1 Other coloring problems

1.1 Edge coloring

Many properties which are traditionally based on vertices (e.g. connectivity) also exist in an edge version, so it is probably not surprising that one can investigate coloring as a property of edges as well as vertices.

Definition 1. A *r-edge-coloring* of a graph G is a function $c : E(G) \to \{1, 2, 3, \ldots, r\}$ such that, if e and f are distinct edges sharing an endpoint, $c(e) \neq c(f)$. If G has an r-edge-coloring, it is called *r-edge-colorable*. The *edge chromatic number* of G, denoted $\chi'(G)$, is the least r such that G is r-edge-colorable.

It is easy to find a lower bound for $\chi'(G)$, and a quite sloppy upper bound:

Proposition 1. *For a nontrivial graph G, $\Delta(G) \leq \chi'(G) \leq 2\Delta(G) - 1$.*

Proof. Let v be a vertex of maximum degree in G, i.e. $d(v) = \Delta(G)$, so v has incident edges $e_1, e_2, \ldots, e_{\Delta(G)}$. All of these edges must be different colors in a valid coloring, so at least $\Delta(G)$ colors are required.

On the other hand, if we enact a greedy coloring of the edges (i.e. coloring the edges, in some order, using the least color possible at each step), then each edge $\{u, v\}$ which we have not yet colored is forbidden to use at most $(d(u) - 1) + (d(v) - 1) \leq 2\Delta(G) - 2$ colors. Thus, there is some color in $\{1, 2, \ldots, 2\Delta(G) - 1$ which is not forbidden, and can be used to color this edge in a greedy coloring. Since this is true of every edge we consider for greedy coloring, the entire greedy coloring can be achieved with $2\Delta(G) - 1$ or fewer colors. □

There is a whole family of graphs that achieves the lower bound seen above:

Theorem 1 (König '18)**.** *If G is bipartite, then $\chi'(G) = \Delta(G)$.*

Proof. We shall prove this via induction on $\|G\|$. When $\|G\| = 0$, both $\chi'(G)$ and $\Delta(G)$ are zero, so the base case is trivially true.

Now, considering an arbitrary G, let us select an edge $e = \{u, v\}$. By our inductive hypothesis, $G - e$ is $\Delta(G)$-colorable (in fact it is possible that one color is not even necessary, but we don't need that fact). Since $d_{G-e}(u) \leq \Delta(G) - 1$ and $d_{G-e}(v) \leq \Delta(G) - 1$, then there must be some color a not utilized on edges incident on u, and some color b not utilized on edges incident to v. If $a = b$, we may extend our coloring of $G - e$, coloring e in the color a to get a simple $\Delta(G)$-coloring of G. The interesting and troublesome case, then, is where the color not utilized at u is utilized at v and vice versa.

Let us denote $v = v_0$, and denote the other endpoint of the edge in color a from v_0 as v_1. Now note that if v_1 has no incident edge in color b, then we could simply change the edge $\{v_0, v_1\}$ to be of color b, and then we would have the situation described above, where the coloring in $G - e$ is easily extended to G. Now, consider the possibility that v_1 does have an incident edge in color b, and denote its other endpoint as v_2. Then, if v_2 had no incident edge in color b, we could toggle *both* the edges $\{v_0, v_1\}$ and $\{v_1, v_2\}$ between colors a and b to get an easily extended coloring.

The above logic can be extended arbitrarily far, so let us consider a maximal walk $v_0 \sim v_1 \sim v_2 \sim v_3 \sim \cdots$ in $G - e$ alternating between edges of color a and b. Note that this walk will in fact be a path: if we suppose v_j is the first revisitation of an earlier vertex v_i, then we see that v_i must either be incident to 3 distinct edges in colors a and b (the edges $\{v_{i-1}, v_i\}$, $\{v_i, v_{i+1}\}$, and $\{v_{j-1}, v_j\}$) or $i = 0$ and v_0 is incident on 2 distinct edges in the color a (the edges $\{v_0, v_1\}$ and $\{v_{j-1}, v_j\}$). In neither case could the coloring be a legal coloring of $G - e$, so this situation cannot occur and the above walk is actually a path (and thus finite).

