MATH CGT: CATALAN NUMBERS AND GRAPH THEORY

HEESUNG YANG

Abstract. This notes covers Chapters 8.1 and most of Chapters 11 and 12 of Brualdi's *Introductory Combinatorics*, 5th edition, the chapters that are part of Dalhousie's combinatorics comprehensive exam syllabus but was not covered in MATH 5370.

1. Chapter 11.1: Basic properties of graphs

Definition 1.1. A *graph* (or *simple graph*) $G = (V, E)$ is an object consisting of two types of objects:

- a finite set called *vertices* $V = \{a, b, c, \dots\}$, and
- a set E of pairs of distinct vertices called *edges*.

The *order* of the graph G is $|V|$, i.e., the number n of vertices in the set V.

Definition 1.2. Suppose that $\alpha = \{x, y\} = \{y, x\} \in E$ (i.e., α is an edge of G). Then we say that x and y are *adjacent* (or α *joins* x *and* y). Then x and y are the *vertices of the edge* α. In this case, x and α are said to be *incident*. Similarly, y and α are also *incident*.

Definition 1.3. A curve is *simple* if the curve is not self-intersecting.

Remark. We can view graphs geometrically. Geometrically speaking, for any point x and a distinct point y, we only connect those two points with a simple curve.

Our current definition of graphs is rather restrictive in the sense that a pair of vertices can only form one edge, so a new definition is required to loosen this restriction.

Definition 1.4. $G = (V, E)$ is a *multigraph* is a graph whose pair of vertices may form more than one edge.

Remark. If G is a multigraph, then E is a multi-set, since there are more than one $\{x, y\}$ whenever x and y have more than one edge.

Definition 1.5. Suppose that $G = (V, E)$ is a multigraph, and that $x, y \in V$. If there are m edges incident to x and y, then m is called the *multiplicity of an edge* $\alpha = \{x, y\}$; we write $m\{x, y\}$ to indicate how many edges there are between x and y in E. If $x = y$, then the edges become of the form $\{x, x\}$. Such edges are called *loops*, A multigraph where loops are allowed is called a *general graph*.

Date: 2 May 2019.

Definition 1.6. A graph of order n is called *complete* if every pair of distinct vertices forms an edge, and we denote such graph K_n. Therefore any complete graph of order n has $\binom{n}{2} = n(n-1)/2$ edges. Conversely, if a graph of order n has no edges, then such graph is called the *null graph* of order n, and it is denoted by N_n.

Before introducing a particular type of graph, we will draw a few complete graphs (say, K_1, K_2, K_3, K_4, K_5).

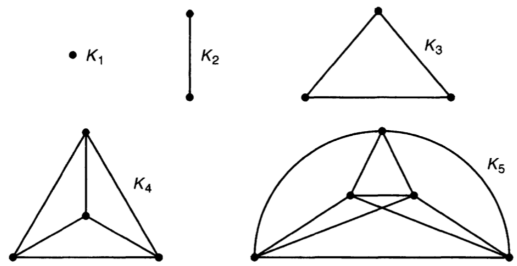

It is not hard to see that, up to $n = 4$, it is possible to draw them so that K_n has no overlapping edges (i.e., no two edges cross each other upon drawing them at a point that is not a vertex). However, there are always at least overlapping edges for K_5. We will categorically define a sub-class of graphs that have no such overlapping edges.

Definition 1.7. A general graph G is *planar* if G can be geometrically represented so that there are no two overlapping edges. Such drawing of G is said to be a *planar graph*, and that graph is called a *planar representation of G*.

Example. K_n is planar if and only if n is one of $1, 2, 3, 4$. For any $n \geq 5$, K_n is not planar.

Recall that we defined the notion of multiplicity on the edges, based on how many edges there are between a pair of vertices. We can also define a similar notion for each vertex.

Definition 1.8. The *degree* of a vertex x in a general graph G is the number of edges that are incident with x, and we write $\deg(x)$. Any loop of x contributes 2 to $\deg(x)$. The list of degrees of each vertex in a graph G in non-increasing order is called the *degree sequence of G*.

Example. The degree sequence of K_n is $(n-1, n-1, \ldots, n-1)$ (repeated n times).

Theorem 1.1. *For any general graph G, the sum $\sum d_i$ of the degrees of all the vertices is of G is always even. Consequently, the number of vertices of G with odd degree is even.*

Proof. Every edge increases the degree by two – one for each of the two vertices that this edge connects. Therefore the sum is always even, and this can happen only when the number of "odd vertices" is even. □

Example. Consider handshaking at a party. One can represent the handshaking in a graph format, by creating an edge between two vertices (guests) if two people shook hands (edges). Applying the above theorem gives that the number of handshaking is always even, and that there must be an even number of people who shook hands odd number of times.

In abstract algebra, one of our main interests is examining structural similarities between two mathematical objects. If they are structurally identical, then the two objects are said to be isomorphic. Thus, it is natural to wonder if the two graphs that look seemingly different geometrically are in fact the same, which prompts the notion of isomorphism in graph-theoretic sense.

Definition 1.9. Let $G = (V, E)$ and $G' = (V', E')$ be two general graphs. Then G and G' are said to be *isomorphic* if there is a bijective map $\theta : V \to V'$ so that for any $x, y \in V$, there are $\theta(x), \theta(y) \in V'$ such that the number of edges joining x and y and the number of edges joining $\theta(x)$ and $\theta(y)$ match. Such map θ is called an *isomorphism* of G and G'.

Just as in abstract algebra, to prove that the two graphs are isomorphic, one can display a bijective map θ, and demonstrate that θ is an isomorphism. To show why two graphs are not isomorphic, one can demonstrate a characteristic not shared by the two graph (e.g. one graph has three vertices of degree 3, whereas the other only has two instead).

It should be noted that having the same degree sequence does not necessarily guarantee an isomorphism. For instance, it is possible to construct two graphs of order 6 with degree sequence $(3, 3, 3, 3, 3, 3)$. Note that the left graph has a cycle of length 3, but the right graph does not have any cycle of length 3.

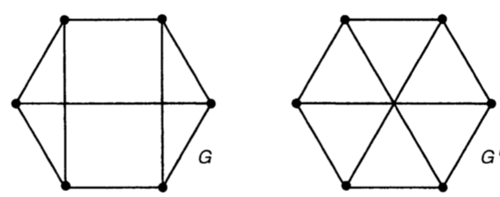

Thus, we can conclude that any two isomorphic graphs must have the same degree sequence, but its converse is false.

Definition 1.10. If $G = (V, E)$ is a general graph, and this sequence consists of edges $\{x_0, x_1\}, \{x_1, x_2\}, \ldots, \{x_{n-1}, x_n\}$ is a *walk of length* n, and this walk is said to *join the vertices* x_0 *and* x_n. We denote the walk by $x_0 - x_1 - \cdots - x_{n-1} - x_n$. In particular, if $x_0 = x_n$, then the walk is said to be *closed*; otherwise, the walk is *open*. Suppose that each edge is distinct. Then that walk is called a *trail*. Furthermore, if a walk has distinct vertices, then that walk is a *path*. A closed path is said to be a *cycle*.

Definition 1.11. Any graph G is called *connected* if there is a walk from any two vertices of G. Otherwise, G is *disconnected*. In other words, if G is disconnected, then there exists a pair of vertices x and y where it is impossible to reach y from x (or vice versa). The length of a shortest walk (hence a path) from x to y is called the *distance* between x and y and is denoted by $d(x, y)$.

Definition 1.12. Let $G = (V, E)$ be a general graph, and $H = (U, F)$ where $U \subseteq V$ and $F \subseteq E$. Then H is a *general subgraph* of G. If F consists of all the edges whose endpoints are the vertices from U, then H is a *subgraph induced by* U, and is denoted G_U. If $V = U$, but $F \subseteq E$, then H is said to be a *spanning subgraph* of G. Suppose that $V_1, \ldots, V_k$ form a partition of V, and that $G_{V_1}, \ldots, G_{V_k}$ are all connected. Then each G_{V_i} is a *connected component* of G.

Now we are ready to state necessary conditions for two general graphs to be isomorphic.

Theorem 1.2. *If G and G' are two general graphs, then the following need to hold in order for G and G' to be isomorphic.*

1. *If G is a graph, then so is G'.*
2. *If G is connected, so is G'. Also, G and G' must have the same number of connected components.*
3. *If G has a cycle of length equal to some integer k, then so must G'.*
4. *If G has an induced general subgraph which is K_n of order n, then so does G'.*

Finally, we make a brief remark that any graph can be represented in matrix form, namely with an *adjacency matrix*. Thus, if v_i and v_j are vertices of a graph G and are adjacent, then in the (ij)-th entry of its adjacency matrix must have the number of edges between these two vertices. Clearly, any adjacency matrix must be symmetric. Furthermore, if A and A' are adjacency matrices of G and G' respectively, then G and G' are isomorphic if and only if there is an invertible matrix D such that $A' = DAD^{-1}$.

