Connectivity

CO 342: Graph Theory

David Duan, 2019 Fall (Prof. Peter Nelson)

Contents

1 Introduction

1.1 Basic Definitions

Def. 1.1.1 A graph $G = (V, E, i)$ is a 3-tuple where

- V is a finite set of *vertices*,
- E is a finite set of *edges* with $V \cap E = \varnothing$,
- $i : V \times E \to \{0, 1, 2\}$ is an *incidence function* such that

$$\forall e \in E : \sum_{v \in V} i(v, e) = 2.$$

Remark. Intuitively, $i(v, e)$ counts the number of times e is incident to v.

- For edge $e = ab$ with $a \neq b$, $i(a, e) = i(b, e) = 1$ and $i(v, e) = 0$ for all other $v \in V$.
- If e is a loop on a, then $i(a, e) = 2$ and $i(v, e) = 0$ for all other $v \in V$.

Def. 1.1.2 Recall the following basic definitions from Math23/49:

- Vertices $u, v \in V$ are *adjacent* if for some $e \in E$,
  - $i(u, e) = i(v, e) = 1$ where $u \neq v$, or
  - $i(u, e) = 2$ where $u = v$, i.e., e is a loop.
- A vertex $u \in V$ and an edge $e \in E$ are *incident* in $G = (V, E, i)$ if $i(u, e) \neq 0$.
- The *degree* of a vertex $v \in V$ is $\deg(v) = \sum_{e \in E} i(v, e)$.
- The *ends* of an edge $e \in E$ are $u, v \in V$ such that $i(u, e) > 0 \land i(v, e) > 0$.

Remark. To see why we define a graph like this, consider the *planar dual* H of a graph G.

If $G = (V, E, i)$ is the primal planar graph with a fixed planar embedding, then $H = (F, E, i')$ is the dual planar graph with F being the faces of the embedding of G and i' the incidence function determined by adjacent faces.

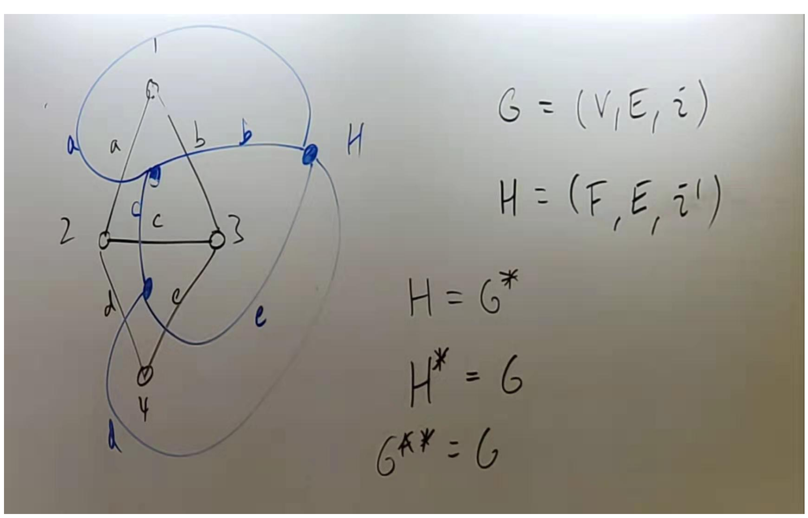

We see that $H = (F, E, i')$ and $G = (V, E, i)$ have different vertex sets and different incidence functions but share the same edge set. Our Def. of a graph makes it easier to work with the dual graph.

In this course, unless otherwise specified, we usually deal with simple graphs.

Def. 1.1.3 A simple graph is a graph with no parallel edges or loops.

1.2 Connectedness

Def. 1.2.1

- A graph G is *connected* if $V(G) \neq \varnothing$ and there is a *walk* from u to v for any $u, v \in V$.
- Two vertices $u, v \in V$ are *connected* if there is a *walk* from u to v.
- In other words, G is connected iff $V(G) \neq \varnothing$ and every pair of vertices is connected in G.

Remark.

- We use *walks* instead of *paths* here because joining two walks always produces a walk but joining two paths doesn't necessary give you a path.
- We want $V(G) \neq \varnothing$ because of **Prop. 1.5.3**; allowing the empty graph to be connected violates this Def..
    - Analogy: Every integer can be uniquely expressed as the product of prime numbers, and 1 is not considered as a prime number.

Prop. 1.3.2 Connectedness is an *equivalence* relation: reflexive, symmetric, transitive.

Proof. By intuition. □

1.3 Subgraph and Induced Subgraph

Def. 1.3.1 A *subgraph* $G = (V, E, i)$ is a graph $H = (V', E', i')$ where $V' \subseteq V, E' \subseteq E'$ and i' is the restriction of i to the domain $V' \times E'$.

Def. 1.3.2 If $X \subseteq V$, then the subgraph $G[X]$ ***induced*** by X is the subgraph (X, E', i') where E' consists of all edges with both ends in X.

Remark. Informally, a subgraph of G is obtained by removing edges and/or vertices arbitrarily; an induced subgraph is obtained by just removing vertices, i.e., an edge in G must also be an induced subgraph of G, given both of its vertices exist in the induced subgraph.

Prop. 1.3.3 Let G be a connected graph. Then there is a sequence $G_1, G_2, \ldots, G_n$ of connected graphs so that $G_n = G$ and, for each $i \in \{1, \ldots, n-1\}$, the graph G_i has i vertices and is an induced subgraph of G_{i+1}.

Proof. Let $k \geq 1$ be maximal so that there exist distinct vertices $v_1, \ldots, v_k$ of G for which $V' := \{v_1, \ldots, v_k\}$ induces a connected subgraph of G for each $i \in \{1, \ldots, k\}$. If $k = n$, we are done. If V' induces a component of G, then G is disconnected. Thus, there is some edge e with an end in V' and an end v outside of V'. It follows that the subgraph H induced by $V' \cup \{v\}$ is connected, since v is connected in H to some v_i and v_i's are pairwise connected in H. Setting $v_{k+1} = v$ gives a contradiction to the maximality of k. $\square$

1.4 Component

Def. 1.4.1 A *component* of G is an induced subgraph of the form $G[X]$ where X is an equivalent class under connectedness.

Intuitively, the following Prop. tells us that components are maximal subgraphs.

