These notes are from https://mathematicaster.org/teaching/graphs2022/extra_03-29.pdf

Let $G = (V, E)$ be a graph. A *matching* is a set of edges which are vertex-disjoint. Equivalently, a matching is an independent set in the line graph $L(G)$. Note that the empty-set is trivially a matching, so every graph has a matching. If $M \subseteq E$ is a matching and $e \in M$, we say that the end-points of the edge e are *matched*.

For a subset $A \subseteq V$, a matching $M \subseteq E$ is said to *saturate* A if every vertex in A is incident to some edge in M. We may also say that M *covers* A, but saturates tends to be the more common jargon. A *perfect matching* is a matching which saturates V; that is, every vertex is an end-point of some edge of the matching. Note that if G has a perfect matching, then this matching has exactly $|V|/2$ many edges; in particular, $|V|$ must be even in order to allow this. Conversely, a matching in G is a perfect matching if and only if it has exactly $|V|/2$ many edges.

We are often interested in the largest matching in a graph, or if the graph has a matching which saturates a particular set. We now introduce a few more terms and notation (some of which may seem unrelated at this point).

1. The *independence number* of G, denoted by $\alpha(G)$ is the size of a largest independent set in G.

2. The *matching number* or *edge-independence number* of G, denoted by $\alpha'(G)$ is the size of a largest matching in G. Equivalently $\alpha'(G)$ is the size of the largest independent set in the line graph $L(G)$how to authenticate.

3. A *vertex-cover* of G is a subset of vertices $B \subseteq V$ such that every edge of G has at least one end-point in B. Equivalently, $B \subseteq V$ is a vertex-cover of G if and only if either $B = V$ or $G - B$ (deleting vertices) has no edges.

   The *vertex-cover number* of G, denoted by $\beta(G)$, is the size of a smallest vertex-cover of G.

4. An *edge-cover* of G is a subset of edges $S \subseteq E$ such that every vertex of G is incident to at least one edge in S. Equivalently, S is an edge-cover of G if and only if the subgraph (V, S) has no isolated vertices.

   The *edge-cover number* of G denoted by $\beta'(G)$, is the size of a smallest edge-cover of G. Note that $\beta'(G)$ only makes sense if G has no isolated vertices since an isolated vertex cannot be covered by any edge.

Generally in graph theory, the use of a prime ($'$) denotes an edge-version of a parameter usually defined based mainly on the vertices of a graph (what this actually means varies from case-to-case). Oftentimes, this is the same parameter of the line graph (for instance, $\alpha'(G) = \alpha(L(G))$), but not always (for instance, $\beta'(G) \neq \beta(L(G))$ in general, e.g. $G = K_{1,3}$).

As a quick example, for K_3, we have $\alpha = \alpha' = 1$ and $\beta = \beta' = 2$. In general, for K_n ($n \geq 2$), we have $\alpha = 1, \alpha' = \lfloor n/2 \rfloor, \beta = n - 1, \beta' = \lceil n/2 \rceil$ (convince yourself that these are true to make sure you understand the concepts).

We will work on relating these four parameters. In particular, we will prove the following:

1. $\alpha'(G) \leq \beta(G) \leq 2\alpha'(G)$ for any graph.

2. $\alpha'(G) = \beta(G)$ if G is bipartite (this is probably the most important in the list and is known as Kőnig's theorem).

3. $\alpha'(G) + \beta'(G) = n$ for any n-vertex graph with no isolated vertices.

4. $\alpha(G) + \beta(G) = n$ for any n-vertex graph (this is a homework exercise).

Items 3 and 4 formalize the intuition that α, β and α', β' are "complementary".

Note that items 1, 3 and 4 together imply that $\alpha(G) \leq \beta'(G) \leq (n + \alpha(G))/2$ if G has no isolated vertices.

We begin with α' and β.

Theorem 1. *For any graph G, we have $\alpha'(G) \leq \beta(G) \leq 2\alpha'(G)$.*

Proof. Let $M \subseteq E(G)$ be a maximum matching, so $|M| = \alpha'(G)$.

Consider letting C consist of all end-points of the edges of M. We claim that C is a vertex-cover of G. Indeed, if this were not the case, then there would be some edge $e \in E(G)$ which has neither end-point in C. But then, $M \cup \{e\}$ would be a strictly larger matching, which is impossible. Thus, $\beta(G) \leq |C| = 2|M| = 2\alpha'(G)$.

On the other hand, let $C \subseteq V(G)$ be a minimum vertex-cover of G, so $|C| = \beta(G)$. Certainly C must contain at least one end-point of each edge in M. Since all of these end-points are distinct, we must have $\beta(G) = |C| \geq |M| = \alpha'(G)$. □

It is possible that $\alpha'(G) = \beta(G)$ (see Kőnig's theorem below). It is also possible that $\beta(G) = 2\alpha'(G)$, e.g. K_n when n is odd (where we have $\beta = n - 1$ and $\alpha' = \lfloor n/2 \rfloor = (n-1)/2$).