2. Chapter 11.2: Eulerian trails

In this section, we will examine what condition(s) must a graph satisfy in order for one to draw the entire graph without lifting a pencil. Specifically, we will prove that this can be done if and only if every vertex has even degree (for a closed Eulerian trail to exist) or exactly two of the vertices is of odd degree (for an open Eulerian trail to exist).

Definition 2.1. A trail of in a general graph is said to be *Eulerian* if that trail contains every edge of G.

Lemma 2.1. *Let $G = (V, E)$ be a general graph, and assume that the degree of each vertex is even. Then each edge of G belongs to a closed trail, and hence to a cycle.*

Theorem 2.1. *Let G be a connected general graph. Then G has a closed Eulerian trail if and only if the degree of each vertex is even.*

Proof. ($\Leftarrow$) Suppose every vertex of G has even degree; we will give a constructive proof. Let $G_1 = (V, E_1) = G$. Pick some $\alpha_1 \in E_1$, and by the previous lemma there is a closed trail containing the edge α_1. Let $E_2 = E_1 \setminus F_1$, where F_1 is the set of edges that belong to γ_1. Thus by the above lemma, removing the edges reduces the degree of each involved vertex by even number. In conclusion, all vertices in $G_2 = (V, E_2)$ are of even degree. Assuming that E_2 is non-empty, pick a vertex v_1 which is part of γ_1, and an edge $\alpha_2 \in E_2$ which is incident to v_1. Again apply the above lemma to obtain a closed trail containing the edge α_2. We can now concatenate γ_1 and γ_2 to obtain a closed trail. First, start from α_1, and travel through γ_1 until v_1 is hit for the first time; now traverse through the closed trail γ_2, and once the trail arrives at v_1 for the second time, travel through the remaining edges of γ_1. Repeating this algorithm until the graph $G_k = (V, E_k)$ satisfies $E_k = \emptyset$ yields a closed Eulerian trail for G.

($\Rightarrow$) Suppose that G has a closed Eulerian trail. Observe that from one edge, we visit a vertex that is not a starting vertex, and then through another edge we leave that vertex. As for the starting vertex, we leave at the beginning, but we arrive back at the end, which concludes the trail (note it is still possible, however, for us to enter the vertex and then leave the vertex in the middle of a trail, just like any other ordinary vertices after the initial departure). Thus for any vertex v, two edges can be paired up: one used to visit v and one used to leave v – hence, every vertex must have even degree. $\square$

Theorem 2.2. *Let G be a connected general graph. Then G has an open Eulerian trail if and only if there are exactly two vertices u and v which is of odd degree. Particularly, every open Eulerian trail joins u and v.*

Proof. Recall that every general graph has even number of vertices of odd degree. Therefore it suffices to argue that there cannot be more than two vertices of degree.

($\Leftarrow$) Suppose that $G = (V, E)$ has exactly two vertices of odd degree u and v. Then observe that $G' = (V, E \cup \{\{u, v\}\})$ is a general graph where every vertex has even degree. Thus there is a closed Eulerian trail starting from v with the starting edge $\{v, u\}$; call this Eulerian trail γ. Then removing the initial edge $\{v, u\}$ gives us an open Eulerian trail from u to v.

($\Rightarrow$) Suppose that $G = (V, E)$ has an open Eulerian trail. Observe that the observation on a pair of edges for each non-starting and non-terminating vertex holds, so every vertex other than the initiating vertex and the terminating vertex must be of even degree. Therefore, an open Eulerian trail must start from u, an odd-degree vertex, and end at v, the other odd-degree vertex. In order for every other vertex to be even except for u and v, it follows that u and v must be adjacent. $\square$

3. Chapter 11.3: Hamiltonian paths and cycles

Definition 3.1. Let G be a graph of order n. Then a cycle of G is a *Hamiltonian cycle* if the cycle hits every single vertex in G (hence is of length n). A path is a *Hamiltonian path* in G if it is a path hitting every single vertex in G (hence is of length $n - 1$).

Remark. It follows from the definition that every edge of a Hamiltonian path or a Hamiltonian cycle is distinct. Furthermore, the existence of a Hamiltonian path or a Hamiltonian cycle wholly depends on which pairs of vertices are adjacent, regardless of multiplicity. Hence, it suffices to consider simple graphs only in this section.

Example. For any complete graph K_n, we see that there are $n!$ Hamiltonian paths (just permute the vertices $v_1, \ldots, v_n$), and that there are $(n-1)!$ Hamiltonian cycles (circular permutations of n vertices).

We first discuss one condition that guarantees a graph *not* to have a Hamiltonian cycle.

Definition 3.2. An edge of a connected graph G is called a *bridge* if the removal of that edge renders the new graph disconnected. In other words, a bridge is an edge whose removal creates more than one connected component.

Theorem 3.1. *A connected graph of order $n \geq 3$ that has a bridge cannot have a Hamiltonian cycle.*

Proof. Suppose that $\alpha = \{x, y\}$ is a bridge of a connected graph G. Let $G' = (V, E \backslash \{\{x, y\}\})$. Then G' has two connected components. Now suppose that G has a Hamiltonian cycle γ. Then γ necessarily starts in one of the connected components in G', eventually travels through the vertices of the other connected component, and then come back to the first connected component to complete the cycle. But this necessitates that γ travel the bridge edge α twice, but this contradicts the fact that γ is a Hamiltonian cycle. □

Now we discuss a sufficient condition for a graph to have a Hamiltonian cycle.

Definition 3.3. Suppose G is a graph of order n. Then we say G *has the Ore property* if for any pair of distinct non-adjacent vertices x and y, we have $\deg(x) + \deg(y) \geq n$.

Lemma 3.1. *Any disconnected graph G cannot satisfy the Ore property.*

Proof. Suppose that G has more than one connected component. Then there is a partition of the set of vertices, say U and W where $|U| = r$ and $|W| = s$. Then each vertex from U can have degree at most $r - 1$ and that from W at most $s - 1$. So if $x \in U$ and $y \in W$, then $\deg(x) + \deg(y) \leq (r - 1) + (s - 1) = (r + s) - 2 = n - 2 < n$. We found a pair of vertices not satisfying the Ore property. □

Theorem 3.2. *Any connected graph G of order n satisfying the Ore property has a Hamiltonian cycle.*

Corollary 3.1. *A graph of order $n \geq 3$ whose vertices have degree at least $n/2$ has a Hamiltonian cycle.*

Proof. The given condition automatically implies that a graph satisfies the Ore property. □

Theorem 3.3. *A graph of order n in which the sum of the degrees of each pair of non-adjacent vertices is at least $n - 1$ has a Hamiltonian path.*

Proof. Exercise – tweak the algorithm used to prove that any graph with the Ore property has a Hamiltonian cycle (see Brualdi, Theorem 11.3.2). □

4. Chapter 11.4: Bipartite multigraphs

Definition 4.1. Let $G = (V, E)$ be a multigraph. Then G is *bipartite* if there is a bipartition of vertices X and Y such that every edge in G connects one vertex in X to one vertex in Y. Thus, any two vertices in the same bipartition are not adjacent. We call X the *left vertices* and Y the *right vertices*.

Remark. It follows from the definition that no bipartite multigraph can contain loops. A bipartite graph G is *complete* with bipartition X and Y if any vertex from X is adjacent to each vertex in Y. The complete bipartite graph with m vertices in X and n vertices in Y is denoted $K_{m,n}$.

Theorem 4.1. *A multigraph is bipartite if and only if each of its cycle has even length.*

Proof. It suffices to assume that G is connected; if G is not connected, then one can apply the argument in this proof to each connected component of G.

($\Leftarrow$) Suppose that each cycle of a connected graph G has even length. Let x be a vertex of G, and let G be connected. Let X be a set of vertices whose distance from x is even; let Y be a set of vertices whose distance from x is odd. We see that X and Y form a bipartition. Clearly, X and Y are disjoint, so we only need to prove that no two vertices in the same set can be adjacent. Suppose that $a, b \in X$, and that $\{a, b\}$ is an edge. Then $d(x, a)$ and $d(x, b)$ are both even. Suppose that α and β are walks of length $d(x, a)$ and $d(x, b)$ respectively. If α and β have no common vertex other than x, then there is a cycle starting at x of length $d(x, a)+d(a, b)+d(b, x) = d(x, a)+1+d(b, x)$. Since $d(x, a)$ and $d(b, x)$ have the same parity, it follows that $d(x, a) + 1 + d(b, x)$ is odd, contradicting the assumption that any cycle of G must have even length. Now suppose that α and β have at least one common vertex other than x. Particularly, let z be the last common vertex of α and β. Break α into two parts: α_1 from x to z and α_2 from z to a. Similarly, break β into two parts: β_1 from x to z and β_2 from z to b. Thanks to the way z is chosen, there cannot be any common vertex besides z between α_2 and β_2. We claim that the length of α_1 and β_1 need be the same. Suppose otherwise – without loss of generality, suppose α_1 is shorter than β_1. If this is the case, then concatenating α_1 and β_2 will create a shorter walk from x to b, and this contradicts the minimality of β. Hence, the length of α_2 and that of β_2 must have the same parity. Note that the concatenation of $\alpha_2, \{a, b\}$, and β_2 form a cycle of odd length, which cannot happen. Thus no two vertices from X can be adjacent, and the similar reasoning shows that no two vertices from Y can be adjacent, as required.