Prop. 1.5.2 A graph H is a component of G iff H is a maximal connected subgraph of G, that is, H is a connected subgraph of G, and there is no connected subgraph H' of G such that H is a subgraph of H' and $H \neq H'$.

Proof.

$\Longrightarrow$: Let G be a graph and H be a component of G. Suppose for a contradiction that exists a connected subgraph H' of G such that H is a subgraph of H' and $H \neq H'$. Since H is induced, $V(H) = V(H') \implies E(H) = E(H') \implies H = H'$, so we must have $V(H) \neq V(H')$, i.e., there exists some $v \in V(H') \neq V(H)$. Since v is not in the same equivalence class as the vertices in $V(H)$, it follows that v is not connected in G or in H' to any vertex in H'. Thus, H' is disconnected, a contradiction. It follows that H is a maximal connected subgraph.

$\Longleftarrow$: Let H be a maximal connected subgraph of G. If there exists an edge $e \in E(G)$ with ends in H but $e \notin E(H)$, then the subgraph obtained from H by adding this edge e is also connected, contradicting the maximality of H. Therefore, H is an induced subgraph of G. By connectedness, the vertices in H are pairwise connected in G. If there is some $w \in V(G) \setminus V(H)$ that is connected in G to a vertex in H, then let P be a wv-path where $v \in V(H)$. Clearly, $H \cup P$ is connected and not equal to H, which again contradicts the maximality of H. Therefore, $V(H)$ is an equivalent class under connectedness and H is a component by definition. $\square$

1.5 Union and Direct Sum of Graphs

Def. 1.5.1 Let $G_1 = (V_1, E_1, i_i)$ and $G_2 = (V_2, E_2, i_2)$.

Suppose that the subgraph obtained from G_1 by restricting to $V_1 \cap V_2$ and $E_1 \cap E_2$ is the same as the subgraph obtained from G_2 by restricting to $V_1 \cap V_2$ and $E_1 \cap E_2$, (i.e., G_1 and G_2 "agree" on their common vertices and edges), then the *union* $G_1 \cup G_2$ is defined to be the graph with vertex set $V_1 \cup V_2$, edge set $E_1 \cup E_2$, in which a vertex v is incident to an edge e iff e and v are incident in either G_1 or G_2.

When $V_1 \cup E_1$ and $V_2 \cup E_2$ are disjoint, the union $G_1 \cup G_2$ is called the *direct sum* of G_1 and G_2 and is written $G_1 \oplus G_2$.

Example 1.5.2 [Union]

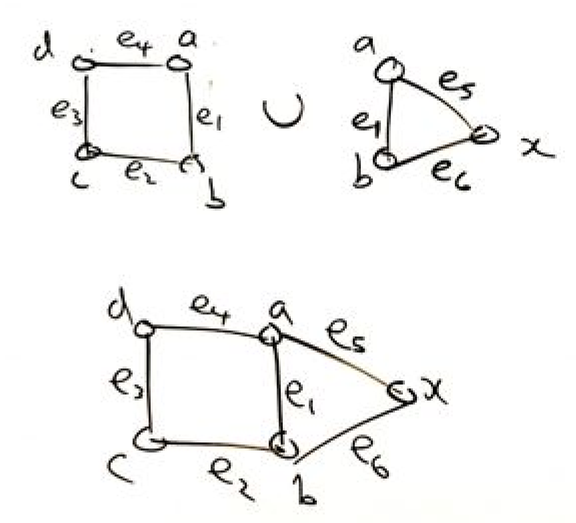

Prop. 1.5.3 Every graph is uniquely the direct sum of connected (sub)graphs.

Proof. By intuition. □

1.6 Subtraction (Removing Edges or Vertices)

Def. 1.6.1 For a set $X \subseteq V \cup E$ and a graph $G = (V, E, i)$, write $G - X$ for the subgraph of G with vertex set $V \setminus X$ and edge set $E \setminus X'$, where X' is the set of edges that are either in X or incident with a vertex in X.

Example 1.6.2 [Subtraction] Observe when we remove an vertex, we must also remove all the edges that are incident to it.

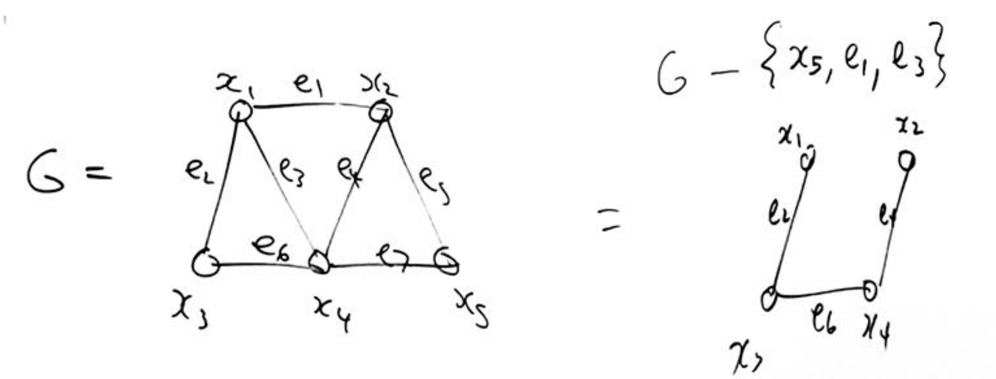

1.7 Other Definitions

Def. 1.7.1 (It is easier to consider paths and circuits as graphs rather than some "part" of a graph satisfying certain properties when talking with connectedness later.)

- A *path* is a graph when edges and vertices form a path.
- The *ends* of a path are its degree-1 vertices (or its only vertex if the path has no edges).
- A *circuit* is a graph when edges and vertices form a circuit.

Def. 1.7.2 (We extend the definition of a path to sets of vertices.)

- Given disjoint sets of vertices A, B in a graph G, an A, B-*path* is a path with one end in A, the other end in B, and all its other vertices in $V(G) \setminus A \cup B$.
- Define an a, B-path or a, b-path similarly where a, b are single vertices.

Def. 1.7.3 (Cut edge, cut vertex, and separator.)

- A set $X \subseteq V \cup E$ *separates* A and B in G if there is no A, B-path in $G - X$.
- e is a *cut edge* or *bridge* if there are vertices u, v of G that are not separated by $\varnothing$ but are separated $\{e\}$.
- A *cut vertex* of G is a vertex v such that there is some pair a, b not separated by $\varnothing$ but separated by $\{v\}$.

