving x in H, so that in total at least $|X|\mu(G)+1$ arcs of H lie in X.

On the other hand, for any $u \in X$ and colour α used on some edge from u to y, if α is responsible for an arc $\overrightarrow{vu}$ in H then by definition of H we have $\alpha \in M(v)$. But if v and w are two distinct vertices of X, we just proved $M(v) \cap M(w) = \emptyset$, so α is responsible from at most one arc in H going from X to u. In other words, the in-degree of u is at most equal to the number of colours used on edges from y to u; this is at most $\mu(G)$ for any $u \in X$, and at most $\mu(G)-1$ for $u = x$. Summing, the number of arcs in H is at most $|X|\mu(G)-1$. This contradiction completes the proof. $\square$

It is fairly easy to deduce Shannon's Theorem 12 from Vizing's Theorem.

Corollary 14 (Vizing, 1964)**.** *For a simple graph G we have $\Delta(G) \leq \chi'(G) \leq \Delta(G)+1$.*

Determining which of the two numbers $\Delta(G), \Delta(G)+1$ is the right value for the chromatic index of a given simple graph G is not an easy task. This problem is known to be NP-complete. Even if all vertices have degree three, deciding if a the simple graph has chromatic index three or four is NP-complete (both results are due to Holyer, 1981)

Edge List Colouring

Also for this section we allow graphs to have multiple edges.

List colourings of edges are defined analogously to vertex list colourings. The essential comparable concepts are *k-edge-list-colourable* or *k-edge-choosable*, and *edge list chromatic number*, *list chromatic index* or *edge choice number*. This latter parameter is denoted $\mathrm{ch}'(G)$. Again we trivially have $\mathrm{ch}'(G) \geq \chi'(G)$. For vertex colouring, we have seen that the chromatic number and the choice number can be arbitrarily far apart. For edge colouring, it is conjectured that there actually is no difference! The following List Colouring Conjecture is attributed to Vizing (1975). (But see the Jensen & Toft book for a discussion about its history.)

Conjecture 15. *For every graph G, $\mathrm{ch}'(G) = \chi'(G)$.*

Using a greedy algorithm, it is again easy to obtain that $\mathrm{ch}'(G) \leq 2\,\Delta(G)-1$. For a long time, not much progress was made beyond that easy observation.

Theorem 16 (Borodin, Kostochka & Woodall, 1997))**.** *For a graph G we have $\mathrm{ch}'(G) \leq \frac{3}{2}\,\Delta(G)$ (so $\mathrm{ch}'(G) \leq \frac{3}{2}\,\chi'(G)$).*

Theorem 16 should be compared with Theorem 12. The proof of Theorem 16 uses the techniques developed by Galvin (see below). The next result says more, but only if the maximum degree is very large.

Theorem 17 (Kahn, 2000)**.** *For every $\epsilon > 0$ there exists a constant D_ϵ, so that if G is a graph with $\chi'(G) \geq \Delta(G) \geq D_\epsilon$, then $\mathrm{ch}'(G) \leq (1+\epsilon)\,\chi'(G)$.*

Kahn's proof is a masterpiece of probabilistic techniques (for which see Lecture 4).

One of the major breakthroughs in the research on edge list colouring was the following result.

Theorem 18 (Galvin, 1995)**.** *For every bipartite graph G,* $\mathrm{ch}'(G) = \chi'(G)$*.*

Galvin's proof relies on a concept called "kernels in directed graphs". An *orientation G^* of an undirected graph G* is an assignment of exactly one of the two possible directions to each edge of G. By the *outdegree* $d^+_{G^*}(v)$ of a vertex v we denote the number of arcs that have v as a tail.

A *kernel* of a directed graph G^* is an independent set $K \subseteq V$, so that for every vertex $v \in V \setminus K$, there is an arc in G^* from v to a vertex in K. (An *independent set* is a set of vertices in which no pair is joined by an edge.)

The first step to proving Theorem 18 is to show that orientations with kernels help in *vertex*-colouring (a result of Bondy, Boppana and Siegel).