Problem 1 (Extra-challenging problem)**.** Prove that $\beta(G) = 2\alpha'(G)$ if and only if each connected component of G is an odd-clique.

I really don't want to lead you astray and waste your time. You don't have the tools necessary to actually prove this (unless I'm missing something). The proof I have in mind (though perhaps there's a simpler one) requires HW9.4 along with the Tutte–Berge formula. We won't cover the Tutte–Berge formula in this class, though it's a lovely theorem. If you have the time and motivation, I encourage you to look into Tutte–Berge and figure out how to use it to prove the above problem. I've included a proof of the above problem which invokes Tutte–Berge (as a black-box) on the last page of these notes in case you care.

Theorem 2 (Kőnig's Theorem)**.** *If G is a bipartite graph, then $\alpha'(G) = \beta(G)$. In other words, for bipartite graphs, the size of a maximum matching is equal to the size of a smallest vertex-cover.*

Proof. Suppose that G has parts A and B. Recall the notion of A-B paths from the notes from 03-03; we will rely on the workhorse of Lemma 1 in those notes.

Since G has parts A and B, we observe that A-B paths correspond exactly to the edges of G. In other words, a set of disjoint A-B paths corresponds exactly a matching in G and so $\alpha'(G) = p_G(A, B)$; the maximum number of disjoint A-B paths.

On the other hand, recall the notation of an A-B separator from the notes from 03-03 which is a set of vertices so that every A-B path contains at least one of these vertices. Observe that $U \subseteq V(G) = A \sqcup B$ is an A-B separator if and only if every edge of G is incident to some vertex of U. In other words, an A-B separator in G is precisely a vertex-cover of G and so $\beta(G) = \kappa_G(A, B)$; the size of a minimum A-B separator.

Applying Lemma 1 from the notes from 03-03, we conclude that

$$\alpha'(G) = p_G(A, B) = \kappa_G(A, B) = \beta(G).$$

□

We will use Kőnig's Theorem next lecture in order to prove, arguably, the most important theorem about matchings: Hall's Marriage Theorem.

But, for now, we seek to prove that $\alpha'(G) + \beta'(G) = n$ for any n-vertex graph with no isolated vertices. Recall that $\beta'(G)$ only makes sense if G has no isolates since we cannot cover an isolated vertex with an edge.

In order to accomplish this feat, we need to understand the structure of a minimum edge-cover.

Lemma 3. *Let $G = (V, E)$ be a graph with no isolated vertices. If $S \subseteq E$ is a minimum edge-cover of G, then (V, S) is a forest with no isolated vertices.*

Proof. Since S is an edge-cover, we cannot have any isolated vertices in (V, S) since an isolated vertex would be an uncovered vertex.

Suppose for the sake of contradiction that (V, S) contains a cycle; label the vertices of one of these cycles as $(v_1, \ldots, v_k)$, of course $k \geq 3$. Consider the edge $e = v_1v_2$, which covers only the vertices v_1 and v_2. However, $v_1v_k \in S$ covers v_1 and $v_2v_3 \in S$ covers v_2. That is to say, $S \setminus \{e\}$ is also an edge-cover of G, contradicting the minimality of S. □

In fact, more is true. One can prove that every connected component of this forest is a star (that is, isomorphic to $K_{1,\ell}$ for some $\ell \geq 1$), but this extra structure will be unnecessary for our arguments. However, it's a good exercise! The key lemma that one needs to establish in order to get this extra structure is the following: If G is a connected graph with no isolated vertices which contains no copy of K_3 nor of P_4, then G is a star.

Theorem 4. *If G is an n-vertex graph with no isolated vertices, then $\alpha'(G) + \beta'(G) = n$.*

Proof. We prove the two inequalities ($\leq n$ and $\geq n$) separately. Before we dive in, let me walk through the intuition.

Suppose that we start with maximum matching (which has size $\alpha'(G)$); if we can somehow use this matching to build a edge-cover with at most $n - \alpha'(G)$ many edges, then also $\beta'(G) \leq n - \alpha'(G)$.

Suppose that we start with a minimum edge-cover (which has size $\beta'(G)$; if we can somehow use this edge-cover to build a matching with at least $n - \beta'(G)$ many edges, then also $\alpha'(G) \geq n - \beta'(G)$.

This same intuition (with changed terms and symbols) will help you work through HW9.4.

To begin, we show that $\alpha'(G) + \beta'(G) \leq n$.