1.8 k-Connectedness

Def. 1.8.1 Let $k \geq 1$, a graph G is k-*connected* if $|V(G)| > k$ and there does not exist a set $X \subseteq V(G)$ with $|X| < k$ such that $G - X$ is disconnected.

Remark. Intuitively, a graph is k-connected when we cannot remove less than k vertices to make the graph disconnected. Note we also need the graph to have enough vertices since saying a graph with 2 vertices is 3-connected is meaningless.

- A graph is 1-connected when it is connected, except when $|V(G)| = 1$ (it is still a connected graph) as it violates the size constraint.
- A graph is 2-connected when it has no cut vertex, except when $G = *$ - $*$ (it does not have a cut vertex) as it violates the size constraint.

Example 1.8.2

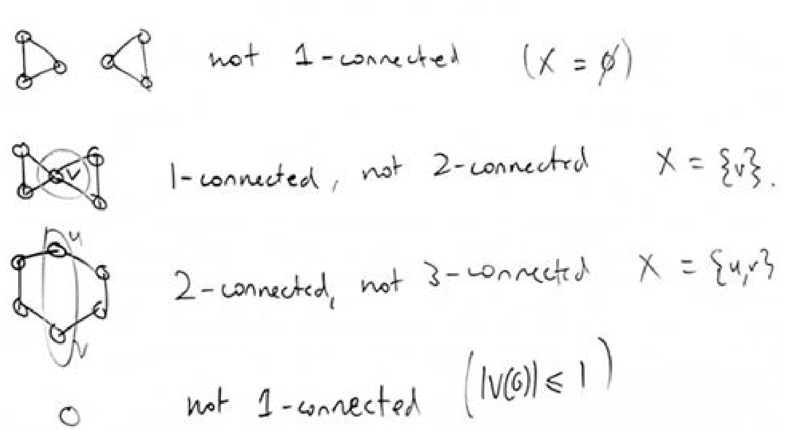

Def. 1.8.3 Let $k \geq 1$, a graph G is *k-edge-connected* if there does not exist a set $X \subseteq E(G)$ with $|X| < k$ such that $G - X$ is disconnected.

Remark. Intuitively, a graph is k-edge-connected when we cannot remove less than k edges to make the graph disconnected.

Prop. 1.8.4 If G is k-connected, then G is k-edge-connected.

Proof. Suppose not. Let G be a k-connected graph so that $G - X$ is disconnected for some $X \subseteq E(G)$ with $|X| < k$. Let C be a component of $G - X$. Note that $|V(G)| \geq k + 1$ (by definition of k-connectedness) and $1 \leq |V(C)| \leq k - 1$ (as C is only one of the components in $G - X$).

We claim there exists some pair a, b of vertices with $a \in V(C)$ and $b \in V(G - C)$ so that no edge in X has ends a and b. Suppose not, i.e., there is an edge in X joining every pair of $a \in V(C)$ and $b \in V(G - C)$. Then to disconnect C from $G - C$ using X, we must remove at least

$$\overbrace{\underbrace{|V(C)|}_{\text{number of } a\text{'s}} \ \underbrace{(|V(G)| - V(C)|)}_{\text{number of } b\text{'s}}}^{\text{total number of such edges}} \geq k > |X|$$

edges from G, contrary to the fact that $|X| < k$. Thus, we can find a pair of vertices a, b where $a \in V(C)$ and $b \in V(G - C)$ such that X does not contain an edge with ends a, b.

For each $e \in X$, let v_e be an end of e that is not equal to a or b. Let $Y = \{v_e : e \in X\}$. Since X does not contain an edge with ends a and b, we can find such v_e for each $e \in X$, so $|Y| = |X| < k$. Since removing v_e necessarily removes e from G, the graph $G - Y$ is a subgraph of $G - X$ having both a and b as vertices. Since a and b are not connected in $G - X$, they are not connected in $G - Y$. But $|Y| < k$, contradicting k-connectedness of G. □

Remark. The converse is not necessarily true. Consider a graph such that two K_{11} (complete graphs with 11 vertices) sharing a cut vertex. This graph is not 2-connected because removing the cut vertex disconnects it, but is 10-edge connected because you need to remove at least ten edges to disconnect the graph since each graph vertex is connected to at least 10 other vertices.

Prop. 1.8.5 If G is k-edge-connected and $e \in E(G)$, then G/e is k-edge-connected. In other words, edge contraction does not destroy k-edge-connectedness. (Move this)

2 Basic Results

Highlights from the previous chapter.

1. G is k-connected if $|V(G)| > k$ and we cannot remove less than k vertices to disconnect G.
2. v is a cut vertex if there exists a connected pair of vertices a, b separated by $\{v\}$.
3. 1-connected $\iff$ connected and $|V| \geq 2$.
4. 2-connected $\iff$ connected with no cut vertices and $|V| \geq 3$.

2.1 Ear Decomposition

Def. 2.1.1 G' arises from G by *adding a path* if there exists a non-trivial path P, i.e., $|V(P)| > 1$, such that $G' = G \cup P$ and $(E(P) \cup V(P)) \cap (E(G) \cup V(G))$ is precisely the set of the two ends of P.

Lemma. 2.1.2 If G is 2-connected and G' is obtained from G by adding a path, then G' is 2-connected.

Proof. Clearly $|V(G')| \geq |V(G)| \geq 2$ by 2-connectedness of G. Let x and y be the ends of path P that was added to G to obtain G'. We show that for all distinct vertices a, b, v of G', there is an a, b-path in $G' - v$.

- If $a, b \in V(G)$, then by 2-connectedness of G, a and b are connected in $G - v$ and therefore also in $G' - v$.
- If exactly one of a and b, WLOG, say a, is a vertex of G, then since P is a xy-path containing b, b is connected to at least one of x and y in $G' - v$. Similarly, both x and y are connected to a in $G - v$ by 2-connectedness of G and thus in $G' - v$. It follows by transitivity that a and b are connected in $G' - v$.
- Finally, if $a, b \notin V(G)$, then since P is a path containing a and b, either a and b are connected in $G' - v$ (in which case the claimed statement holds), or one of a and b is connected in $G' - v$ to x and the other is connected to y in $G' - v$. Since x and y are connected in $G' - v$, it follows by transitivity that a and b are connected in $G' - v$.