Lemma 19. *Let G be a graph and L a vertex list assignment of G. Suppose there exists an orientation G^* of G, such that $|L(v)| \geq d^+_{G^*}(v) + 1$ for each vertex v and such that every induced subgraph of G^* has a kernel. Then G is L-colourable.*

Proof. We use induction on $v(G)$. Given G and an assignment L of lists with $|L(v)| \geq d^+_{G^*}(v) + 1$ for each v, choose a colour c which appears in at least one list, and let S be the set of vertices v such that $c \in L(v)$. Let K be a kernel of $G^*[S]$. By induction, we can properly colour $V(G) \setminus K$ such that for each $v \in V(G) \setminus K$ we assign to v a colour in $L(v) \setminus \{c\}$. The reason is that for $v \in S$ we have $|L(v) \setminus \{c\}| = |L(v)| - 1$ but v has an outneighbour in K (since K is a kernel) while for $v \notin S$ we have $|L(v) \setminus \{c\}| = |L(v)|$ by definition of S. Now since K is independent and all vertices of K have lists containing c, we can colour all vertices of K with colour c, obtaining the desired colouring. □

Recall that the line graph $L(G) = (V_L, E_L)$ is the graph that has the edges of G as vertices: $V_L = E$; and two edges are adjacent in the line graph if they have a common endvertex in G.

For a graph G, for each vertex v choose a linear ordering $\leq_v$ of the edges incident with v. Then we can translate this to an orientation of the line graph $L(G)$ as follows: If two edges e, f have a common endvertex v, and $e \leq_v f$ in the chosen linear ordering around v, then orient the edge ef in $L(G)$ from e to f. If e and f are parallel edges, then it is possible that we have both an arc from e to f and an arc from f to e (if e and f have different ordering around each of their two common endvertices). This causes no problems in what follows.

We call any orientation of $L(G)$ obtained from a system $\leq_v$ of linear orderings as above a *normal orientation*. Notice that an induced subgraph of a line graph with a normal orientation is again a line graph with a normal orientation.

So what does a kernel in a line graph $L(G)^*$ with normal orientation look like? First we observe that an independent set in a line graph $L(G)$ corresponds to a matching in G. (A *matching* is a set of edges so that no two have a common endvertex.) Next assume the normal orientation of the line graph originated from linear orderings $\leq_v$ of the edges incident with each vertex v. So a kernel in $L(G)^*$ is a matching M in G so that for each edge $e \in E \setminus M$ there is an arc from e to some $f \in M$ in $L(G)^*$. In other words, for each edge $e = uv \in E$ with $e \notin M$, we have that there is an edge $uw = f_1 \in M$ with $e \leq_u f_1$, or an edge $wv = f_2 \in M$ with $e \leq_v f_2$.

The following lemma, together with Lemma 19, is most of the proof of Galvin's Theorem 18.

Lemma 20. *Let G be a bipartite graph and let $L(G)^*$ be a normal orientation of the line graph of G. Then $L(G)^*$ has a kernel.*

Proof. Assume the normal orientation of the line graph originated from linear orderings $\leq_v$ of the edges incident with each vertex v. And denote the two parts of G by X and Y.

We prove the lemma by induction on the number of edges of G. If there is only one edge, then we can just use that edge as the kernel. So assume there is more than one edge. For each $x \in X$, let e_x be the edge incident with x that is maximal for the linear ordering $\leq_x$. Take $N = \{e_x \mid x \in X\}$.

If N is a matching, then it is a kernel in $L(G)^*$, since for each other edge $e = xy$ with $x \in X$, we have that $e \leq_x e_x$, hence there is an arc in $L(G)^*$ from e to $e_x \in N$.

So suppose N is not a matching, hence there exists $x, x' \in X$, $x \neq x'$, and $y \in Y$ so that $e_x = xy$ and $e_{x'} = x'y$. Without loss of generality we can assume $e_x \leq_y e_{x'}$. Now remove e_x from G to form G^-, and leave the orderings of the edges around each vertex the same. By induction, $L(G^-)^*$ has a kernel. This kernel corresponds to a matching M in G^-.