Let $M \subseteq E(G)$ be a maximum matching of G, so $|M| = \alpha'(G)$. Let A be the set of end-points of M and let $B = V(G) \setminus A$. Note that $|A| = 2|M| = 2\alpha'(G)$ and so $|B| = n - 2\alpha'(G)$. Now, since no vertex of B is isolated, we can arbitrarily select one edge incident to each vertex in B to build a set $S \subseteq E(G)$ with $|S| \leq |B|$ such that S covers B.[1] Since M covers A, S covers B and $V(G) = A \sqcup B$, we thus know that $M \sqcup S$ is an edge-cover of G. Therefore,

$$\alpha'(G) + \beta'(G) \leq \alpha'(G) + |M \sqcup S| = \alpha'(G) + |M| + |S| \leq \alpha'(G) + \alpha'(G) + (n - 2\alpha'(G)) = n.$$

We now show that $\alpha'(G) + \beta'(G) \geq n$, which will conclude the proof.

Let $S \subseteq E(G)$ be a minimum edge-cover of G, so $|S| = \beta'(G)$. Note: we know that S exists since G has no isolated vertices. Applying Lemma 3, we know that (V, S) is a forest with no isolated vertices; let $G_1, \ldots, G_k$ denote the connected components of this forest. Since (V, S) has

[1] One can actually show that $|S| = |B|$ since otherwise we could extend the matching M (why?). However, simply knowing that $|S| \leq |B|$ suffices for our arguments.

no isolated vertices, we know that each G_i has at least one edge. We may therefore select one edge from each G_i to build a matching with k edges (since the connected components are vertex-disjoint), so $\alpha'(G) \geq k$. Additionally, (V, S) is a forest with k connected components and so it has exactly $n-k$ many edges, so $|S| = n-k$. Therefore,

$$\alpha'(G) + \beta'(G) \leq k + (n-k) = n. \qquad \square$$

As promised, here is a proof of Problem 1, which uses HW9.4 and the Tutte–Berge formula as black-boxes.

To begin, we should state the Tutte–Berge formula. For a graph G, let $\text{odd}(G)$ denote the number of connected components of G which have an odd number of vertices.

Theorem 5 (Tutte–Berge)**.** *Let $G = (V, E)$ be any graph. Then*

$$\alpha'(G) = \frac{1}{2} \cdot \min_{U \subseteq V} \Big(|V| + |U| - \text{odd}(G - U) \Big).$$

This is yet another case where the "obvious" necessary condition is sufficient. Can you see what this "obvious" necessary condition is? (I don't claim that it's literally obvious, just "obvious" in the sense that if you think about matchings in just the right way, it'll pop out.)

Proof of Problem 1. We have already remarked that $\beta(G) = 2\alpha'(G)$ if $G \cong K_n$ where n is odd. If G has connected components $G_1, \ldots, G_k$, then $\beta(G) = \sum_{i=1}^k \beta(G_i)$ and $\alpha'(G) = \sum_{i=1}^k \alpha'(G_i)$ and so $\beta(G) = 2\alpha'(G)$ if every connected component of G is an odd-clique.

Now suppose that $\beta(G) = 2\alpha'(G)$ and that G has n vertices.

Now, the Tutte–Berge formula tells us that there is some $U \subseteq V$ such that

$$\alpha'(G) = \frac{1}{2} \Big(n + |U| - \text{odd}(G - U) \Big).$$

HW9.4 tells us that $\alpha(G) + \beta(G) = n$ and so $\alpha(G) = n - \beta(G) = n - 2\alpha'(G)$ by assumption. Therefore,

$$\alpha(G) = n - 2|M| = n - 2\alpha'(G) = \text{odd}(G - U) - |U|.$$

Suppose that the connected components of $G - U$ are $G_1, \ldots, G_\ell$; of course $\ell \geq \text{odd}(G - U)$ (we could have a strict inequality here if some components have an even number of vertices). Of course, there can be no edges in G between these ℓ components, so we find that $\alpha(G) \geq \sum_{i=1}^\ell \alpha(G_i)$ since we could simply combine independent sets from the individual components. Combining these observations, we find that

$$\text{odd}(G - U) - |U| \geq \sum_{i=1}^{\ell} \alpha(G_i) \implies |U| \leq \text{odd}(G - U) - \sum_{i=1}^{\ell} \alpha(G_i).$$

Certainly $\alpha(G_i) \geq 1$ for each $i \in [\ell]$ and so $\sum_{i=1}^\ell \alpha(G_i) \geq \ell \geq \text{odd}(G - U)$. This implies that $|U| \leq 0 \implies |U| = 0$. In particular, $U = \varnothing$ and so $G - U = G$; thus $G_1, \ldots, G_\ell$ are the connected components of G. Since certainly $|U|$ cannot be negative, we, in fact, have shown that $\sum_{i=1}^\ell \alpha(G_i) = \ell = \text{odd}(G)$. In particular, each of $G_1, \ldots, G_\ell$ have an odd number of vertices. Furthermore, $\alpha(G_i) = 1$ for all $i \in [\ell]$ and so each G_i is a clique. This concludes the proof. $\square$