Math 530 Spring 2022, Lecture 27: Menger's theorems

website: `https://www.cip.ifi.lmu.de/~grinberg/t/22s`

1. More about paths

In this and the next lecture, we will learn a few more things about paths in graphs and digraphs.

1.1. Menger's theorems

We begin with a series of fundamental results known as **Menger's theorems** (named after Karl Menger, who discovered one of them in 1927 as an auxiliary result in a topological study of curves[1]).

Imagine you have 4 different ways to get from Philadelphia to NYC, all using different roads (i.e., no piece of road is used by more than one of your 4 ways). Then, if 3 arbitrary roads get blocked, then you still have a way to get to NYC.

This is obvious (indeed, each blocked road destroys at most one of your 4 paths, so you still have at least one path left undisturbed after 3 roads have been blocked). A more interesting question is the converse: If the road network is sufficiently robust that blocking 3 arbitrary roads will not disconnect you from NYC, does this mean that you can find 4 different ways to NYC all using different roads?

Menger's theorems answer this question (and various questions of this kind) in the positive, in several different setups. Each of these theorems can be roughly described as "the maximum number of pairwise independent paths from some place to another place equals the minimum size of a bottleneck that separates the former from the latter". Here, the "places" can be vertices or sets of vertices; the word "independent" can mean "having no arcs in common" or "having no intermediate vertices in common" or "having no vertices at all in common"; and the word "bottleneck" can mean a set of arcs or of vertices whose removal would disconnect the former place from the latter. Here is a quick overview of all Menger's theorems that we will prove:[2]

[1]See [Schrij03, §9.6e] for more about its history.

[2]All undefined terminology used here will be defined further below.

- for directed graphs:

Theorem ...	the places are ...	the paths must be ...	the bottleneck consists of ...
1.1.6	vertices	arc-disjoint	arcs
1.1.10	vertices	arc-disjoint	arcs of a cut
1.1.18	sets of vertices	arc-disjoint	arcs of a cut
1.1.32	vertices	internally vertex-disjoint	vertices $\in V \setminus \{s, t\}$
1.1.36	sets of vertices	internally vertex-disjoint	vertices $\in V \setminus (X \cup Y)$
1.1.39	sets of vertices	vertex-disjoint	vertices $\in V$

- for undirected graphs:

Theorem ...	the places are ...	the paths must be ...	the bottleneck consists of ...
1.1.22	vertices	edge-disjoint	edges
1.1.25	vertices	edge-disjoint	edges of a cut
1.1.42	vertices	internally vertex-disjoint	vertices $\in V \setminus \{s, t\}$
1.1.44	sets of vertices	internally vertex-disjoint	vertices $\in V \setminus (X \cup Y)$
1.1.45	sets of vertices	vertex-disjoint	vertices $\in V$

(I could state more, but I don't want this to go on forever.)

1.1.1. The arc-Menger theorem for directed graphs

We begin with the most natural setup: a directed graph (one-way roads) with roads being arcs. The following definitions will help keep the theorems short:

Definition 1.1.1. Two walks **p** and **q** in a digraph are said to be **arc-disjoint** if they have no arc in common.

Example 1.1.2. The following picture shows two arc-disjoint paths **p** and **q** (they can be told apart by their labels: each arc of **p** is labelled with a "**p**", and likewise for **q**):

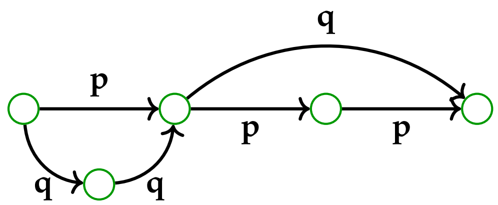

.

The following picture shows two paths **r** and **s** that are **not** arc-disjoint (the common arc is marked with "**r**, **s**"):

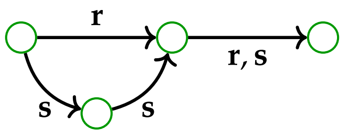

.

Definition 1.1.3. Let $D = (V, A, \psi)$ be a multidigraph, and let s and t be two vertices of D. A subset B of A is said to be an ***s*-*t*-arc-separator** if each path from s to t contains at least one arc from B. Equivalently, a subset B of A is said to be an ***s*-*t*-arc-separator** if the multidigraph $\left(V,\ A \setminus B,\ \psi\mid_{A\setminus B}\right)$ has no path from s to t (in other words, removing from D all arcs contained in B destroys all paths from s to t).

Example 1.1.4. Let $D = (V, A, \psi)$ be the following multidigraph:

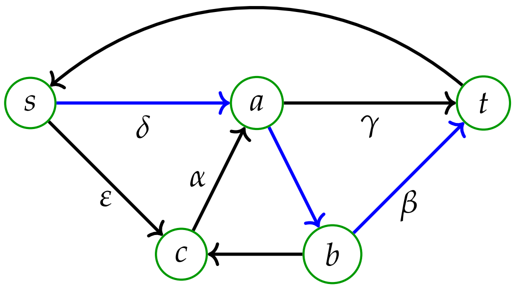

.