Sicne this path is maximal, its endpoint v_n has only one incident edge in color a or b (which one depends on the parity of n), so we could swap the roles of colors a and b on this path and retain a valid coloring of $G - e$. The only possible impediment to thereby freeing up color a for use on edge e is the possibility that this path passes through u, so that toggling the colors on it would leave an edge of color a incident on e. Such a situation could only transpire if u were in fact the endpoint of the path, since it only has an incident edge of color n, and none of color a. Furthermore, in order for the edge $\{v_{n-1}, v_n\}$ to be color b, n would need to be even, since the colors on the path alternate. Thus, in order for u to be on the path, there would need to be a path in $G - e$ from v to u of even length, which together with e forms an odd cycle in G, violating bipartiteness, so the aforementioned path does not pass through v, and thus a toggle-and color approach serves to give G a valid $\Delta(G)$-coloring. □

Incidentally, the above proof is explicitly constructive: not only does it assert that a $\Delta(G)$-coloring exists for any bipartite graph, but that a slightly modified greedy algorithm is an effective way to find such a coloring.

We might be curious about the achievability of the upper bound as well; surprisingly, except in the trivial case of $\Delta(G) = 1$ (where it equals the lower bound) and $\Delta(G) = 2$ (where is is achieved by any graph containing an odd cycle). In fact there is a much tighter upper bound, which mostly puts to rest questions about the edge chromatic number.

Theorem 2 (Vizing '64)**.** $\chi'(G) \leq \Delta(G) + 1$.

Proof. We prove this by induction on $\|G\|$. In the $\|G\| = 0$ case it is trivial, as seen in the previous theorem. For the induction step, as above we note that given an edge e, $G - e$ must have a $(\Delta(G) + 1)$-coloring by the inductive hypothesis. It is important to note that this is true for *every* choice of e.

In the previous theorem, we showed how, given an edge e, we know some color is absent from each endpoint in the coloring of $G - e$, and that the only possible impediment to creating a slight modification of this coloring and extending it to G is a path in alternating colors in $G - e$ between the endpoints of e (note that in the above theorem, bipartiteness guaranteed such a path donesn't exist, but here we have no such guarantee).

Thus, the only way Vizing's theorem can be false is if the following remarkable circumstance occurs: there is a graph G such that for every edge $e = \{u, v\}$ and every coloring of $G - e$ which leaves a color a unused at u and a color b unused at v, there is an alternating path in colors a and b between u and v.

This is an extremely restrictive condition, and it should come as no surprise that no graph actually satisfies it. We shall show it by arbitrarily picking a vertex u and one of its neighbors v_0, denoting $e_0 = \{u, v_0\}$, and considering a $(\Delta(G) + 1)$-coloring c_0 on $G - e_0$. There are at least 2 colors unused incident on v_0 in this coloring; if either of them are unused incident to x we have a trivial coloring-extension, so let us consider instead the situation where $e_1 = \{u, v_1\}$ is an edge such that $c_0(e_1)$ is unused incident to v_0. There is at least one unused color incident to v_1, so if that color is incident to u we shall further construct $e_2 = \{u, v_2\}$ such that v_2 is distinct from v_1 and v_0 such that $c_0(e_2)$ is unused incident to v_0. We continue in such a fashion until we have a maximal sequence of distinct vertices $v_0, \ldots, v_k$ and associated edges $e_0, \ldots, e_k$ such that $c_0(e_k)$ is unused incident to v_{k-1}.

The advantage of such a sequence is that a very minor modification of the coloring c_0 of $G - e_0$ can serve to color $G - e_i$. We define the coloring

$$c_i(e) = \begin{cases} c_0(e_{j+1}) \text{ if } e = e_j \text{ for } j < i \\ \quad c_0(e) \text{ otherwise} \end{cases}$$

which shuffles every edge-color, starting at e_j downwards one place in the sequence. This will be a proper coloring of $G - e_i$. because it uses the same colors at u as c_0 itself, and at each v_i it introduces a color which was unused at v_i under the scheme in c_0. Note that except on the edges $e_0, \ldots, e_k$, all these colorings agree.

Now, let us consider an unused color in c_0 at v_k and denote this color as a. Since our sequence of distinct vertices adjacent to u is maximal, there is no possible v_{k+1}, so either the color a is also unused at u, or some edge e_i for $i < k$ is of color b. In the first case, where a is unused at both v_k and u in c_0, that fact will remain true for c_k, so c_k is trivially extendable to a coloring of G by coloring e_k in color a.

So now we must only consider the case where a is unused at y_k, and there is some i such that $c_0(e_i) = a$. Let b be a color unrepresented at x in c_k; our extremely rigorous condition for non-color-extendability above guarantees that there must be a path alternating between the colors a and b in c_k from y_k to x. Since $c_0(e_i) = a$, it follows that $c_k(e_{i-1}) = a$, and the edge e_{i-1} must be used in the aforementioned path, since it is the only edge in color a or b incident on u. Note that every edge in this path except the last one is not an e_j edge, so this alternating path is, except for its last section, present in every coloring.