If $e_{x'} \in M$, then, since $e_x \leq_y e_{x'}$, there is an arc in $L(G)^*$ from e_x to $e_{x'}$, so M is also a kernel in $L(G)^*$.

If $e_{x'} \notin M$, then there is an edge $f \in M$ so that there is an arc $L(G)^*$ from $e_{x'}$ to f. But since $e_{x'}$ was the maximal element around x', this arc must arise since $e_{x'}$ and f both have y as a common endvertex, and $e_{x'} \leq_y f$. As we also have $e_x \leq_y e_{x'}$, this means $e_x \leq_y f$, and hence also this time we can conclude that M is a kernel in $L(G)^*$. $\square$

Proof of Theorem 18. Take $k = \chi'(G)$, and let $\varphi : E \to \{1, \ldots, k\}$ be an edge colouring of G. Denote the two parts of G by X and Y.

We form the following linear orderings of the edges around a vertex v. If $x \in X$, and e_1, e_2 have x as an endvertex, then we set $e_1 \leq_x e_2$ if $\varphi(e_1) \leq \varphi(e_2)$. While if $y \in Y$, and e_1, e_2 have y as an endvertex, then we set $e_1 \leq_y e_2$ if $\varphi(e_1) \geq \varphi(e_2)$.

Form the orientation $L(G)^*$ of $L(G)$ using the linear orderings above. What can we say about $d^+_{L(G)^*}(e)$ of an edge $e = xy$ of G? This is the number of edges incident with x that have a colour larger than e plus the number of edges incident with y that have a colour smaller than e. Since all edges incident with the same vertex have different colours, and there are k colours in total, this means that for all edges e we have $d^+_{L(G)^*}(e) \leq k - 1$ (we subtract one for the colour e has itself).

So if we give each edge e a list $L(e)$ of $k = \chi'(G)$ colours, then by Lemma 19 we know that the edges are L-colourable, proving the theorem. $\square$

Colouring C_5-free graphs

All the proofs in the last sections are fairly 'hands-on'. Except for Galvin's Theorem, we construct a colouring step by step, perhaps having to play with a colouring which almost works until it does. Galvin's Theorem is a little different — we take bigger steps, using the kernel method. But still this is in some sense a step-by-step colouring.

There are lots of other methods of colouring, though. Here is one.

Theorem 21. *For each $\alpha > 0$, there exists $C = C(\alpha)$ such that every n-vertex graph G which has minimum degree at least αn and does not contain C_5 satisfies $\chi(G) \leq C$.*

This is originally a theorem of Thomassen (with a very nice proof). The proof here is due to Łuczak and Thomassé.

Proof. We can assume $\alpha \leq \frac{1}{2}$, since any graph with five or more vertices and minimum degree bigger than $n/2$ contains C_5 (this is an easy exercise). We choose $\varepsilon = \alpha/100$ (which is easily small enough for the proof; there is no good reason to try to optimise constants), and we set $K = \lfloor \log_{1+\varepsilon} 2/\alpha \rfloor + 1$. We choose

$$C(\alpha) = 100^K \alpha^{-K-10}\,.$$

If $n \le C(\alpha)$, we simply colour each vertex of G with a different colour. So suppose $n > C(\alpha)$. We begin by taking a maximum cut (X, Y) of G; that is, $Y = V(G) \setminus X$ where X is chosen to maximise the number of edges between X and Y.

Without loss of generality, we can assume $\chi(G[X]) \ge \chi(G[Y])$. So it is enough to show $\chi(G[X]) \le C/2$. Observe that for each $x \in X$ we have $d(x;Y) \ge \alpha n/2$ (otherwise we could move x to Y and increase the number of edges crossing). Here $d(x;Y)$ means the number of edges from x to Y.

Given $X' \subset X$ and $Y' \subset Y$, and a partition $Y' = Y_1 \cup Y_2 \cup Y_3$, for $i = 1, 2, 3$ let

$$X_i := \left\{ x \in X' : x \notin X_j \text{ for } j < i \text{ and } \frac{d(x;Y_i)}{|Y_i|} \ge (1+\varepsilon)\frac{d(x;Y')}{|Y'|} \right\}.$$

Let $X_4 := X' \setminus (X_1 \cup X_2 \cup X_3)$. If $|Y_i| \ge \frac{\alpha}{10}|Y'|$ for each $i = 1, 2, 3$, and X_4 is independent then we have an *ε-booster for* (X', Y').