Having considered all cases, we conclude the proof. □

Prop. 2.1.2 (Ear Decomposition) A loopless graph G is 2-connected if and only if there exists graphs $G_1, \ldots, G_k$ such that

1. G_1 is a circuit; $G_k = G$,
2. For each $1 \leq i < k$, G_i is 2-connected and G_{i+1} arises from G_i by adding a path.

Intuition for $\Longleftarrow$. Go as far as you can, then look at what's stopping you -- there must be an extra vertex. But then you could still add a path. Contradiction.

Proof.

$\Longleftarrow$: We can use **Lemma. 2.1.2** to show that $G_k = G$ obtained from a 2-connected graph G_{k-1} is also 2-connected.

$\Longrightarrow$: Let G be a *loopless* 2-connected graph, and let ℓ be maximal so that there exists subgraphs $G_1, \ldots, G_\ell$ of G so that each G_i is 2-connected and arises from G_{i-1} by adding a path, while G_1 is a circuit. Note that $\ell \geq 1$ because every 2-connected graph has a circuit (acyclic graphs have degree 1 vertices whose neighbours are cut vertices, i.e., can't be 2-connected).

Since adding a single new edge between vertices of G_ℓ is an example of adding a path (* -- *), the maximality of ℓ implies that every edge of G between two vertices of G_ℓ is an edge of G_ℓ, i.e., G_ℓ is an induced subgraph of G.

If $V(G_\ell) = V(G)$ then $G_\ell = G$ and the theorem holds, so we must have $V(G_\ell) \neq V(G)$. Since G is connected, there must be an edge from a vertex $u \notin V(G_\ell)$ to $v \in V(G_\ell)$.

Because G is 2-connected, there exists a path from u to $V(G_\ell)$ in the graph $G - v$, call it P. Now Pv is a path of G with both ends in $V(G_\ell)$ and no other vertices in $V(G_\ell)$.

Now the graph $G_\ell \cup (Pv)$ is a subgraph of G obtained from G_ℓ by adding a path. It is also 2-connected because of $\Longleftarrow$. This contradicts the maximality of ℓ. $\square$

Cor. 2.1.3 For any two vertices u and v in a 2-connected graph G, there is a circuit of G containing u and v.

Proof. Observe this is the $k = 2$ case of a version of Menger's theorem. We will provide a proof without using Menger's.

Define $G_1, \ldots, G_k$ as in **Prop. 2.1.2**, and let $\ell \in \{0, 1, \ldots, k\}$ be maximal so that every pair of vertices of G_ℓ are contained in a circuit. Since G_1 is a circuit, $\ell \geq 1$, and we may assume that $\ell < k$ since otherwise the result holds.

Let P be a path with ends x, y so that $G_{\ell+1}$ is obtained from G_ℓ by adding P. Let $a, b \in V(G_{\ell+1})$. We will show that some circuit of $G_{\ell+1}$ contains a and b.

- If $a, b \in V(G_\ell)$, then by the choice of ℓ, a circuit C of G_ℓ contains a and b; this C is also a circuit of $G_{\ell+1}$, as required.
- If $a, b \in V(P)$, then let Q by an xy-path in G_ℓ. Now P and Q are both xy-paths in $G_{\ell+1}$ that intersect only in $\{x, y\}$. Thus, $P \cup Q$ is a circuit of $G_{\ell+1}$ containing a and b.
- Finally, suppose $a \in V(G_\ell)$ and $b \notin V(G_\ell)$, so $b \in V(P)$. By the choice of ℓ, there is a circuit C_0 of G_ℓ containing a and x, and by the connectedness of G_ℓ, there is a yC_0-path Q_0 of G_ℓ. It is clear that $C_0 \cup Q_0$ contains an xy-path Q that contains a. Now $P \cup Q$ is a circuit of $G_{\ell+1}$ containing a and b as required.

Having considered all cases, we conclude the proof. $\square$

2.2 Block Graph

Recall the following definitions from the first chapter:

- A *separator* of G is a set $X \subseteq V(G) \cup E(G)$ s.t. $G - X$ is non-empty and disconnected.
- A *cut vertex* is a vertex v s.t. $\{v\}$ is a separator.

Def. 2.2.1 A *block* of G is a maximal connected subgraph of G with no cut vertex. That is, B is a block if $|B| \leq 2$ or B is 2-connected.

In a sense, blocks are the 2-connected analogues of components (which are 1-connected).

Def. 2.2.2 The *block graph* of a graph G is a simple bipartite graph with bipartition (B, K) s.t.

- B is the set of blocks of G,
- K is the set of cut vertices of G,
- A block $b \in B$ is adjacent in the block graph to a cut vertex w of G iff $w \in b$.

Example 2.2.3 (Block Graph) Red = Block, Blue = Cut Vertices.

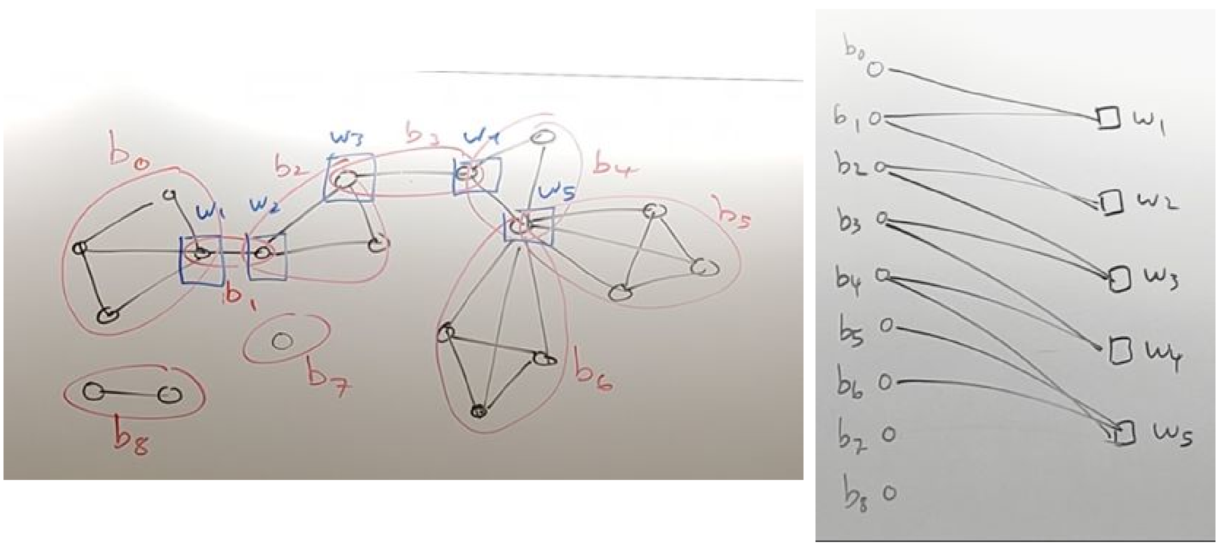

Observe the block graph "encodes" how the blocks are pieced together to form G; we can reconstruct the original graph from its block graph.

