The Theory of Computation

Bernard M.E. Moret

Addison-Wesley-Longman, 1998

Solution Manual
Chapter 8, Problem 8.14

We want to prove Vizing's theorem: the chromatic index of a graph is either the maximum degree of the graph or one more than that. That the chromatic index has to be at least equal to the degree of the graph is obvious, so we need only prove that it need not exceed that degree by more than one.

Let d be the maximum degree of the graph. We prove the theorem by induction on d. The basis is clearly true for $d = 0$: such a graph has no edge at all, so that its chromatic index is 0. The proof does not involve any deep reasoning, but it has a lot of case analysis.

Assume then that the theorem holds for all graphs of degree up to some d and let us consider an arbitrary graph G of degree $d + 1$. Remove from this graph enough edges to reduce its degree to d, for instance in the following manner: start at any vertex of degree $d + 1$; if one of its neighbors is a vertex of degree $d + 1$, remove the edge between the two, otherwise remove any incident edge; then move on to any remaining vertex of degree $d + 1$. By inductive hypothesis, the resulting subgraph G' can be edge-colored with d or $d + 1$ colors. Now consider adding back (or, equivalently, coloring) the removed edges. In particular, consider coloring edge $\{u, v\}$; we proceed to show that we can color this edge from our set of $d + 2$ colors. We can then move on to any remaining uncolored edge and eventually color the whole graph, thereby completing our proof.

Denote by $d(w)$ the set of colors (between 1 and $d + 2$) *not* used among the colored edges incident upon w. Obviously, if we have $d(u) \cap d(v) \neq \emptyset$, we can assign to the new edge any color from the intersection. Consider then the cases where $d(u)$ and $d(v)$ do not share any color. By hypothesis, we have $d(u) \geq 2$ and $d(v) \geq 2$, because at least on edge incident upon each is as yet uncolored and each can have at most $d + 1$ incident edges. For any neighbor w of u, we must also have $d(w) \geq 1$, again because no vertex has degree exceeding $d + 1$; in case the edge $\{u, w\}$ is not yet colored, then we must in fact have $d(w) \geq 2$. Thus we can think of assigning to each colored edge $\{u, w_i\}$ incident upon u an alternate color picked from $d(w_i)$, a color that differs from the color used for the edge.

Let c_1 be a color in $d(v)$; by our assumption, it cannot be a color in $d(u)$, so that there must be an edge incident upon u that is colored c_1; denote this edge $\{u, w_1\}$ and denote the alternate color assigned to this edge by c_2—thus we have both $c_2 \neq c_1$ and $c_2 \notin d(w_1)$. Now either c_2 is in $d(u)$ or it is not. If it is, we are done: we recolor edge $\{u, w_1\}$ with color c_2, leaving us free to color edge $\{u, v\}$ with color c_1.

If it is not, we start an inductive process. If c_2 is not in $d(u)$, then there is some vertex, call it w_2, such that the edge $\{u, w_2\}$ is colored with c_2. Let c_3 be the alternate color that was assigned to this edge. Again, either c_3 is in $d(u)$ or it is not. Again, if it is, we are done, this time with a slightly longer chain: we recolor edge $\{u, w_2\}$ with color c_3, then recolor edge $\{u, w_1\}$ with color c_2, leaving us free to color edge $\{u, v\}$ with color c_1. If it is not, then there must be some vertex w_3 such that edge $\{u, w_3\}$ is colored with c_3. The process must either terminate at some stage i with c_i in $d(u)$, in which case we successively recolor all edges involved, freeing us to color $\{u, v\}$ with color c_1, or it terminates at some stage i because we have a "loop"—that is, because c_{i+1}, while not in $d(u)$, is equal to an earlier c_j, in which case we must have $j < i$. Thus we have just this last case to handle: c_{i+1} is not in $d(u)$ and $c_{i+1} = c_j$ for some $j < i$. We now distinguish two cases: $j = 1$ and $j > 1$.

If j equals 1, i.e., if c_{i+1} equals c_1, then vertex w_i cannot be vertex v, because edge $\{u, v\}$ is not (yet) colored, whereas edge $\{u, w_i\}$ already is. Consider then the three distinct vertices u, v, and w_i. Let c be come color in $d(u)$; we know that c_1 is a color in both $d(v)$ and $d(w_i)$. Look at the subgraph of G that uses only edges colored with c or c_1. In this graph, every vertex has degree at most 2 and any vertex of degree 2 has incident upon it one edge colored c and one edge colored c_1; a path in this graph has edges of alternating colors. Two subcases arise: if there is no path from u to v in this graph, follow any maximal (i.e., not a subset of another) path or cycle from a; any such path or cycle must alternate edge colors, starting with a c_1 edge, moving to a c edge, etc. (Of course, there may be must one edge or even no edge at all.) Exchange the colors on this path or cycle; the result is still a legal edge-coloring in G' and now avoids c_1 at u, leaving us free to use it for coloring edge $\{u, v\}$. Otherwise, there is no path from w_i to one (or both) of u and v (because c_1 is used at neither w_i nor v), so follow any maximal path or cycle from w_i and exchange the colors c and c_1 in this path or cycle. Then the color c is no longer used at w_i, so that we can recolor the edge $\{u, w_i\}$ with color c and thus use color c_1 to color edge $\{u, v\}$.

If j is larger than 1, we have a comparable situation: vertex w_i cannot be vertex w_{j-1} (otherwise we would have two edges from u assigned the same color), and thus we can consider the three distinct vertices u, w_i, and w_{j-1}. Let c be come color in $d(u)$; we know that c_j is a color in both $d(w_i)$ and $d(w_{j-1})$. Again, we have two subcases (is u connected to w_{j-1} in the 2-colored subgraph) and again we follow a maximal path or cycle and exchange all colors along it to free up a color at the beginning of the path.