($\Rightarrow$) Suppose that a connected graph G is bipartite with bipartition X and Y; let γ be a cycle. Furthermore, without loss of generality, assume that γ starts from a vertex in X. Then any walk γ alternates between X and Y, so any walk of odd length ends up at a vertex in Y, and any walk of even length ends up at a vertex in X. Thus, any cycle of a bipartite multigraph must be of even length. ☐

Theorem 4.2. *Let G be a bipartite graph with bipartition X and Y. If $|X| \neq |Y|$, then G does not have a Hamiltonian cycle. If $|X| = |Y|$, then G does not have a Hamiltonian cycle beginning at a vertex X and ends at a vertex in X. If $|X|$ and $|Y|$ differ by at least 2, then G cannot have a Hamiltonian path. If $|X| = |Y| + 1$, then G does not have a Hamiltonian path that begins at X and ends at Y, or vice versa.*

5. Chapters 11.5 & 11.7: Trees

Definition 5.1. A connected graph G is a *tree* if removing any one edge results in a disconnected graph. In other words, G is a tree if every edge of G is a bridge.

The remainder of this section is devoted to examining different ways of characterizing trees.

Theorem 5.1. *A connected graph of order n has at least $n-1$ edges. Moreover, for each positive integer n, there exist connected graphs with exactly $n-1$ edges. Removing any edge from a connected graph of order n with exactly $n-1$ edges leaves a disconnected graph, and hence each edge is a bridge.*

Proof. Start with n vertices but with no edge. Adding an edge decreases the number of connected components by at most 1. If the two already connected vertices are connected, then there is no reduction in the number of connected components; otherwise, the number decreases by 1. Thus, in order to decrease the number of connected components from n to 1, there must be at least $n-1$ edges in order for a graph of order n to be connected. □

In light of the following theorem, we obtain the first alternative characterization of trees.

Theorem 5.2. *A connected graph G is a tree if and only if G has exactly $n-1$ edges.*

Proof. ($\Leftarrow$) This is immediate from the previous theorem.

($\Rightarrow$) Suppose G is a tree of order n. We prove by induction on n. If $n=1$, then G has no edges, so is indeed a tree. Now suppose that the claim holds for all $1 \leq k < n$. Let G be a tree of order n, and let α be an edge of G. Let G' be the graph of G but with α removed. Since G is a tree, removing α results in two connected components; say the two connected components have k vertices and l vertices respectively. Then by the inductive hypothesis, the two connected components have $k-1$ and $l-1$ edges, respectively. Hence, G has $(k-1)+(l-1)+1=(k+l)-1$. But then $k+l=n$, so G has $n-1$ edges as required. □

One can also characterize trees in terms of cycles. We first need to prove a lemma before stating this characterization.

Lemma 5.1. *If G is a connected graph with an edge $\alpha = \{x, y\}$, then α is a bridge if and only if there is no cycle of G containing α.*

Proof. ($\Rightarrow$) Suppose that α is a bridge. Then G must consist of two connected components that are only connected by α; so if there were to be a cycle including α, then the two connected components must be held together by another edge in that cycle since a cycle cannot have repeated edges. Therefore α cannot be contained in any cycle.

($\Leftarrow$) Suppose that α is not a bridge. Then even when α is removed, the new graph still remains connected. Therefore, for any x and y, there is a path in the new graph (hence in the original graph), say $x - \cdots - y$. Add α at the end, i.e., $x - \cdots - y - x$; this creates a cycle containing α, as desired. □

Theorem 5.3. *A connected graph G is a tree if and only if G has no cycles.*

Proof. ($\Rightarrow$) Suppose that G is a tree, and let α be an edge of G. Since α is a bridge, it follows that there is no cycle of G containing α. Since this applies to any edge in G, it follows that G has no cycle.

($\Leftarrow$) Suppose that G has no cycle. Then by the previous lemma, any edge is a bridge, so G is a tree. □

Theorem 5.4. *A graph G is a tree if and only if every pair of distinct vertices x and y is joined by a* unique *path. This path is necessarily a path of length $d(x, y)$, i.e., it is a shortest path.*

Proof. ($\Rightarrow$) Suppose G is a tree. Then G is connected, so for any x and y there is a path connecting the two vertices. Suppose that there is more than one path connecting x and y. Then there is some u, the first vertex in which the two paths begin to diverge; let v be first vertex in which the two diverging walks meet together again. Then there are two distinct paths from u to v, so concatenating these two paths gives us a cycle, which contradicts the fact that G is a tree. Thus there can only be one path from x to y.

($\Leftarrow$) Suppose that any two vertices are joined by a unique path, which implies that G is connected. Since there is only one path from any two points, there cannot be any cycle. (Otherwise, this will imply that these two points are joined by more than one paths, which is a contradiction.) Therefore G is a tree. □

We summarize all the equivalent characterizations of trees.

Theorem 5.5. *The following statements are equivalent.*

(i) G is a tree.
(ii) G is a connected graph with exactly $n - 1$ edges.
(iii) G is a connected graph that has no cycles.
(iv) G is a graph such that every pair of vertices x and y of G is joined by a unique path.

Now we look at some properties of trees.

Definition 5.2. A *pendent vertex* (or a *leaf*) of G is a vertex whose degree is equal to 1. The unique edge incident to a leaf is called a *pendent edge.*

Theorem 5.6. *Let G be a tree of order $n \geq 2$. Then G has at least two leaves.*

Proof. Since there are $n-1$ edges, if $(d_1, d_2, \ldots, d_n)$ is a degree sequence of the vertices, then

$$d_1 + d_2 + \cdots + d_n = 2(n - 1).$$

Suppose at most one vertex equals 1. Then the remaining $n-1$ d_i's has degree at least 2, so

$$d_1 + \cdots + d_n \geq 2(n - 1) + 1,$$

but this is a contradiction, so at least two of the d_i's must equal to 1. The claim follows. □

Definition 5.3. Let G be a graph. If H is a spanning subgraph of G that is also a tree, then H is a *spanning tree* of G.

Theorem 5.7. *Every connected graph has a spanning tree.*

Proof. Start with the set of entire edges, and remove all the edges that are not bridges. Once this procedure is completed, then every edge is a bridge, and the new subgraph remains connected. Thus this new graph is a spanning tree. □

The remainder of this section will explore some algorithms related to trees.

5.1. Dijkstra's algorithm for a distance tree for a vertex

Definition 5.4. Suppose $G = (V, E)$ is a graph, and for each edge $\alpha = \{x, y\}$ there is a *weight* associated with α, say $c(\alpha) = c\{x, y\}$. Then G is called a *weighted graph* with weight function $c : E \to \mathbb{R}_{\geq 0}$. If $\gamma : x_0 - x_1 - \cdots - x_k$ is a walk, then the *weight of a walk* γ is $c(\gamma) = c\{x_0, x_1\} + c\{x_1, x_2\} + \cdots + c\{x_{k-1}, x_k\}$. If γ is a walk from x_0 to x_k of the smallest weight, then $c(\gamma)$ becomes the *weighted distance* between x_0 and x_k denoted by $d_c(x_0, x_k)$. If $x_0 = x_k = x$, then $d_c(x, x) = 0$. If there is no walk from x to y, then $d_c(x, y) = \infty$.

Definition 5.5. Let $G = (V, E)$ be a graph, and u a vertex of G. Let H be a spanning tree rooted at u so that the weight of a walk between u and x is equal to $d_c(u, x)$ for any $x \in V$. Then H is called a *distance tree for u*.

Let $G = (V, E)$ be a weight graph of order n; let u be an arbitrary vertex of G. The following algorithm, called *Dijkstra's algorithm*, starts with a vertex, pick a vertex whose edge is of the smallest weight. After that, we keep track of the vertices that were accounted for, and keep adding edges of the minimum weight from one of the vertex covered by the algorithm already to another vertex not yet accounted for. Terminate the algorithm once no more edge can be added. See p444 of Brualdi (Theorem 11.7.4) for the proof that this algorithm gives us a distance tree for u.