Now, we consider the coloring c_{i-1} of the graph $G - e_{i-1}$. By the construction criterion for our sequence of vertices, we know that $c_0(e_i) = a$ is unrepresented at v_{i-1} in c_0, and thus also in c_{i-1}. Since a is unrepresented at v_{i-1} and b is unrepresented at x, the alternating path in colors a and b in c_{i-1} from v_{i-1} must terminate at u. But most of this path we already know, from inspecting it in c_k: we know an alternating path in colors a and b goes from v_{i-1} to v_k, ending in color b. We know the color a is unused at v_k in c_0 (and thus also in c_{i-1}, so our path does not actually reach all the way to u, violating our condition of unextensibility of colorings. □

So it turns out that edge chromatic numbers are very closely bounded by degree: every graph G has an edge chromatic number of either $\Delta(G)$ or $\Delta(G) + 1$. While specific graphs have been classified (e.g. bipartite graphs, planar graphs of large degree), it is in general provably difficult (in the algorithmic complexity sense) to determine which of these two chromatic numbers an arbitrary graph has.

1.2 List coloring

There are a number of other variations on proper colorings, which we won't dwell much on here, but introduce their descriptions and basic facts about them:

Definition 2. For a graph G and a set S_v for each $v \in V(G)$, a *list coloring* of G is a selection of colors $c(v) \in S_v$ such that adjacent vertices have distinct colors. The *list-chromatic number* is the least k such that if each $|S_v| = k$, then regardless of the choice of S_v, G has a list coloring.

The list-chromatic number is obviously at least equal to the regular chromatic nubmer, since a list coloring with the specific set $\{1, 2, 3, \ldots, k\}$ at each vertex is identical to a traditional k-coloring. However, the list chromatic number can, surprisingly, exceed the chromatic number: consider, for example, $K_{3,3}$, which is 2-colorable. If each part contains a vertex with list $\{1, 2\}$, $\{1, 3\}$, and $\{2, 3\}$, then two colors must appear in each part, guaranteeing two adjacent vertices of the same color, so $K_{3,3}$ is not 2-list-colorable. This construction can actually be extended to produce bipartite graphs of arbitrarily high list-chromatic number: if $m = \binom{2k-1}{k}$, there is a choice of lists of size k for the vertices of $K_{m,m}$ such that the resulting structure is not list-colorable: assign to each vertex a distinct k-element subset of $\{1, 2, 3, \ldots, 2k-1\}$. Then every selection from the top is guaranteed to include k distinct colors: if only $k-1$ disticnt colors are selected on top, which color was selected from the node whose list is the set of k unselected colors? Likewise on the bottom part, so since $k + k > 2k - 1$, some color is used on both the top and bottom, so this coloring is not valid.

There is also a list-coloring parameter for edges, which remarkably seems to have the same value as the traditional edge-coloring parameter, but little progress has been made towards proving this equality.

2 Line graphs

We have thus far drawn a number of correspondences between vertex parameters and associated edge parameters. We could generalize this construction and save ourself some symbols by defining a graph in which the vertices stand in for edges of our original graph.

Definition 3. The *line graph* of G, denoted $L(G)$, is a graph such that $V(L(G)) = E(G)$ and distinct e and f are adjacent in $L(G)$ if and only if they share an endpoint in G.

Some results translate naturally: a proper edge-coloring of G is a vertex-coloring of $L(G)$ and vice versa, so $\chi'(G) = \chi(L(G))$. Likewise, any paths in G can be converted to paths in $L(G)$ (although not necessarily the reverse), so $L(G)$ is connected if G is either connected or consists of a connected component and isolated vertices, and similarly, $L(G)$ is Hamiltonian if G is Eulerian. Some correspondences are a bit surprising: a matching in G consists of edges not sharing endpoints, which correspond to non-mutually-adjacent vertices in $L(G)$, which is to say, an independent set. Similarly, a clique in $L(G)$ corresponds to a collection of edges all of which share endpoints; if the clique has size greater than 3 they must all share the same endpoint, so if $\omega(L(G)) > 3$, then $\omega(L(G)) = \Delta(G)$.

On that note, line graphs tend to be rather clique-heavy, since any vertex G with several neighbors forms a clique in $L(G)$. We can quantify this by noting that some sucstructures cannot occur in a line graph:

Proposition 2. *For any graph G, $L(G)$ contains no induced $K_{1,3}$.*

Proof. Suppose $L(G)$ contained an induced $K_{1,3}$; then it contains vertices e and f_1, f_2, f_3 such that e is adjacent to each f_i but the f_i are mutually nonadjacent. Thus, in G there are edges e, f_1, f_2, f_3 such that e shares an endpoint with each of f_1, f_2, f_3, but that the f_i have mutually exclusive sets of endpoints. Let $e = \{u, v\}$. The above assertion guarantees that each f_i is incident to either u or v. Since there are 3 of them, the Pigeonhole Principle guarantees two of them are adjacent to the same one, and thus two of the f_i must share an endpoint. □

Graphs without an induced $K_{1,3}$ are usually called *claw-free*, and this is considered to be in some ways a generalization of the concept of a line graph.