We generate partitions $\mathcal{X}$ and $\mathcal{Y}$ of X and Y respectively, together with a relation 'in correspondence', as follows. We begin with $\{X\}, \{Y\}$ and we say that X and Y are in correspondence. We now iteratively do the following. Pick $X' \in \mathcal{X}$ and $Y' \in \mathcal{Y}$ which are in correspondence. If there is a ε-booster Y_1, Y_2, Y_3 for (X', Y'), we replace Y' with Y_1, Y_2, Y_3 and we replace X' with X_1, X_2, X_3, X_4 (if some of these sets are empty we simply do not add them); we say X_i and Y_i are in correspondence for each $i = 1, 2, 3$. So X_4 is not in correspondence with any set of the new $\mathcal{Y}$. We repeat this procedure until any remaining (X', Y') which are in correspondence do not have ε-boosters.

There is a natural way to draw a tree representing this process: we start with the root labelled X, then add children labelled with the sets into which X is split by finding an ε-booster for (X, Y), and then to each of those, children for their ε-boosters, and so on.

We claim that when this process terminates we have $\mathcal{X}$ with at most 4^K parts. Indeed, suppose that we have in the above tree a path from the root with $K+1$ vertices. Then we have $Y = Y'_0, Y'_1, \ldots, Y'_K$, where each Y'_i is obtained by finding an ε-booster in Y'_{i-1}. Let x be a vertex in the corresponding set of $\mathcal{X}$ (the end of the path). Then we have

$$\frac{d(x;Y'_K)}{|Y'_K|} \ge (1+\varepsilon)\frac{d(x;Y'_{K-1})}{|Y'_{K-1}|} \ge \cdots \ge (1+\varepsilon)^K \frac{d(x;Y'_0)}{|Y'_0|} \ge (1+\varepsilon)^K \frac{2}{\alpha} > 1\,.$$

But x cannot have more than $|Y'_K|$ neighbours in Y'_K; this is a contradiction. So our tree, in which each node has at most four children, also has depth at most K; it has at most 4^K leaves, and the leaves are precisely the elements of $\mathcal{X}$.

This also justifies that any set $Y' \in \mathcal{Y}$ has size at least $\left(\frac{\alpha}{10}\right)^{K+1} n$; we start with $|Y| \ge \frac{\alpha}{2} n$, and by definition each time we find a booster our sets decrease in size by at most a $\frac{\alpha}{10}$ factor.

Now consider a set X' in the final $\mathcal{X}$. If it does not have a corresponding Y', it is independent. If it does have a corresponding Y', then (X', Y') does not have any ε-booster. Suppose $a, b, c \in X'$ and $ab \in E(G)$. Each of a, b, c have at least $\alpha|Y'|/2$ neighbours in Y' by construction. Let Z_1 be a set of $\alpha|Y'|/4$ vertices in Y' which are neighbours of a; and let Z_2 be a disjoint set of $\alpha|Y'|/4$ neighbours of b in Y'. Let $Z_3 := Y' \setminus (Z_1 \cup Z_2)$. Note that each of Z_1, Z_2, Z_3 contains more than $\frac{\alpha}{10}|Y'|$ vertices. Since there is no booster for (X', Y'), in particular this partition of Y' does not give an ε-booster. So in the corresponding partition X'_1, X'_2, X'_3, X'_4 of X' the set X'_4 is not independent.

We need to separate two cases. First, it is possible that we have $a, b \in X'_4$. But this can only happen if both a and b have more than $(1-\varepsilon)|Y'|$ neighbours in Y' (because we need $(1+\varepsilon)\frac{d(a;Y')}{|Y'|} > 1$ to avoid $a \in X'_1$, and similarly for b). But in this case, since c has at least $\alpha|Y'|/2$ neighbours in

Y, and since $n > C(\alpha)$, we can find distinct vertices $d, e \in Y'$ such that d is adjacent to a and c, and e to b and c. Now a, d, c, e, b gives a copy of C_5 in G, a contradiction.