(1) Begin with $U = \{u\}, D(u) = 0, F = \emptyset$, and $T = (U, F)$.

(2-i) Suppose that $x \in U$ and $y \notin U$. If there is no edge from x to y for any x and y, then the algorithm terminates.

(2-ii) Suppose there is at least one edge from a vertex in U to a vertex not in U. Pick an edge $\alpha = \{x, y\}$ where $x \in U$ and $y \notin U$ whose weight is the minimum weight. Once this edge is chosen, do the following:

- (a) Add y to U.
- (b) Add $\alpha = \{x, y\}$ into F.
- (c) Let $D(x) + c\{x, y\} = D(y) \to D(x)$, and then go back to (2-i) to determine if the algorithm must terminate. Otherwise, go to Step (2-ii).

5.2. **Algorithms for a minimum weight spanning tree**

In this section we present two algorithms that can be used to find a spanning tree of a graph whose sum of weights is the smallest.

Definition 5.6. For any subgraph H, we can define the *weight of a subgraph H of G* as the sum of weights of all the edges of H, which we denote $c(H)$. If H is a spanning tree such that

$$c(H) = \min\{c(J) : J \text{ a spanning tree of } G\},$$

then H is said to be a *minimum weight spanning tree*.

5.2.1. *Greedy algorithm*

The following greedy algorithm gives a minimum weight spanning tree. See p446-447 (Theorem 11.7.5) for the proof that this algorithm indeed yields a minimum weight spanning tree for $G = (V, E)$ of order n, a weighted connected graph with weight function c.

(1) Start with $F = \emptyset \subseteq E$.

(2) Let α be an edge not in F such that $F \cup \{\alpha\}$ does not contain any cycle. Out of those edges, let α be an edge of minimum weight. Add this chosen edge to F.

(3) Repeat the second step until $|F| = n - 1$; output $T = (V, F)$ upon termination of this algorithm.

5.2.2. *Prim's algorithm*

See Theorem 11.7.6 from Brualdi (p448-449) for the proof that Prim's algorithm outlined below yields a minimum weight spanning tree. Let $G = (V, E)$ be a weighted connected graph with weight function c, and let $u \in V$.

(1) Let $i = 1, U_1 = \{u\}, F_1 = \emptyset$, and $T_1 = (U_1, F_1)$.

(2) While $i \leq n - 1$, do the following steps:

   (a) Let $x \in U_i$ and $y \notin U_i$. Let $\alpha_i = \{x, y\}$ be an edge of smallest weight.

   (b) Let $U_{i+1} := U_i \cup \{y\}$ and $F_{i+1} := F_i \cup \{\alpha_i\}$. Let $T_{i+1} = (U_{i+1}, F_{i+1})$.

   (c) Increase i by 1.

(3) Once the previous step terminates, output $T_{n-1} = (U_{n-1}, F_{n-1})$. Note that necessarily $U_{n-1} = V$.

6. Chapter 12.1: Chromatic numbers

Definition 6.1. For any graph $G = (V, E)$, a *vertex colouring* of G is an assignment of a colour to each of the vertices of G so that adjacent vertices are assigned different colours. If k different colours were used to obtain a colouring, then such colouring is called a *k-vertex colouring*. If there is a k-colouring of G, then G is *k-colourable*. The smallest number k such that G is k-colourable is called the *chromatic number* of G, which is denoted by $\chi(G)$.

Theorem 6.1. *If G is a graph of order $n \geq 1$, then $1 \leq \chi(G) \leq n$ for any graph G. Furthermore, $\chi(G) = n$ if and only if $G = K_n$, and $\chi(G) = 1$ if and only if $G = N_n$. In particular, $\chi(G) = 2$ if and only if G is a bipartite graph.*

Proof. Clearly, we need at least one colour to get any vertex colouring of G, and we need at most n colours to colour all the vertices of G. Suppose that $\chi(G) = 1$. Since every vertex is of the same colour, this implies that no two vertices are adjacent to each other, so $G = N_n$.

Conversely, suppose that $\chi(G) \neq 1$. Thus $\chi(G) > 1$, so we need at least two colours to get a vertex colouring. This is possible only when there is at least one edge, say from x to y so that x and y are forced to get different colours. Thus $G \neq N_n$ as required.

Suppose that $G = K_n$. Then every pair of vertices is adjacent, so no two vertices can have the same colour. Hence $\chi(K_n) = n$. Suppose that $G \neq K_n$. Then there exist x and y such that x and y are not adjacent. This means we can assign x and y the same colour, so $\chi(G) \neq n - 1$.

Suppose that G is bipartite. Then there exist a bipartition X and Y so that no two vertices in the same bipartition is adjacent. Thus we can colour all the vertices in X with the first colour, and all the vertices in Y with the second colour. Hence $\chi(G) = 2$. Conversely, if $\chi(G) = 2$, then the set of vertices with the first colour and the set of vertices with the second colour form a bipartition, so G is bipartite. □

Corollary 6.1. *Suppose that G is a graph that contains an induced subgraph that is isomorphic to K_p. Then $\chi(G) \geq \chi(K_p) = p$. More generally, if H is a subgraph of G, then $\chi(G) \geq \chi(H)$.*

Corollary 6.2. *Let $G = (V, E)$ be a graph of order n, and suppose that q be the largest number satisfying the following condition: G contains an induced subgraph isomorphic to N_q. Then $\chi(G) \geq \lceil n/q \rceil$.*

Proof. Suppose $\chi(G) = k$. Partition the vertices into $V_1, \ldots, V_k$ so that any vertices in the same partition have the same colour. By assumption we have $|V_i| \leq q$ for any i, so we have

$$n = |V| = \sum_{i=1}^{k} |V_i| \leq qk.$$

Thus $\chi(G) = k \geq n/q$, and the claim follows upon noting that $\chi(G) \in \mathbb{N}$. □

In general, it is not straightforward to obtain a chromatic number of any graph G in general; in fact, there is no good known algorithm to obtain the chromatic number of an arbitrary graph G. However, there are some algorithms to obtain a vertex colouring which can help in estimating the chromatic number. We shall take a look at one of the algorithms called the greedy algorithm.

Algorithm 6.1 (Greedy algorithm for vertex colouring)**.** *Let $G = (V, E)$ with $V = \{x_1, \ldots, x_n\}$.*

(1) Assign the first colour to vertex x_1.

(2) For each $i = 2, \ldots, n$, define p to be the smallest colour such that none of the vertices $x_1, \ldots, x_{i-1}$ adjacent to x_i is coloured p. Assign the p-th colour to x_i.

In essence, in the i-th step, the greedy algorithm the following:

(1) Consider the vertices that already received a colour (i.e., $x_1, \ldots, x_{i-1}$).

(2) Only consider the vertices from that list that is adjacent to x_i.

(3) See how many colours are used amongst these vertices. Say $p-1$ colours were used to colour those vertices. Then colour x_i with the p-th colour.

Hence, if the maximum degree of a vertex is Δ, at most Δ vertices are adjacent to x_i for any i. If colour p is assigned to x_i, then $p \leq \Delta + 1$, since p hits the maximum if and only if each of the Δ vertices adjacent to x_i is coloured with a different colour. This guarantees that a $(\Delta+1)$-colouring always exists. Hence $\chi(G) \geq \Delta + 1$. This observation proves the following theorem.

Theorem 6.2. *Let G be a graph for which the maximum degree of a vertex is Δ. Then the greedy algorithm produces a $(\Delta+1)$-colouring of the vertices of G. Hence $\chi(G) \leq \Delta + 1$.*

It cannot be emphasized enough that the greedy algorithm *may* give $\chi(G)$, but not always. Also, the greedy algorithm itself does not tell us how far off a vertex colouring is from the actual chromatic number. However, in some cases, it is possible to get a tighter bound on $\chi(G)$. The following theorem indicates when this is possible, but the proof is beyond the scope of this notes.

Theorem 6.3 (Brooks)**.** *Let G be a connected graph for which the maximum degree of a vertex is Δ. Suppose also that G is neither a complete graph K_n nor a graph of odd cycle C_n. Then $\chi(G) \leq \Delta$.*

The discussion on vertex colouring with k colours raises another question – how many ways can the vertices of G be colours if we are given k colours? This prompts another definition.

Definition 6.2. Let $k \in \mathbb{N}$, and G a graph of order n. Suppose that $p_G(k)$ denotes the total number of available k-colourings of G. In fact, $p_G(k)$ is a polynomial function of k (see Brualdi for more information) that gives the number of available k-colourings of G. $p_G(k)$ is called the *chromatic polynomial* of the graph G.

In order to prove that $p_G(k)$ is indeed a polynomial function in k, we first note a useful observation in proving this fact. In fact, the following formula is also useful in obtaining the chromatic polynomial of a graph.

