These notes are from https://mathematicaster.org/teaching/graphs2022/extra_04-05.pdf

Let G be a graph. A (vertex-)coloring of G is simply a function $f\colon V(G) \to C$ where C is some set. The set C is often referred to as a "palette" of colors. Intuitively, we are "coloring" each vertex of G with some color in the "palette" C. Mostly, this just gives us a colloquial way to discuss arbitrary functions on $V(G)$. We (kinda) pecified "vertex-coloring", since later we will discuss "edge-colorings" which are simply functions from $E(G)$ to some set. In general, when anyone in math mentions a "coloring" of an object, they are literally just discussing a function from said object to an arbitrary set.

For a positive integer t, a *t-coloring* (or t-vertex-coloring if such a distinction is necessary later) is a coloring where the "palette" has t colors. Since the actual nature of the "palette" doesn't usually matter, usually, a t-coloring is a function $f\colon V(G) \to [t]$. Note that, technically, not every color must be used in a t-coloring. In particular, a t-coloring is also a $(t+1)$-coloring, technically.

So far, colorings are simply functions with domain $V(G)$ and have nothing to do with the actual graph G; let's fix this. A *proper coloring* of G is a function $f\colon V(G) \to C$, where C is some set, such that $xy \in E(G) \implies f(x) \neq f(y)$. That is, no two end-points of an edge receive the same color. It is very rare that we will discuss colorings which are not proper (since we care about the graph, afterall).

Reasonably, a *proper t-coloring* of G is a proper coloring of G which uses at most t colors. Intuitively, if G has a lot of edges, then G requires quite a few colors in order to have a proper coloring. We will see that this is not exactly true, but it's a reasonable intuition.

Definition 1. *Let G be a graph. The* chromatic number *of G, denoted by $\chi(G)$, is the smallest integer t such that G has a proper t-coloring.*

A quick observation:

Observation 2. *$\chi(G) = 1$ if and only if G has no edges.*

Technically this is only true if G has at least one vertex. Your book assumes that any graph satisfies this (and I generally do as well), but sometimes the "null-graph" (the graph with no vertices) can show up if necessary. Note that the null-graph has chromatic number 0.

A couple more quick observations:

Observation 3. *If G is an n-vertex graph, then $\chi(G) \leq n$. Furthermore, $\chi(G) = n$ if and only if $G \cong K_n$.*

Observation 4. *If H is a subgraph of G, then $\chi(H) \leq \chi(G)$.*

From here, we get a simple lower-bound on the chromatic number of a graph in terms of it's clique number.

Definition 5. *The* clique-number *of a graph G, denoted by $\omega(G)$, is the size of the largest clique contained within G. Note that $\omega(G) = \alpha(\overline{G})$.*

By definition, G contains a copy of $K_{\omega(G)}$, which has chromatic number $\omega(G)$, so:

Observation 6. *$\chi(G) \geq \omega(G)$.*

This bound is far from tight. It is know that there are graphs with $\omega(G) = 2$, yet $\chi(G)$ is arbitrarily large.

Let $f\colon V(G) \to C$ be any proper coloring of G. A *color class* of this coloring is the set $f^{-1}(c)$ for any $c \in C$; i.e. the set of vertices which receive color c. Since no two adjacent vertices receive the same color, we observe that every color class is an independent set. In other words, the color classes of a proper coloring $f\colon V(G) \to C$ partition $V(G)$ into $|C|$ many independent sets. Conversely, if $V(G) = A_1 \sqcup \cdots \sqcup A_t$ is a partition wherein each A_i is an independent set of G, then we can define a proper coloring $f\colon V(G) \to [t]$ via $f(v) = i$ iff $v \in A_i$. Thus, provided we don't care about the actual nature of the colors, proper colorings are equivalent to partitions into independent sets.

Observation 7. *$\chi(G)$ is the smallest integer t such that $V(G)$ can be partitioned into t independent sets. In other words, $\chi(G)$ is the smallest integer t for which G is t-partite.*

In particular, $\chi(G) = 2$ if and only if G is bipartite and has at least one edge.

This observation leads to the following lower bound on the chromatic number:

Theorem 8. *For any n-vertex graph G,*

$$\chi(G) \geq \frac{n}{\alpha(G)}.$$

Proof. By definition, we can partition $V(G) = A_1 \sqcup \cdots \sqcup A_{\chi(G)}$ where each A_i is an independent set in G. Thus, $|A_i| \leq \alpha(G)$ since $\alpha(G)$ is the size of a largest independent set in G. We therefore bound

$$n = \sum_{i=1}^{\chi(G)} |A_i| \leq \sum_{i=1}^{\chi(G)} \alpha(G) = \chi(G)\alpha(G) \qquad \Longrightarrow \qquad \chi(G) \geq \frac{n}{\alpha(G)}.$$ □

Again, this bound is far from tight in general.

Let's now discuss an upper bound. Recall the notion of degeneracy from DS3.5.

Definition 9. *For a graph G, the* degeneracy *of G is defined to be*

$$d(G) = \max\big\{\delta(H) : H \text{ is a subgraph of } G\big\}.$$

Certainly $d(G) \leq \Delta(G)$ always. Also, we discussed that $d(G) \leq 1$ if and only if G is a forest, so degeneracy can sometimes be much smaller than the max degree.

Theorem 10. *$\chi(G) \leq d(G) + 1$. In particular, $\chi(G) \leq \Delta(G) + 1$.*

We give two proofs of this fact.