The 'interesting' case is that at most one of a, b is in X_4'. Since X_4' is not independent, it contains at least two vertices; we can take c to be some vertex of X_4' other than a, b. Now, as above, if c has neighbours d in Z_1 and $e \in Z_2$, we get a copy of C_5 on vertices a, d, c, e, b. But suppose c has no neighbours in Z_1. Then we have $d(c; Z_2 \cup Z_3) = d(c; Y')$, and so

$$\frac{d(c; Z_2 \cup Z_3)|Y'|}{|Z_2 \cup Z_3| d(c; Y')} = \frac{|Z_2 \cup Z_3|}{|Y'|} \geq \tfrac{1}{1-\alpha/10} > 1 + \varepsilon \,,$$

and by averaging for at least one of $i = 2, 3$ we have

$$\frac{d(c; Z_i)|Y'|}{|Z_i| d(c; Y')} > 1 + \varepsilon \,,$$

which is a contradiction to the assumption $c \in X_4'$. The same calculation gives a contradiction if c has no neighbours in Z_2; again we reached a contradiction.

In conclusion, if in the final $\mathcal{X}$ we have a set X' in correspondence with some Y' then either X' is independent, or it has at most two vertices; in either case we can colour it with at most two colours. So we can colour $G[X]$ using at most $2 \cdot 4^K$ colours. By the same argument we can colour Y with at most $2 \cdot 4^K$ colours. Putting these together, we can colour G with at most $4^{K+1} < C(\alpha)$ colours, as desired. □

This 'booster-tree method' is a personal favourite of the current lecturer—but it's one we don't really know how to use. It's been used in proofs broadly similar to the above, but it probably should be much more widely useful.

Exercises

Exercise 1. *Complete the proof of Brooks' Theorem.*

Exercise 2. *For $n \geq 3$, a cycle C_n is a graph with vertex set $\{v_1, \ldots, v_n\}$ and edge set $\{v_1v_2, v_2v_3, \ldots, v_{n-1}v_n, v_nv_1\}$.*

Let C_n, $n \geq 3$, be a cycle. Let us assign lists $L(v)$ of two colours to each vertex v of C_n. Show that there is an L-colouring, except if n is odd and all lists are identical.

Let C_{2k}, $k \geq 2$, be an even cycle. Prove that $\chi(C_{2k}) = \mathrm{ch}(C_{2k}) = \chi'(C_{2k}) = \mathrm{ch}'(C_{2k}) = 2$, using only the definitions.

Let C_{2k-1}, $k \geq 2$, be an odd cycle. Prove that $\chi(C_{2k-1}) = \mathrm{ch}(C_{2k-1}) = \chi'(C_{2k-1}) = \mathrm{ch}'(C_{2k-1}) = 3$.

Exercise 3. *Determine the chromatic index $\chi'(K_n)$ of the complete graphs K_n.*

Exercise 4. *Let G be a planar graph and suppose it is drawn in the plane so that at least one face has more than three edges on its boundary. Show that you can add an edge to G so that the larger graph is again a planar, simple, graph. (The issue here is that you need to prevent adding an edge between two vertices that are already connected by an edge. In other words: make sure that after adding an edge the larger graph is still simple.)*

Exercise 5. *Explain how Shannon's Theorem 11 follows easily from Vizing's Theorem 12.*

Exercise 6. *Give an infinite family of directed graphs without a kernel.*

Exercise 7. *A* total colouring *of a graph $G = (V, E)$ is an assignment of colours to vertices* and *edges $V \cup E$ so that adjacent or incident elements get different colours. The smallest k such that G has a total colouring with k colours is called the* total chromatic number, *denoted $\chi''(G)$.*

Find lower and upper bounds of $\chi''(G)$ in terms of $\Delta(G)$ and/or $\deg(G)$.

Determine the total chromatic number of cycles C_n and complete graphs K_n.

Show that if the List Colouring Conjecture 14 is true, then for simple graph G we have $\chi''(G) \leq \Delta(G) + 3$.