Theorem 6.4 (Edge deletion-contraction formula)**.** *Let G be a graph of order n, and x and y adjacent vertices of G. Let G_1 be the graph obtained by deleting the edge $\{x, y\}$. Let G_2 be the graph obtained by deleting the edge $\{x, y\}$, followed by contracting x and y into one vertex. Then*

$$p_G(k) = p_{G_1}(k) - p_{G_2}(k).$$

Proof. Observe that x and y in G_1 may be coloured with two different colours or with the same colour. In G, x and y must be coloured differently. In G_2, x and y (now contracted into one vertex) are necessarily coloured with the same colour. Therefore, we have

$$p_{G_1}(k) = p_G(k) + p_{G_2}(k),$$

so the claim follows. □

Theorem 6.5. *For any graph G, $p_G(k)$ is a polynomial function in k.*

Proof. See the informal discussion on p469 of Brualdi and the formal algorithm for computing the chromatic polynomial of a graph outlined in p470 of Brualdi. □

Now we compute the chromatic polynomial of some familiar graphs.

Example. Let $G = K_n$. Then the first vertex has k options to choose from; the second vertex can be coloured with anything except for the first colour, so there are $k-1$ options. Continuing this way, we see that

$$p_{K_n}(k) = k(k-1)(k-2)\cdots(k-(n-1)) = [k]_n.$$

Example. $p_{N_n}(k) = k^n$ since any vertex can be coloured with any colour.

Theorem 6.6. *Suppose T is a tree of order n. Then $p_T(k) = k(k-1)^{n-1}$.*

Proof. We prove by induction on the number of vertices. Suppose that the order of T is 1. Then $p_T(k) = k = k(k-1)^{1-1}$, so the claim follows. Now assume that the claim holds for all trees of order $n-1$. Any tree of order n can be obtained by adding a vertex to some tree of order $n-1$. The added nth vertex is necessarily a leaf, so there are $k-1$ options for this vertex (any colour except for the colour that the adjacent vertex received). Thus $p_T(k)$ is the product of $k-1$ and the chromatic polynomial for a tree of order $n-1$. Hence $p_T(k) = k(k-1)^{n-2}(k-1) = k(k-1)^{n-1}$ as required. □

Example. Let $G = C_5$ the cycle graph of order 5. We will use the edge deletion-contraction formula to compute $p_{C_5}(k)$. Note that any edge deletion gives us a simple tree order 5. On the other hand, if an edge is deleted and then the two vertices are fused, then we get the cycle graph of order 4. Suppose T_m is the simple tree of order m. Then we have

$$p_{C_5}(k) = p_{T_5}(k) - p_{C_4}(k).$$

Apply the edge deletion-contraction formula again to C_4 to get

$$\begin{aligned} p_{C_5}(k) &= p_{T_5}(k) - (p_{T_4}(k) - p_{C_3}(k)) = p_{T_5}(k) - p_{T_4}(k) + p_{K_3}(k) \\ &= k(k-1)^4 - k(k-1)^3 + k(k-1)(k-2) = k(k-1)(k-2)(k^2-2k+2). \end{aligned}$$

One efficient way to compute a chromatic polynomial is recognizing an induced subgraph isomorphic to K_r for some r. This establishes a lower bound on the chromatic number. Furthermore, computing the smallest k that makes the chromatic polynomial non-zero gives us the chromatic number. We first introduce a new definition and a lemma in order to state the next key result.

Definition 6.3. Let $G = (V, E)$ be a graph of order n that is not equal to K_n. If $U \subsetneq V$ is a subset such that G_{V-U} is disconnected, then U is said to be an *articulation set* of G.

Remark. If G is not complete, then there exists two non-adjacent vertices a and b, so $U = V \setminus \{a, b\}$ is an articulation set of G. Therefore, if G is complete, then no such a and b can be chosen, so G has no articulation set in this case.

Lemma 6.1. *Let G be a graph, and assume that G contains a subgraph H isomorphic to K_r. Then $[k]_r \,|\, p_G(k)$.*

Proof. Note that the vertices in H are all coloured differently, and we can extend the colouring for other remaining vertices not in H after the vertices in H are coloured. That is, we have $p_G(k) = [k]_r q(k)$, where $q(k)$ denotes the number of k-colourings for the vertices not in H. □

Now we state the key result.

Theorem 6.7. *Let U be an articulation set of G and suppose that the induced subgraph G_U is a complete graph K_r. Let the connected components of G_{V-U} be $G_{U_1}, \ldots, G_{U_t}$. For each $1 \leq i \leq t$, define $H_i := G_{U \cup U_i}$. Then*

$$p_G(k) = \frac{p_{H_1}(k) \times \cdots \times p_{H_t}(k)}{([k]_r)^{t-1}},$$

and $\chi(G) = \max\{\chi(H_1), \ldots, \chi(H_k)\}$.

Proof. Observe that H_i and H_j for any $i \neq j$ only have the vertices in U as common vertices. Then for each i, H_i has $p_{H_i}(k)$ colourings. Thus colouring each of H_i covers the all the vertices. However, note that for a colouring of G, the colouring of H_i must have the same colouring for all $1 \leq i \leq t$. Since U is coloured t times, it follows that $p_{H_1}(k) \times \cdots p_{H_t}(k)$ must be divided by $p_{K_r}(k)^{t-1}$ so that the vertices in U are only coloured once. □

7. Chapter 12.2: Plane and planar graphs

Recall that any general graph $G = (V, E)$ is planar if it is possible to draw G without any of the edges overlapping at non-vertex points. Suppose that r denotes the number of regions divided by the edges of G; let n be the order of G, and e the number of edges.

Theorem 7.1 (Euler's formula)**.** *Let G be a connected planar graph of order n with e edge-curves. Then $r - e + n = 2$.*

Proof. Suppose that G is a tree. Then $e = n - 1$ and $r = 1$ (note that G cannot divide the infinite region because otherwise this will imply that G has a cycle). Thus $1 - (n - 1) + n = 1 + 1 = 2$ as required. Now suppose that G is not a tree. Then G has a spanning tree T, and clearly we have $r' - e' + n' = 2$. Since T is a tree, adding any new edge-curve will create a cycle, thereby dividing an existing region into two parts. So every time a new edge is being added, r must increase by 1, and so must e while n remains invariant. Therefore $r - e + n = 2$ still holds even after all the remaining edges are added. □

Using Euler's formula, we can derive one important property of planar graphs.

Theorem 7.2. *Let G be a connected planar graph. Then there is a vertex whose degree is at most* 5.

Proof. Suppose G is a planar graph. The only way for two regions to be bordered by a single edge is to be bordered by loops, which G does not have. There are no multiple edges between any of the two vertices, to no region can be bordered by two edges. So we need at least three edges to divide a region into two parts. If $f_1, \ldots, f_r$ are the number of edges of each subdivided region, then indeed $f_1 + \cdots + f_r = 2e$. Since $f_i \geq 3$, we have $3r \leq 2e$, or $3(e - n + 2) = 3e - 3n + 6 \; leq 2e$. So $e \leq 3n - 6$. Now recall that if $(d_1, \ldots, d_n)$ is a degree sequence of G, then $d_1 + d_2 + \cdots + d_n = 2e$. It follows that

$$\frac{d_1 + \cdots + d_n}{n} = \frac{2e}{n} \leq \frac{2(3n-6)}{n} < 6,$$

so the average of the degrees of vertices is strictly less than 6, and this is possible only when there is a vertex of degree 5 or less. □

However, a converse of the above theorem is not true, as we will see in the next few examples.

Example. A complete graph K_n is planar if and only if $n \leq 4$. It is a straightforward verification to check that K_1, K_2, K_3, and K_4 are all planar. Now consider K_5. As shown in the previous theorem, if K_5 were to be a planar graph, we need to have $e \leq 3n - 6$. But note that $\binom{5}{2} = 10 = e \not\leq 3(5) - 6 = 9$, so K_5 cannot be planar. Since K_n for any $n \geq 5$ contains K_5 as its subgraph, none of K_n can be planar for any $n \geq 5$. Observe that K_5 consists of vertices whose degrees are all less than 5, but K_5 is not planar.

Example. A complete bipartite graph $K_{m,n}$ is planar if and only if $m \leq 2$ or $n \leq 2$. It is straightforward to draw a planar representation of $K_{m,n}$ if either one of them is at most 2. Now consider $K_{3,3}$. Any bipartite graph can only have even cycles; and since $K_{m,n}$ cannot have any multiple edges, any region can be created by a minimum of four edges rather than three. So in this case $4r = 4e - 4n + 8 \leq 2e$, so $e \leq 2n - 4$ if $K_{m,n}$ is to be planar. But note that $K_{3,3}$ has 9 edges, so $9 \not\leq 2(6) - 4 = 8$. Therefore since any $K_{m,n}$ with $m, n \geq 3$ contains $K_{3,3}$ as its subgraph, it follows that $K_{m,n}$ is not planar for any $m, n \geq 3$. Note that all vertices of $K_{3,3}$ has degree 3 but is not planar.