Since individual blocks are maximal 2-connected subgraphs, the block graph is a tree/forest, i.e., there is no cycle.

Prop. 2.2.4 Each graph G is the union of its blocks.

Proof. It suffices to show that each vertex and edge of G is contained in a block of G. This is true because for every edge e from x to y, the subgraph on x, y, e has no cut vertex, so is contained in some maximal subgraph with no cut vertex, i.e., a block. □

Prop. 2.2.5 Any two blocks intersect in at most one vertex (and no edges).

Proof. Let B_1, B_2 be blocks, $x, y \in V(B_1) \cap V(B_2)$ be distinct. By maximality of B_1 and B_2, the subgraph $B_1 \cup B_2$ has a cut vertex w. Let a, b be vertices separated by w. Let $z \in \{x, y\} \setminus \{w\}$. Since B_1 and B_2 have no cut vertex, each of a and b is connected to z in either $B_1 - w$ or $B_2 - w$ and hence in $B_1 \cup B_2 - w$. Then a, b are connected in $B_1 \cup B_2 - w$. Contradiction. □

Cor. 2.2.6 If a vertex w of G is contained in more than one block of G, then w is a cut vertex of G.

Proof. Suppose $v \in V(G)$ is contained in distinct blocks B_1 and B_2 of G. Since B_1 and B_2 are distinct connected induced subgraphs that intersect in at most one vertex, there is a neighbour v_1 of v in B_1 and a neighbour v_2 of v in B_2. We show that v separates v_1 and v_2 in G.

Suppose not, let P be a v_1v_2-path in $G - v$. By choice of v_1 and v_2, there is a circuit C of G with vertex set $V(P) \cup \{v\}$. Now C is a cut-vertex-free subgraph (it is a circuit) of G, so C is contained in a block B' of G (by **Prop. 2.2.4**). But $V(B_1) \cap V(B')$ contains both v and v_1, contradicting (**Prop. 2.2.5**) the fact that two blocks intersect in at most one vertex. □

Cor. 2.2.7 The block graph of G is a forest.

Proof. Suppose not, so the block graph contains a smallest circuit C. Since the block graph is bipartite, it follows that the vertices of C can be enumerated in order $v_1, B_1, v_2, B_2, \ldots, v_k, B_k$ for some $k \geq 2$, where the B_i's are blocks and the v_i's are cut vertices, where $v_i \in V(B_{i-1}) \cap V(B_i)$ for each i and $B_0 = B_k$ by definition. Since each block B_i is connected, it contains a v_iv_{i+1}-path P_i (where $v_{k+1} = v_1$ by definition).

We now argue that the P_i intersect only when expected. If there are distinct P_i and P_j that intersect at some internal vertex w of P_i, then w is contained in B_i and B_j, so by **Cor. 2.2.6** is a cut vertex of B_i, contradicting the fact that B_i is a block. It follows that the union of the P_i is a circuit of G. But this circuit is a 2-connected subgraph of G that intersects the block B_1 in both v_1 and v_2, contradicting the fact (**Prop. 2.2.5**) that the intersection of two blocks is at most one vertex. □

Prop. 2.2.8 A graph is bipartite if and only if its blocks are bipartite.

Proof. Let G be a minimal counterexample. That is, each block of G is bipartite but G is not. Clearly $|V(G)| > 1$. Let H_0 be a block of G with at most one neighbour in the block graph of G. Such block exists because the block graph of G is a forest (**Cor. 2.2.7**) and every non-trivial forest has a leaf.

Let X be the set of vertices of H_0 contained in another block of G. **Cor. 2.2.6**, each vertex of X is a cut vertex of G and H_0 contains at most one cut vertex of G, so $|X| \leq 1$.

Let $H = G - (V(H_0) \setminus X)$. Now $G = H \cup H_0$ and $V(H) \cap V(H_0) = X$. Moreover, each block of H is a block of G, so it is bipartite. Also, there is no edge from a vertex w in $H_0 - X$ to a vertex in $H - X$, because such an edge would be contained in a block of G other H_0 which intersects H_0 in the vertex w that is not a cut vertex of G, contradicting (**Prop. 2.2.5**).

Let (A_0, B_0) be a bipartition of H_0 and (A, B) be a bipartition of H, chosen so that $A \cap A_0$ contains X (this can be done as $|X| \leq 1$ and the sides of a bipartition can be swapped). Since there are no edges from $V(H - X)$ to $V(H_0 - X)$, it follows that $(A \cup A_0, B \cup B_0)$ is a bipartition of G, contrary to the choice of G as a counterexample. □

2.3 Edge Contraction

Def. 2.3.1 Given an edge e in a graph $G = (V, E, i)$ with ends u, v, let G/e denote the graph with the edge set $E \setminus \{e\}$, vertex set $(V \setminus \{u, v\}) \cup \{x_{uv}\}$ where x_{uv} is a new vertex, in which each edge with an end equal to u or v in G now has an end equal to the new vertex x_{uv} replacing it.

Remark.

1. Any edge parallel to e becomes a loop at the new vertex.
2. If e is a loop, delete it and keep the vertex with the same name.
3. Edge contraction does not break connectedness, i.e., if G is connected, then so is G/e.
4. Edge contraction may break 2-connectedness. Consider a cycle with an edge connecting two non-adjacent vertices; contracting this edge would create a cut vertex.

Lemma 2.3.2 Let $k \geq 1$. If G is k-connected and $X \subseteq V(G)$ with $|X| = k$, then each vertex in X has a neighbour in each component of $G - X$.