Proof. We prove the claim by induction on $n = |V(G)|$.

If $n = 1$, then $d(G) = 0$ and $\chi(G) = 1$.

Now consider $n \geq 2$ and let $v \in V(G)$ be a vertex with $\deg v = \delta(G)$; note that $\deg v \leq d(G)$. Set $H = G - v$. Since H is a subgraph of G, we have $d(H) \leq d(G)$ and so the induction hypothesis tells us that $\chi(H) \leq d(H) + 1 \leq d(G) + 1$. Thus, suppose that $f\colon V(H) \to C$ is a proper coloring of H where C is some set with $|C| = d(G) + 1$ (note that, perhaps, many colors could be unused). We show that we can extend f to a proper coloring of G. Indeed, since $|C| = d(G) + 1$ and $\deg v \leq d(G)$, there is some color $c \in C$ which is unused by the neighbors of v. Thus, setting $f(v) = c$, we know that f is now a proper coloring of G using at most $d(G) + 1$ many colors and so $\chi(G) \leq d(G) + 1$. □

Proof. In DS3.5.2, you showed that there is an ordering $V(G) = \{v_1, \ldots, v_n\}$ so that $|N(v_i) \cap \{v_1, \ldots, v_{i-1}\}| \leq d(G)$ for all $i \in [n]$. We define a proper coloring $f\colon V(G) \to [d(G)+1]$ greedily based on this ordering.

Start by setting $f(v_1) = 1$. For $i \in \{2, \ldots, n\}$, assuming that f has already been defined on $\{v_1, \ldots, v_{i-1}\}$, define $f(v_i)$ to be the smallest color not used by the already-colored neighbors of v_i (that is, the smallest color not used in $N(v_i) \cap \{v_1, \ldots, v_{i-1}\}$). Since there are $d(G)+1$ many available colors and $|N(v_i) \cap \{v_1, \ldots, v_{i-1}\}| \leq d(G)$, this is always possible.

The resulting coloring f is a proper coloring. Indeed, if $v_iv_j \in E(G)$ where $i < j$, then $v_i \in N(v_j) \cap \{v_1, \ldots, v_{j-1}\}$ and so $f(v_i) \neq f(v_j)$ by construction.

Thus, since f uses at most $d(G)+1$ many colors, we've shown that $\chi(G) \leq d(G)+1$. □

Degeneracy can oftentimes be difficult to work with theoretically, so it's usually best to just keep the bound $\chi(G) \leq \Delta(G) + 1$ in mind.

Let's finish today by proving a general upper bound on the chromatic number of a triangle-free graph in terms of its number of vertices. A graph is triangle-free if it has no copy of K_3; equivalently $\omega(G) \leq 2$.

Theorem 11. *If G is an n-vertex, triangle-free graph, then $\chi(G) \leq \sqrt{2n}$.*

Proof. We prove the claim by induction on n.

If $n = 1$, then $\chi(G) = 1 \leq \sqrt{2}$, so suppose that $n \geq 2$.

If $\Delta(G) \leq \sqrt{2n} - 1$, then we would have $\chi(G) \leq \Delta(G) + 1 \leq \sqrt{2n}$ as needed; thus we may suppose that $\Delta(G) > \sqrt{2n} - 1$. Fix any vertex $v \in V(G)$ with $\deg v > \sqrt{2n} - 1$.

Since G is triangle-free, we know that $N(v)$ must be an independent set in G. If $V(G) = N(v) \cup \{v\}$, then G is bipartite and so $\chi(G) \leq 2 \leq \sqrt{2n}$ since $n \geq 2$, so we may suppose that this is not the case.

Set $H = G - (N(v) \cup \{v\})$ (deleting vertices). We claim that $\chi(G) \leq 1 + \chi(H)$.

Let $A_1, \ldots, A_{\chi(H)}$ be a partition of $V(H)$ into $\chi(H)$-many independent sets; note that $\chi(H) \geq 1$ since H has some vertices. H is an induced subgraph of G and so each A_i is also an independent set in G. Now, v has no neighbors in H and so $A_1 \cup \{v\}$ is also an independent set. Finally, $N(v)$ is an independent set, and so $A_1 \cup \{v\}, A_2, \ldots, A_{\chi(H)}, N(v)$ is a partition of $V(G)$ into $\chi(H)+1$ many independent sets and so $\chi(G) \leq 1 + \chi(H)$ as claimed.

Now, H is a subgraph of G and so is also triangle-free. Furthermore, H has $n - \deg v - 1 < n$ many vertices, so the induction hypothesis tells us that

$$\chi(H) \leq \sqrt{2(n - \deg v - 1)} < \sqrt{2(n - \sqrt{2n})},$$

since $\deg v > \sqrt{2n} - 1$. Therefore, since $n \geq 2$,

$$\chi(G) \leq 1 + \chi(H) < 1 + \sqrt{2(n - \sqrt{2n})} < 1 + \sqrt{2n - 2\sqrt{2n} + 1} = 1 + \sqrt{(\sqrt{2n} - 1)^2} = \sqrt{2n}. \quad \square$$

Note that our proof actually showed that $\chi(G) < \sqrt{2n}$, unless $G \cong K_2$. It turns out that the bound in Theorem 11 is not too far from the truth. It is know that if n is large, then

$$C_1\sqrt{\frac{n}{\log n}} \leq \max_{G \text{ triangle-free on } n \text{ vertices}} \chi(G) \leq C_2\sqrt{\frac{n}{\log n}},$$

for some constants C_1, C_2.