The above two graphs (K_5 and $K_{3,3}$)) will play an important role in characterizing all non-planar graphs as we will see shortly.

Definition 7.1. Let $G = (V, E)$ be any graph, and $\{x, y\} \in E$. Suppose z is a new vertex added on $\{x, y\}$ in order to obtain two new edges $\{x, z\}, \{z, y\}$. Let this new graph be G'. Adding such z is called *subdividing the edge* $\{x, y\}$, and G' is called a *subdivision* of a graph G.

Evidently, if G is not planar, then any of its subdivisions cannot be planar either. If H were a subdivision of G that is planar, then we can obtain a planar representation of G by deleting the vertices used to subdivide edges from a planar representation of H, contrary to G being non-planar. This insight implies that G is not planar it contains a subdivision of

K_5 or $K_{3,3}$. The converse is the more challenging direction, and Kuratowski proved that the converse holds as well.

Theorem 7.3 (Kuratowski). *A graph G is planar if and only if it does not have a subgraph which is a subdivision of K_5 or $K_{3,3}$.*

We present an alternative characterization of planarity which still involves K_5 and $K_{3,3}$.

Definition 7.2. For any graph G, a *contraction* of a graph G is a graph that can be obtained by successively contracting edges (i.e., fuse the two vertices of an edge into one point).

Theorem 7.4 (Wagner-Harary-Tutte). *A graph G is planar if and only if it does not contain a subgraph that contracts to K_5 or $K_{3,3}$.*

8. Chapter 12.4: Independence number and clique number

Definition 8.1. Let $G = (V, E)$, and $U \subseteq V$. Then U is an *independent set* if no two vertices from U are adjacent. Equivalently, U is an independent set if G_U is a null graph. Let I be a largest independent set of G. Then the *independence number* of G is the cardinality of I, and we denote $\alpha(G)$.

Remark. Clearly, any subset of an independent set is also an independent set.

Example. Let $V_1, \ldots, V_k$ be the colour partition of $G = (V, E)$. That is, any two vertices from the same V_i are of the same colour. Then $V_1, \ldots, V_k$ are independent sets each. $\chi(G)$ is precisely the smallest k such that V can be partitioned into k independent sets. In fact, we can re-state the inequality in Corollary 6.2 in the following way: $\chi(G) \geq \lceil n/\alpha(G) \rceil$.

Example. $\alpha(N_n) = n$ and $\alpha(K_n) = 1$. Also, $\alpha(K_{m,n}) = \max\{m, n\}$.

Definition 8.2. Suppose $G = (V, E)$ is a graph, and let $U \subseteq V$ a subset of V so that any vertex not in U is adjacent to some vertex in U. Then U is called a *dominating set* of G. Let D be a set of smallest size such that D is a dominating set. Then the *domination number*, denoted by $\mathrm{dom}(G)$, is the size of D.

Remark. If W is a dominating set, then any set containing W is also a dominating set.

Example. $\mathrm{dom}(N_n) = n$ and $\mathrm{dom}(K_n) = 1$. Also, $\mathrm{dom}(K_{m,n}) = 2$ if $m, n \geq 2$.

Like independent numbers, computing a dominating number of a graph is very difficult. However, if G is connected, we can obtain a simple inequality.

Theorem 8.1. *Let G be a connected graph of order $n \geq 2$. Then* $\mathrm{dom}(G) \leq \lfloor n/2 \rfloor$.

Proof. We prove by induction on n. Any connected graph has a spanning tree; let T be a spanning tree of G. Then $\mathrm{dom}(G) \leq \mathrm{dom}(T)$, so we can reduce this problem by considering trees of order $n \geq 2$ only. If $n = 2$, then T can only be K_2, so $\mathrm{dom}(T) = 1 = \lfloor 2/2 \rfloor$. Now suppose that $n \geq 3$. Suppose that y is a vertex which is adjacent to a leaf x of T. If T^* is a subgraph of T obtained by removing y, then T^* is a forest, at least one of which is a tree of order 1 (since x is a leaf). Let $T_1, \ldots, T_k$ be the connected components of T^* such that every

T_i is of order at least 2 (say that n_i is the order of T_i). Then $n_1 + \cdots + n_k \geq n - 2$, so each T_i has a dominating set by the induction hypothesis, and any dominating set is of size at most $\lfloor n_i/2 \rfloor$. Therefore, the union of these dominating sets and $\{y\}$ is indeed a dominating set of T. The claim follows upon noting that

$$1 + \left\lfloor \frac{n_1}{2} \right\rfloor + \cdots + \left\lfloor \frac{n_k}{2} \right\rfloor \leq 1 + \left\lfloor \frac{n_1 + \cdots + n_k}{2} \right\rfloor \leq 1 + \left\lfloor \frac{n-2}{2} \right\rfloor = \left\lfloor \frac{n}{2} \right\rfloor. \qquad \square$$

Definition 8.3. A *clique* in a graph G is a subset U of vertices whose subgraph induced by U is a complete graph. The size of a clique of largest size is the *clique number* of G, which we denote $\omega(G)$.

Definition 8.4. For any graph G, the *complementary graph* of G is $\overline{G} = (V, \overline{E})$. That is, $\overline{G}$ consists of the same vertices V, and two vertices are adjacent in $\overline{G}$ if and only if the two vertices are not adjacent in G.

Remark. If U is a clique in a graph G, then U is an independence set in $\overline{G}$, the complement graph of G. Similarly, if J is an independent set in a graph G, then J is a clique in $\overline{G}$. Hence, it follows that $\alpha(G) = \omega(\overline{G})$ and $\omega(G) = \alpha(\overline{G})$. Clearly, $\chi(G) \geq \omega(G)$ since $\chi(G) \geq p$ if G contains a subgraph isomorphic to K_p. Let G be a bipartite graph with at least one edge. Then $\chi(G) = \omega(G) = 2$, so this is one example where the chromatic number and the clique number match. However, it is possible to find an example where the strict inequality holds. Let $G = C_n$ for some odd n. Then $\chi(C_n) = 3$ but $\omega(C_n) = 2$.

Example. $\omega(N_n) = 1, \omega(K_n) = n$, and $\omega(K_{m,n}) = 2$.

Observe that there is a complementary relationship between the clique number and the chromatic number. Indeed, the chromatic number denotes the smallest number of partitions $V_1, \ldots, V_k$ such that two vertices in the same V_i are not adjacent. We can also partition the vertices so that any two vertices in the same partitions are adjacent also (say $W_1, \ldots, W_r$, each of which is a clique in G).

Definition 8.5. The partition $W_1, \ldots, W_r$ is called a *clique-partition* of a graph G. The smallest such r is said to be the *clique-partition number* of G, and we denote it by $\theta(G)$.

Remark. As with the independence number and the clique number, there is a complementary relationship between the clique-partition number and the chromatic number. Particularly, we have $\chi(G) = \theta(\overline{G})$ and $\theta(G) = \chi(\overline{G})$. As with the chromatic number and the clique number, one can derive the following inequality: $\theta(G) \geq \alpha(G)$. Why is this true? This follows from the observation that no two non-adjacent vertices can be in the same clique. Thus, each of the vertices in an independence set of size $\alpha(G)$ must be in a separate clique, so there must be at least $\alpha(G)$ clique-partitions for G.

It is natural to wonder for which G we have $\chi(G) = \omega(G)$ and/or $\theta(G) = \alpha(G)$. Suppose that G is a graph with two connected components consisting of H and K_p such that $\omega(H) \leq p$. Then $\chi(G) = p$ and $\omega(G) = p$. In fact it is possible to impose the equality for any induced subgraph of G. We can do the same thing with $\theta(G) = \alpha(G)$, which prompts the following definitions.

Definition 8.6. A graph G is *χ-perfect* if $\chi(H) = \omega(H)$ for any subgraph H of G. A graph G is *θ-perfect* if $\theta(H) = \alpha(H)$ for any subgraph H of G. If G is both χ-perfect and θ-perfect, then G is a *perfect graph*.

In fact, being χ-perfect implies being θ-perfect and vice versa; therefore, there is only one kind of perfection, so we can just talk about perfect graphs without any prefix before the word perfect.

Theorem 8.2. *A graph G is χ-perfect if and only if G is θ-perfect. Equivalently, G is χ-perfect if and only if $\overline{G}$ is χ-perfect. Thus, G is perfect if and only if $\overline{G}$ is perfect.*

9. Chapter 12.5: Matching number

Definition 9.1. Let $G = (V, E)$ be a graph. Then $M \subseteq E$ is a *matching* if no two edges in M have any vertex in common. We say the matching M *meets a vertex x* if there is an edge in M with one of the endpoints x. Furthermore, if M meets every vertex in x, then M is called a *perfect matching*. The *matching number* of a graph G is the largest number of edges in a matching in G; we denote it by $\rho(G)$.