Proof. If there were a vertex $x \in X$ and a component C of $G - x$ so that x has no neighbour in C , then $X \setminus \{x\}$ would be a set of $< k$ vertices so that $G - (X \setminus \{x\})$ is disconnected. This contradicts the k-connectedness of G. □

Prop. 2.3.3 If e is an edge of G with ends u, v and $X \subseteq V \cup E$ containing neither e, u, nor v, then $(G - X)/e = (G/e) - X$.

Intuition. The order of contradiction and deletion does not matter as long as they do not affect the same edge. Proof omitted. □

Prop. 2.3.4 If G is k-edge-connected and $e \in E(G)$, then G/e is k-edge-connected. In other words, edge contraction does not destroy k-edge-connectedness.

Proof. Let e be an edge of a k-edge-connected graph G such that G/e is not k-edge-connected. Let a, b be a disconnected pair of vertices of $(G/e) - X$ for some set $X \subseteq E(G/e)$ with $|X| < k$. Since G is k-edge-connected, there is an a, b-path P in $G - X$. If $e \notin E(P)$, then $E(P)$ is the edge set of an a, b-path of G/e. If $e \in E(P)$, then P/e is a subgraph of G/e and so P/e is an a, b-path in G/e, contrary to the choice of a and b. □

2.4 2-Connectedness

Prop. 2.4.1 If G is 2-connected, $e \in E(G)$, and $|V(G)| > 3$, then either G/e or $G - e$ is 2-connected.

Proof. Suppose that neither G/e nor $G - e$ is 2-connected.

Let y, z be vertices so that $(G/e) - y$ and $G - \{e, z\}$ are disconnected graphs.

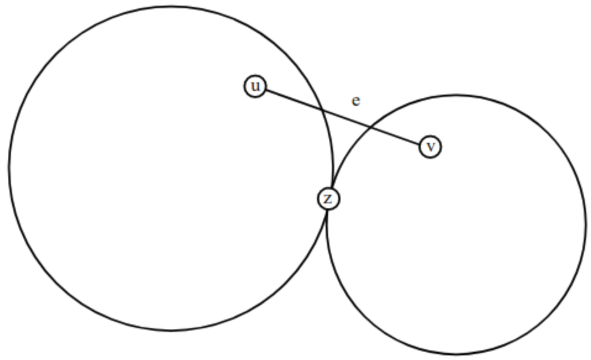

Let u, v be the ends of e (with $u \neq v$). If $y \neq x_{uv}$, then $(G/e) - y = (G - y)/e$ but $G - y$ is connected so $(G - y)/e$ is connected, a contradiction to the choice of y. Therefore $y = x_{uv}$.

Since $y = x_{uv}$, we have $(G/e) - y = G - \{u, v\}$. We will show that $G - e$ must be 2-connected. To do this, we will show that z is not a cut vertex of $G - e$.

First, note that if $z \in \{u, v\}$, then $G - e - z = G - z$ which cannot be disconnected since G is 2-connected. Thus, $z \notin \{u, v\}$.

By assumption, $(G/e) - y$ is disconnected. Let C be a component of $(G/e) - y = G - \{u, v\}$ that does not contain z. By the lemma above, C contains a neighbour u' of u and v' of v. Since C does not contain z, u, or v, removing $\{e, z\}$ from G does not affect the connectivity within C, i.e., u' and v' are connected in $G - \{e, z\}$. It follows from transitivity that u and v are connected in $G - \{e, z\}$.

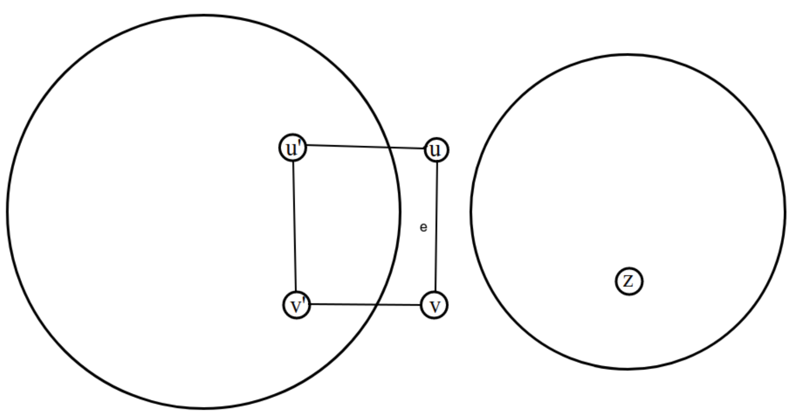

Let a be a vertex of $G - z$. Since $G - z$ is connected, there exists a path from a to u or v in $G - \{z, e\}$. Therefore,

1. a is connected to u or v in $G - \{z, e\}$,
2. u and v are connected in $G - \{z, e\}$, and
3. Thus, a and u are connected in $G - \{z, e\}$.

Since a is arbitrary, $G - \{z, e\}$ is connected. This contradicts the choice of z. The proof is complete. □

Lemma. 2.4.2 If a graph contains a degree-2 vertex, then it is not 3-connected.

Intuition. Removing its two neighbours disconnects the graph. □

Cor. 2.4.3 For a k-connected graph, the minimum degree must have at least k. Let $\delta(G)$ denote the min degree of a graph and $k(G)$ denote the vertex connectivity, then $\delta(G) \geq k(G)$.

Intuition. Removing the vertex with minimum degree necessarily disconnects the graph. □

Remark. The "subtract-contraction" proposition (**Prop 2.4.1**) works well for 2-connectedness, but we cannot hope for a version of it for 3-connectedness. Consider the following example.

Ex. 2.4.4 An n-wheel is a graph with n vertices on the perimeter and one vertex in the middle. For example, the 3-wheel is a complete graph on 4 vertices and a 4-wheel can be viewed as a square with four corners connected to the center.

Observe that an n-wheel for $n \geq 3$ is 3-connected, but its "spoke" edges can be neither delete nor contracted while maintaining 3-connectedness. Thus, we cannot hope for a version of the above proposition for 3-connectedness.

2.5 3-Connectedness

Thm. 2.5.1 (Tutte) If $G = (V, E)$ is a 3-connected graph with $|V(G)| > 4$, then G has an edge e such that G/e is 3-connected.