Since we have shown that certain structures cannot occur in a line graph $L(G)$, two questions arise: given a graph H, can we verify whether or not it is a line graph? And if it is a line graph, can we reconstruct the graph G whose line graph it is? The answer to the second question is clearly "no", if we demand that G be unique; isolated vertices in G have no effect on the line graph, but more problematically, the graphs K_3 and $K_{1,3}$ both have line graph equal to K_3.

However, aside from this peculiar special case, line graphs are in fact uniquely determined and easily identified.

Theorem 3 (Whitney, '32)**.** *A graph H is the line graph of some G if and only if there is a collection $H_1, H_2, \ldots, H_n$ of cliques of size 1 or larger in H such that each edge of H lies in exactly one H_i and each vertex of H lies in exactly two of the H_i. The graph G is determined by $V(G) = \{v_1, v_2, \ldots, v_n\}$ such that $v_i \sim v_j$ if the cliques share a vertex.*

Furthermore, the collection $H_1, \ldots, H_n$ is uniquely determined by H unless some component of H is a K_3, so that graph G without isolated vertices is uniquely determined.

Proof. Without loss of generality, we may consider connected H (and connected G); if H has distinct components, this procedure can be perforemd on each of them individually, and likewise if G has distinct nontrivial components, they each have their own line graph.

We shall consider several possible substructures of graphs, assuming that at least one clique-decomposition as above is possible, and then showing that *only* one such decomposition is possible. Specifically, we will look at an individual edge e in H, and demonstrate that that edge must belong to a highly specified partition in order to yield a clique-decomposition.

Case I: e lies in no clique of size ≥ 2. Then the only clique which contains e is the K_2 consisting of e itself.

Case II: e lies in a disconnected K_3. This is the one special case where the clique-decomposition is not unique: One may choose 3 K_2s consisting of the edges of the K_3, or choose the entire K_3 and 3 K_1s consisting of the vertices.

Case III: e lies in exactly two distinct K_3s whose identifying vertices are non-adjacent. Let $e = \{u, v\}$, and let x, y be the non-adjacent vertices in the two K_3 subgraphs incident on e. By construction, u is incident on the edges $\{u, v\}$, $\{u, x\}$, and $\{u, y\}$. These cannot each be a K_2 clique in the decomposition, as that would cover u three times. Thus, these edges must be covered by a K_3 (which covers $\{u, v\}$ and one of the other two) and a K_2 (which covers the remaining edge). Likewise, the remaining edge incident on v must be covered by a K_2, and at this point we have covered all vertices except either x or y twice. Now, if the entire graph is this 4-vertex system, then although we have two choices of decomposition, they are isomorphic. If x has other neighbors, we are forced to use $\{u, v, y\}$ as a K_3, and vice versa if y has other neighbors. The case where both x and y have other neighbors cannot, we see by the above argument, be a line graph.

Case IV: e lies in three or more distinct K_3s whose identifying evertices are non-adjacent. If $e = \{u, v\}$, and $\{u, v, a\}$, $\{u, v, b\}$, and $\{u, v, c\}$ are cliques, and a, b, c are mutually nonadjacent, so $\{u, a, b, c\}$ is an induced $K_{1,3}$, which cannot lie in a line graph.

Case V: e lies in exactly one clique of size K_r for $r > 3$ but no K_{r+1}, and every clique containing e is a subgraph of the aforementioned K_r. We shall see that this clique of size K_r must be the partition containing e. Suppose e lies in a smaller clique than K_r. Now either $e = \{u, v\}$ is a K_2

partition in its own right, or partitioned into a clique K of size 3 or larger. In the latter case, there is a vertex v in the K_r outside of K which has 3 edges to K which still need covering; these three edges could only be covered with K_2s, which would over-cover v. On the other hand, if $e = \{u, v\}$ is covered with a K_2, there are vertices x and y in $K_r - e$ such that $\{u, x\}, \{u, y\}, \{v, x\}, \{v, y\}$ must all be covered. Since u and v have already been covered once, and they have 2 incident edges still requiring covering, they must be covered with K_3s or larger, but this will double-cover $\{x, y\}$. $\square$