Remark. Any edges contain two vertices, so if G has n vertices then a matching M necessarily have at most $n/2$ edges. Also any graph with a perfect matching necessarily is of even order.

We can think of a matching problem as an SDR problem. Suppose that $\mathcal{A} = (A_1, \ldots, A_n)$ is a collection of subsets of $Y = \{y_1, \ldots, y_m\}$. Let each A_i be represented by a vertex x_i; let each elements y_j be represented by a vertex y_j, with x_i and y_j being adjacent if $y_j \in A_i$. Then the graph G generated by this relation is a bipartite graph. Notice that we can also create a collection of subsets based on a given matching. Thus a matching problem for a bipartite graph and an SDR problem are two equivalent ways of looking at the same problem. Therefore, if G has a perfect matching, then its corresponding $\mathcal{A}$ has an SDR.

Theorem 9.1. *Let $G = (V, E)$ be a bipartite graph with bipartition X, Y with associated family $\mathcal{A}_G$ of subsets of Y. If t is a positive integer such that a subfamily $(A_{i_1}, \ldots, A_{i_t})$ has an SDR $(e_{i_1}, \ldots, e_{i_t})$, then there is a matching $M = \{(x_{i_1}, e_{i_1}), \ldots, (x_{i_t}, e_{i_t})\}$.*

Conversely, from a matching $M = \{(x_{i_k}, e_{i_k}) : 1 \leq i \leq t\}$, one can obtain a subfamily of $\mathcal{A}_G$ with an SDR. Therefore, G has a perfect matching if and only if $\mathcal{A}_G$ has an SDR.

Definition 9.2. If $G = (V, E)$ is a graph, then $W \subseteq V$ is a *cover* of G if every edge of G contains a vertex from W. The smallest number of vertices that form a cover of G is denoted by $c(G)$.

Clearly, if W is a cover, then its complement $V \setminus W$ must be an independent set, since otherwise there will be an edge connecting two vertices from $V \setminus W$, which means that W is not a cover. Conversely, if U is an independent subset of V, then every edge from U must be connected to a vertex in $V \setminus U$, making $V \setminus U$ a cover of G. Thus we can conclude that W is a cover of G if and only if its complement is an independent set.

Theorem 9.2 (König-Egerváry theorem)**.** *If $G = (V, E)$ is a bipartite graph, then the largest number of edges in a matching equals the smallest number of vertices in a cover, i.e., $\rho(G) = c(G)$.*

The bipartite assumption in König-Egerváry is crucial since it is possible to find a non-bipartite graph such that $c(G) < \rho(G)$ – for instance, consider K_n for any $n > 3$. Then $\rho(G) = \lceil n/2 \rceil$ whereas $c(G) = n - 1$. However for non-bipartite graphs it is possible to express $\rho(G)$ in terms of other numbers of graph-theoretic significance. To state further results on $\rho(G)$, we introduce a few additional notions.

Definition 9.3. Let G be a disconnected graph. Then an *odd component* of G is a connected component of G consisting of an odd number of vertices. The number of odd components of G is denoted by $\mathrm{oc}(G)$.

Theorem 9.3. *A graph $G = (V, E)$ has a perfect matching if and only if $\mathrm{oc}(G_{V \setminus U}) \leq |U|$ for every $U \subseteq V$. Therefore, G has a perfect matching if and only if removing a set of vertices does not create more odd components than the number of vertices removed.*

Theorem 9.4 (Berge-Tutte formla)**.** *Let $G = (V, E)$ be a graph with n vertices. Then*

$$\rho(G) = \min\{n - (\mathrm{oc}(G_{V \setminus U}) - |U|)\}$$

where the minimum is taken over all $U \subseteq V$.

10. Chapter 12.6: Connectivity

Every graph is either connected or disconnected. However, one can discuss how much or how well a graph is connected. Intuitively speaking, for any graph of order n, K_n is the most connected graph compared to any other graphs of order n. Indeed, more edges need to be removed from K_n than any other graphs of order n, and more vertices need to be removed from K_n to render the new graph disconnected.

Definition 10.1. Let $G = (V, E)$ be a graph of order $n \geq 2$. Then the *vertex-connectivity* of G is

$$\kappa(G) = \min\{|U| : G_{V \setminus U} \text{ is disconnected}\},$$

i.e., the smallest number of vertices whose removal renders the new graph disconnected. Equivalently, $\kappa(G)$ is the smallest size of an articulation set. The *edge-connectivity* of G is

$$\lambda(G) = \min\{|F| : G' = (V, E \setminus F) \text{ is disconnected}\},$$

i.e., the smallest number of edges whose removal renders the new graph disconnected.

Remark. It follows from the definition that $\kappa(G) = \lambda(G) = 0$ if G is disconnected. Note that $\lambda(G) = 1$ if and only if G has a bridge.

Example. $\kappa(K_n) = \lambda(K_n) = n - 1$; $\kappa(N_n) = \lambda(N_n) = 0$.

Theorem 10.1. *If G is a graph of order $n \geq 2$, then*

$$0 \leq \kappa(G) \leq n-1,$$

with equality on the left if and only if G is disconnected, and with equality on the right if and only if G is a complete graph.

Proof. The equality is evident, so it suffices to show the remaining inequality. Suppose G is a non-complete connected graph. Then there are two vertices a and b that are non-adjacent. So removing all the vertices except for a and b will definitely result in a disconnected graph, so $\kappa(G) \leq n-2$ for any non-complete connected graph of order n. □

Theorem 10.2. *For any graph G, we have*

$$\kappa(G) \leq \lambda(G) \leq \delta(G),$$

where $\delta(G)$ is the smallest degree of a vertex of G.

Proof. The second equality follows upon noting that if all the edges adjacent to any vertex of degree $\delta(G)$ (say x) are removed, then the new graph is disconnected (there will be no walk from x to any other vertices). Thus $\lambda(G) \leq \delta(G)$.

Suppose $G = K_n$. Then $\kappa(G) = \lambda(G) = n-1$. If G is disconnected, then $\kappa(G) = \lambda(G) = 0$. Thus we may assume that G is a connected non-complete graph. Suppose F is a subset of edges if size $\lambda(G)$ whose removal results in a disconnected graph (say H); let U be a subset of vertices whose removal results in a disconnected graph. Since G is connected, and F is of smallest size to render the new graph disconnected, it follows that H has exactly two connected components; suppose that V_1 and V_2 form two connected components. Then $|V_1| + |V_2| = n$. Suppose that F consists of every possible edge joining vertices in V_1 to vertices in V_2. But this means $|F| \geq n-1$, or equivalently $\lambda(G) \geq n-1$. This forces $\lambda(G) = n-1$, which is impossible to happen since $G \neq K_n$. Hence there is $a \in V_1$ and $b \in V_2$ that are not adjacent. With this observation, we will try to construct U so that $G_{V \setminus U}$ is disconnected. Suppose that $\alpha \in F$ is an edge with a being one of the vertices of α. If this is the case, then the other vertex must be in V_2 (since $\alpha \in F$); add this vertex in V_2 to U. If this is not the case, then we add the vertex of α that is in V_1. Each operation adds at most one new vertex, and there are $|F|$ edges, so it follows that $|U| \leq |F|$. Note that removing all the vertices in U renders the graph disconnected since there will be no path from a to b. In conclusion,

$$\kappa(G) \leq |U| \leq |F| = \lambda(G),$$

as required. □

We can formulate connectedness using the connectivities we introduced so far. For instance, G is connected if and only if its vertex-connectivity satisfies $\kappa(G) \geq 1$.

Definition 10.2. Let G be a graph such that $\kappa(G) \geq k$. Then G is said to be *k-connected*. If $\{v\}$ is an articulation set of G, then v is said to be an *articulation vertex* of a graph G.

Clearly if G is k-connected, then G is also m-connected for any $k > m$ since $\kappa(G) \geq k > m$.

Theorem 10.3. *Let G be a graph of order $n \geq 3$. Then the following are equivalent.*

(1) G is 2-connected.

(2) G is connected and does not have an articulation vertex.

(3) For each triple of vertices a, b, c, there is a path joining a and b that does not contain c.

Proof. ((1) $\Rightarrow$ (2)) If $\kappa(G) \geq 2$, then at least two vertices must be removed to make the new graph disconnected. Therefore G cannot have any articulation vertex. Also, being 2-connected implies being 1-connected, so necessarily G is connected.

((2) $\Rightarrow$ (3)) Suppose a, b, c are three vertices. Since G is connected, there is a path from a to b. Note that removing c will keep the graph connected since G has no articulation vertex, meaning one can find a path from a to b not containing c.