Lemma. 2.5.2 If Tutte's Theorem does not hold, then we claim that for every edge e with distinct ends x, y, there is a vertex $z \notin \{x, y\}$ such that $G - \{x, y, z\}$ disconnected, and each of the vertices x, y, z has a neighbour in each component of $G - \{x, y, z\}$.

Proof. Let a_{xy} be the vertex of G/e created by the contraction. Since G/e is not 3-connected (assumption), there exists a set of vertices $Z \subseteq V$ such that $|Z| \leq 2$ and $(G/e) - Z$ is disconnected.

Suppose $a_{xy} \notin Z$. Since $|Z| \leq 2$ and G is 3-connected, $G - Z$ is connected. By **Prop. 2.3.3**, we can exchange the order of subtraction and edge contraction, so $(G/e) - Z = (G - Z)/e$ is connected, as edge contraction does not destroy connectedness. Contradiction. Thus, $a_{xy} \in Z$.

Let z be the other element of Z. The disconnected graph $(G/e) - Z$ is now equal to $G - \{x, y, z\}$. The other part follows from what we proved in **Prop. 2.3.3**. ■

Proof. (Tutte) Suppose Tutte's theorem is false. Choose e with ends x, y so that the number of vertices of a smallest component C of $G - \{x, y, z\}$ is as small as possible, where z is given by **Lemma. 2.5.2**. Let v be a neighbour of z in C.

By Lemma, there exists $w \in V$ so that $G - \{v, z, w\}$ is disconnected and each of v, z, w has a neighbour in each component in $G - \{v, z, w\}$. Since x and y are adjacent in G (so they must be in the same component afterwards) and $G - \{v, z, w\}$ has ≥ 2 components, there is a component D of $G - \{v, z, w\}$ containing neither x nor y. (i.e., D is also a component of $G - \{v, z, w, x, y\}$.)

We argue that $V(D) \subseteq V(C)$. This will lead to a contradiction since $v \in V(C) \setminus V(D)$ and C was chosen to be as small as possible. To see this, let $b \in V(D)$. Since v has a neighbour in D and D is a component of $G - \{v, z, w, x, y\}$, there is a path from b to v in the graph of $G - \{z, w, x, y\}$.

Then there exists a path from b to v in $G - \{x, y, z\}$, so b is in the same component as v in the graph $G - \{x, y, z\}$, and thus $b \in V(C)$. This proves $V(D) \subseteq V(C)$. The proof is complete. $\square$

3 Menger's Theorem

3.1 Menger's Theorem for Sets of Vertices

Thm. 3.1.1 (Menger's Theorem For Sets of Vertices) Given $A, B \subseteq V(G)$ and k being the cardinality of a minimal set X such that $G - X$ has no A, B-paths, then there are k vertex-disjoint A, B-paths in G.

Proof. Suppose the statement is false. Let G, A, B, k specify a counterexample where $|E(G)|$ is as small as possible.

If every edge of G is a loop, then every A, B-separator contains $A \cap B$ and $A \cap B$ itself is an A, B -separator, so $k = |A \cap B|$ is the size of a smallest A, B-separator. But each vertex in $A \cap B$ is an A, B-path, so there are k vertex-disjoint A, B-paths. Therefore G is not a counterexample. It follows that G must contain an edge e with ends u, v with $u \neq v$.

Let x_e be the identified vertex of G/e. Define

$$A' = \begin{cases} A, & u, v \notin A \\ (A \setminus \{u, v\}) \cup \{x_e\} & \{u, v\} \cap A \neq \varnothing \end{cases}$$
$$B' = \begin{cases} B, & u, v \notin B \\ (B \setminus \{u, v\}) \cup \{x_e\} & \{u, v\} \cap B \neq \varnothing \end{cases}$$

Then A', B' both are sets of vertices of G/e. We will show that if there are k disjoint A', B'-paths in G/e, then there are k disjoint A, B-paths in G.

Claim 1. Let H be a subgraph of G containing e. Let $H' = H/e$. Then there exists an A, B-path in G if and only if there exists an A', B'-path in H'.

Proof. Let C be component of H and C' be the corresponding component of H'. That is, $C = C'$ if $e \neq C$ and $C' = C/e$ otherwise. Then C contains a vertex in A iff C' contains a vertex in A' by definition of A'. The same goes for B and B'. Then H contains an A, B-path iff some component C of H contains a vertex in A and a vertex in B iff some component C' of H' contains a vertex in A' and a vertex in B' iff H' contains an A', B'-path. ■

Claim 2. There does not exist k disjoint A', B'-paths in G/e.

Proof. Suppose that disjoint A', B'-paths $P_1, \ldots, P_k$ existed in G/e. Each path P_i not containing x_e is also an A, B-path in G. If none of the P_i contains x_e, then G has k disjoint A, B-paths, contradicting the choice of G as a counterexample.

Now suppose one of the P_i, say P_1, contains x_e. Let H be the subgraph of G with vertex set $(V(P_1) \setminus \{x_e\}) \cup \{u, v\}$ and edge set $E(P_1) \cup \{e\}$. Now $H/e = P_1$. By Claim 1, since P_1 has/is an A', B'-path, H contains an A, B-path $\hat{P}_1$. Since $P_1, \ldots, P_k$ are vertex-disjoint, the paths $\hat{P}_1, \ldots, P_k$ are vertex-disjoint A, B-paths in G, again a contradiction. ■

By the minimality of $E(G)$, G/e is not a counterexample. It follows that the smallest A', B'-separator X' in G/e has size less than k.

Claim 3. Let X' be a minimal A', B'-separator in G/e. Then $x_e \in X'$.

Proof. If $x_e \notin X'$, then $(G/e) - X' = (G - X')/e$, so $(G - X')/e$ contains no A', B'-path, as X' is a separator. Applying Claim 1 with $H = G - X'$, we see that $H = G - X'$ has no A, B-path. Thus, X' is a separator of size $< k$ in G, a contradiction. Therefore, $x_e \in X'$. ■

Let us "uncontract" e from G: define $X = (X' \setminus \{x_e\}) \cup \{u, v\}$. Note that $(G/e) - X' = G - X$, so X separates A from B in G. This gives $|X| \geq k$ (assumption). Since $|X'| < k$ and $|X| = |X'| + 1$, we get $|X| = k$. Let $X = \{x_1, \ldots, x_k\}$.