((3) $\Rightarrow$ (1)) For any triple of vertices, we see that there is a path from a to b, so G is connected. If c is an articulation vertex, then removing c results in a disconnected graph. Therefore removing c results in no path from a so b, which contradicts our assumption. Hence G cannot have any articulation vertex. If G has no articulation vertex, then we need to remove more than one vertex to disconnect G. Therefore $\kappa(G) > 1$, so G is 2-connected as required. □

Definition 10.3. Fo any connected graph G, a *block* of G is a maximal induced subgraph of G that is connected and has no articulation vertex.

Suppose that U is a subset of vertices such that G_U is a block. If $U \subsetneq W \subseteq V$, then G_W is either disconnected or has an articulation vertex. Thus if $|U| \geq 3$, then G_U is 2-connected. If $|U| = 2$, then $G_U = K_2$. We conclude this section with a theorem that are useful in determining the blocks of a graph, and a theorem that gives an alternative characterization of 2-connectedness.

Theorem 10.4. *Let $G = (V, E)$ be a connected graph of order $n \geq 2$, and let*

$$G_{U_i} = (U_i, E_i) \text{ for all } 1 \leq i \leq r$$

so that each G_{U_i} is a block of G. Then $E_1, \ldots, E_r$ form a partition of the set E of edges of G, so each edge of G belong to exactly one block. Furthermore, each pair of blocks has at most one vertex in common.

Theorem 10.5. *Let $G = (V, E)$ be a graph of order $n \geq 3$. Then G is 2-connected if and only if, for each pair a, b of distinct vertices, there is a cycle containing both a and b.*

Corollary 10.1. *Let G be a graph of order $n \geq 3$. Then G is 2-connected if and only if for each pair a, b of distinct vertices, there are two paths joining a and b whose only common vertices are a and b.*

Theorem 10.6 (Menger's theorem, special case)**.** *Let k be a positive integer, and let G be a graph of order $n \geq k+1$. Then G is k-connected if and only if, for each pair a, b of distinct vertices, there are k paths joining a and b such that each pair of paths has only the vertices a and b in common.*

11. Chapter 8.1: Catalan numbers

Definition 11.1. The *Catalan sequence*, named after Eugéne Catalan, is the sequence $C_0, C_1, \ldots, C_n, \ldots$ of the form

$$C_n = \frac{1}{n+1}\binom{2n}{n},$$

and C_n is called the *n-th Catalan number.*

Interestingly, Catalan numbers show up in combinatorial contexts, as well will see.

Theorem 11.1. *The number of sequences $a_1, \ldots, a_{2n}$ of $2n$ terms that can be formed by using exactly n $+1$'s and exactly n -1's whose partial sums are always non-negative (i.e., $a_1 + a_2 + \cdots + a_k \geq 0$ for any $1 \leq k \leq 2n$) equals the n-th Catalan number*

$$C_n = \frac{1}{n+1}\binom{2n}{n}.$$

Proof. Let S be the set of sequences of length $2n$ consisting of exactly n $+1$'s and exactly n -1's. Then we partition S into two sequences – let A_n be the set of sequences whose partial sums are always non-negative, and let U_n be the set of sequences whose partial sum is negative for some $1 \leq k \leq 2n$. So $|A_n| + |U_n| = |S| = \binom{2n}{n}$. Suppose that the sequence $a_1, \ldots, a_{2n}$ is in U_n. Then there exists first k such that $a_1 + \cdots + a_k < 0$. Since each term is ± 1, it follows that necessarily $a_k = -1$ and $a_1 + a_2 + \cdots + a_{k-1} = 0$. Thus amongst the remaining $2n - k$ terms, there are more 1's than -1's.

Now consider a new sequence, with the sign of the first k terms reversed (say $a'_1, \ldots, a'_{2n}$). Then $a'_i = -a_i$ for any $1 \leq i \leq k$ and $a'_i = a_i$ otherwise. Suppose the set of such new sequences is V_n. Observe that there is a one-to-one correspondence between the set of sequences in U_n and V_n, so $|U_n| = |V_n|$. Any sequence is V_n has one more 1 than -1 in the first k terms; as observed previously, there are one more 1's than -1's amongst the remaining $2n - k$ terms. Hence any sequence in V_n has $n+1$ 1's and $n-1$ -1's. Therefore,

$$|U_n| = |C_n| = \frac{(2n)!}{(n+1)!(n-1)!},$$

so we have

$$\begin{aligned}
|A_n| = \binom{2n}{n} - |U_n| &= \frac{(2n)!}{n!n!} - \frac{(2n)!}{(n+1)!(n-1)!} \\
&= (2n)!\left(\frac{1}{n!n!} - \frac{1}{(n+1)!(n-1)!}\right) \\
&= (2n)!\left(\frac{n+1}{(n+1)!n!} - \frac{n}{(n+1)!n!}\right) = \frac{(2n)!}{n!(n+1)!} \\
&= \frac{1}{n+1}\binom{2n}{n} = C_n.
\end{aligned}$$

□

Therefore, the theorem above comes in handy when a combinatorial problem can be reduced into counting the number of sequences of length $2n$ consisting of the equal number of $+1$'s and -1's. Some examples are presented in Brualdi Chapter 8, Exercises #1, #2, and #36.

We conclude this section with the linear recurrence formula for C_n.

Definition 11.2. $C_n^* := n!C_{n-1}$ is called the *n-th pseudo-Catalan number.*

Observe that, from the definition of C_n, we have

$$C_n = \frac{4n-2}{n+1}C_{n-1},$$

so we have $C_n^* = (4n-6)C_{n-1}^*$, with $C_1^* = 1$. With these facts in mind, consider the following combinatorial problem. Let $a_1, \ldots, a_n$ be n numbers, and we want to find the number of multiplication schemes i.e., the number of ways to carry out the multiplication of the n numbers. We need to multiply two numbers $n-1$ times in total, each of which is either one of the a_i's or some partial product of the a_i's thereof. Let h_k be the number of multiplication schemes for k numbers. We may express each multiplication scheme with the multiplication symbol and the parentheses to indicate the order of multiplications.

So for example, we have $h_1 = 1, h_2 = 2, h_3 = 12$, and $h_4 = 120$. In fact, we can look at this inductively.

(1) Pick a multiplication scheme for $a_1, \ldots, a_{n-1}$, and we will compute the number of multiplication schemes for $a_1, \ldots, a_n$ where a_n is not the last number to be multiplied. Then a_n belongs to one of the parentheses already present. There are $n-2$ multiplication operations in any of the schemes we picked, so there are $n-2$ pairs of parentheses. Now with a_n in one of the parentheses, note that we can add a_n in four possible ways: to the left of the first term, to the right side of the first term (and add parentheses to group the first term and a_n), to the left side of the second term, or to the right side of the second term (and add parentheses to group the second term and a_n). Each scheme gives $2 \cdot 2 \cdot (n-2) = 4n-8$ schemes.

(2) Pick a multiplication scheme for $a_1, \ldots, a_{n-1}$ where a_n is the final number to be multiplied. That is, a_n is placed either to the left of the chosen multiplication scheme with $n-1$ numbers or to the right. Thus every scheme gives us two schemes.

Hence, we see that $h_n = (4n-6)h_{n-1}$. But note that $h_1 = C_1^*$, and that h_n satisfies the identical recurrence relation satisfied by C_n^*, from which we have $C_n^* = h_n$.

Now, we will add one more additional restriction: we are only interested in multiplication schemes such that $a_1, \ldots, a_n$ appear in this particular order if read from left to right. If h_n' is the number of multiplication schemes with this restriction, we note that $h_n' = h_n/n!$. Indeed, the total number of multiplication schemes with a particular arrangement of parentheses is $n!$, and we are only interested in one of them, so

$$h_n' = \frac{h_n}{n!} = \frac{C_n^*}{n!} = C_{n-1}.$$

Finally, observe that the total number of multiplication schemes of n numbers with that desired restriction is the sum of $h'_k h'_{n-k}$, and that $h'_1 = h_1/1! = 1 = C_0$, so

$$C_{n-1} = h'_n = h'_1 h'_{n-1} + h'_2 h'_{n-2} + \cdots + h'_{n-1} h'_1 = C_0 C_{n-2} + C_1 C_{n-3} + \cdots + C_{n-2} C_0.$$

Therefore, we have the following theorem.

Theorem 11.2. *For any $n \geq 1$, the n-th Catalan number C_n satisfies the following linear recurrence relation:*

$$C_n = C_0 C_{n-1} + C_1 C_{n-2} + \cdots + C_{n-1} C_0 = \sum_{k=0}^{n-1} C_k C_{n-1-k}.$$

Department of Mathematics and Statistics, Dalhousie University, 6316 Coburg Rd, Halifax, NS, Canada B3H 4R2

E-mail address: hsyang@dal.ca