If there were an (A, X)-separator Y in $G - e$ with $|Y| < k$, then $(G - e) - Y$ would contain no A, X paths. Since every A, B-path contains an A, X-path, this implies that Y is an A, B-separator, so $|Y| < k$ gives a contradiction. Thus, the smallest size of an A, B-separator is $\geq k$. Similarly, the smallest size of an X, B-separator is $\geq k$.

Since $G - e$ is not a counterexample, it follows there are k disjoint A, X-paths $P_1, \ldots, P_k$ in $G - e$. Say that x_i is the end in X of P_i. Similarly, we can find k disjoint X, B-paths $Q_1, \ldots, Q_k$, where x_i is the end in X of Q_i.

If there were a vertex $z \in (P_i - X) \cap (Q_j - X)$ for some i and j, then let a_i be at the end of P_i in A and b_j be the end of Q_j in B, then a and z are connected in $G - X$, and z and b are connected in $G - X$, so a and b are connected in $G - X$, contradicting the fact that X is a separator. Thus, the paths P_i, Q_j intersect only at their ends in X for any i and j. Thus, $P_i x_i Q_i : 1 \leq i \leq k$ gives a collections of k disjoint A, B-paths in G, a contradiction. □

3.2 Menger's Theorem for Two Vertices

Def. 3.2.1 Let $a, b \in V(G)$ and $a \neq b$. a, b-paths $P_1, \ldots, P_k$ are *internally disjoint* if the sets $V(P_1) \setminus \{a, b\}, \ldots, V(P_k) \setminus \{a, b\}$, and if a, b are adjacent $E(P_1), \ldots, E(P_k)$ are disjoint sets.

Thm. 3.2.2 (Menger's Theorem for Two Vertices) If a, b are nonadjacent vertices of G and k is the size of a smallest a, b-separator X with $a, b \notin X$, then there are k internally disjoint a, b-paths in G.

Intuition. Apply Menger's theorem on the neighbour sets of a and b. Note that the neighbour sets can be overlapping. This gives you k internally disjoint A, B-paths. Some of the paths might be trivial. You can then glue a and b onto both ends of the paths and get k internally disjoint paths as required.

Proof. Let $A, B \subseteq V(G)$ be the neighbour sets of a and b, respectively. Let k' be the smallest size of an A, B-separator X' in G. Since $G - X'$ contains no A, B-paths, and every a, b-path contains an A, B-path, X' is necessarily an a, b-separator, so $k' = |X'| \geq k$. Applying Menger's Theorem for Sets of Vertices, there are k vertex-disjoint A, B-paths $P_1, \ldots, P_k$. It follows that there are k internally disjoint a, b-paths of the form $aP_1b, \ldots, aP_kb$. □

3.3 Fan Lemma

Lemma 3.3.1 (Fan Lemma) If a is a vertex of a graph G and $B \subseteq V(G)$ with $a \notin B$, then one of the following holds:

1. There exist k paths $P_1, \dots, P_k$, each starting at a, all disjoint except their intersection at a, and all intersecting B in precisely the last vertex.
2. There is a set $X \subseteq V(G) \setminus \{a\}$ so that $|X| < k$ and $G - X$ has no a, B-path.

Proof. Let A be the set of vertices of G adjacent to a. By Menger's theorem for sets, there is either a set of k vertex-disjoint A, B-paths, or there is a set $X \subseteq V(G)$ with $|X| < k$ for which $G - X$ has no A, B-paths. For the second case, since $a \notin B$, every a, B-path contains an A, B-path, so $G - X$ has no A, B-paths implies $G - X$ has no a, B-paths, as required.

It remains to prove the first case. Let $P_1, \dots, P_k$ be vertex-disjoint A, B-paths. If none of these paths contain a, then $aP_1, \dots, aP_k$ is a collection of k vertex-disjoint a, B-paths that intersect only at a. If one of the paths, say P_i, contains a, then let x be the vertex occurring after a in P_i. Since $a \notin B$, this choice is well-defined, and we have $x \in A$ by the definition of A. Thus, P_i contains an A, B-path Q which does not contain a. Then the paths $aP_1, \dots, aP_{i-1}, aP_{i+1}, aP_k, aQ$ form a collection of k distinct a, B-paths that intersect only at a, as required. □

Prop. 3.3.2 If G is a k-connected graph and $A \subseteq V(G)$ with $|A| \leq k$, then G has a circuit containing each vertex in A.

Proof. Consider a circuit C containing as many vertices from A as possible. (Recall 2-connected graphs have circuits, so such C must exist.) Let $a \in A \setminus V(C)$. Since G is k-connected, there is no set $X \subseteq V(G) \setminus \{a\}$ of size less than $\min(|C|, k)$ such that there are no a, C-paths in $G - X$. Therefore there are at least $\min(|C|, k)$ paths from a to C.

Let $P_1, \dots, P_\ell$ be the paths formed by C between the elements of $V(C) \cap A$. Since $|A \cap V(C)| < k$, we have $\ell = |A \cap C| < k$.

By the Fan Lemma, there exists $\min(|V(C)|, k)$ paths from a to C that only intersect at a. In In either case (that is, if $k < |V(C)|$ or $k \geq |V(C)|$), there exists some i s.t. P_i contains the end of C of these two paths Q, Q'.

Now there is a circuit C' contained in $C \cup Q \cup Q'$, containing all vertices $V(C) \cap A$ but also the vertex a. Thus, contradicts the maximality of C. □

3.4 Other Versions of Menger's Theorem

Thm. 3.4.1 If G is a graph with $|V(G)| \geq k$, then the following are equivalent:

1. G is k-connected.
2. For every $a, b \in V(G)$ with $a \neq b$, there are k internally disjoint a, b-paths in G.

Thm. 3.4.2 Let $a, b \in V(G)$ with $a \neq b$. For $k \geq 1$, exactly one of the following holds:

1. There exists k internally disjoint a, b-paths in G.

2. There exists a set $X \subseteq V(G) \setminus \{a, b\}$ such that $|X| \leq k$ and $G - X$ contains no a, b paths.

Thm. 3.4.3 Let $A, B \subseteq V(G)$. For $k \geq 1$, exactly one of the following holds:

1. There are k vertex-disjoint A, B-paths in G.
2. There exists $X \subseteq V(G)$ with $|X| < k$ such that $G - X$ has no A, B-paths.