Then, the set $\{\alpha, \gamma\}$ is not an s-t-arc-separator (since the path drawn in blue contains no arc from this set). However, the set $\{\beta, \gamma\}$ is an s-t-arc-separator, and so is the set $\{\delta, \varepsilon\}$. Of course, any set that contains any of $\{\beta, \gamma\}$ and $\{\delta, \varepsilon\}$ as a subset is therefore an s-t-arc-separator as well.

Example 1.1.5. Let D be a multidigraph. Let s and t be two vertices of D. Then, the empty set $\varnothing$ is an s-t-arc-separator if and only if D has no path from s to t. This degenerate case should not be forgotten!

We can now state the first Menger's theorem:

Theorem 1.1.6 (arc-Menger theorem for directed graphs, version 1)**.** Let $D = (V, A, \psi)$ be a multidigraph, and let s and t be two distinct vertices of D. Then, the maximum number of pairwise arc-disjoint paths from s to t equals the minimum size of an s-t-arc-separator.

Example 1.1.7. Let D be the multidigraph from Example 1.1.4. Then, the minimum size of an s-t-arc-separator is 2 (indeed, $\{\beta, \gamma\}$ is an s-t-arc-separator of size 2, and it is easy to see that there are no s-t-arc-separators of smaller size). Hence, Theorem 1.1.6 yields that the maximum number of pairwise arc-disjoint paths from s to t is 2 as well. And indeed, we can easily find 2 arc-disjoint paths from s to t, namely the red and the blue paths in the following figure:

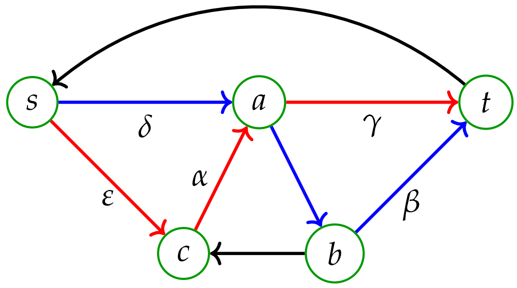

.

Before proving Theorem 1.1.6, let me state another variant of this theorem, which is closer to the proof. First, some notations:

Definition 1.1.8. Let $D = (V, A, \psi)$ be a multidigraph, and let s and t be two distinct vertices of D.

(a) For each subset S of V, we set $\overline{S} := V \setminus S$ and

$$[S, \overline{S}] := \{a \in A \mid \text{the source of } a \text{ belongs to } S, \\ \text{and the target of } a \text{ belongs to } \overline{S}\}.$$

(These are the same definitions that we introduced for networks in Lecture 26.)

(b) An ***s*-*t*-cut** means a subset of A that has the form $[S, \overline{S}]$, where S is a subset of V that satisfies $s \in S$ and $t \notin S$. (This was just called a "cut" back in Lecture 26.)

An s-t-cut is called this way because its removal would cut the vertex s from the vertex t. More precisely:

Remark 1.1.9. Let $D = (V, A, \psi)$ be a multidigraph, and let s and t be two distinct vertices of D. Then, any s-t-cut is an s-t-arc-separator.

Proof. Let B be an s-t-cut. We must prove that B is an s-t-arc-separator. In other words, we must prove that each path from s to t contains at least one arc from B.

We know that B is an s-t-cut. In other words, $B = [S, \overline{S}]$, where S is a subset of V that satisfies $s \in S$ and $t \notin S$. Consider this subset S.

Each path from s to t starts at a vertex in S (since $s \in S$) and ends at a vertex outside of S (since $t \notin S$). Thus, each such path has to escape the set S at some point – i.e., it must contain an arc whose source is in S and whose target is outside of S. But such an arc must necessarily belong to $[S, \overline{S}]$ (by the definition of $[S, \overline{S}]$). Thus, each path from s to t must contain an arc from $[S, \overline{S}]$. In other words, each path from s to t must contain an arc from B (since $B = [S, \overline{S}]$). In other words, B is an s-t-arc-separator (by the definition of an s-t-arc-separator). This proves Remark 1.1.9. $\square$

Theorem 1.1.10 (arc-Menger theorem for directed graphs, version 2)**.** Let $D = (V, A, \psi)$ be a multidigraph, and let s and t be two distinct vertices of D. Then, the maximum number of pairwise arc-disjoint paths from s to t equals the minimum size of an s-t-cut.

Example 1.1.11. Let D be the following multidigraph:

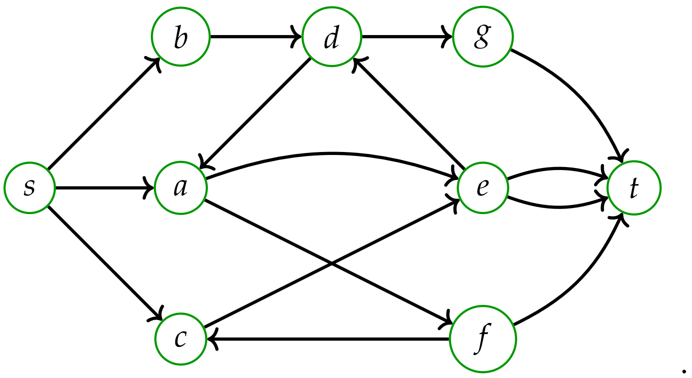

Then, the maximum number of pairwise arc-disjoint paths from s to t is 3. Indeed, the following picture shows 3 such paths in red, blue and brown, respectively:

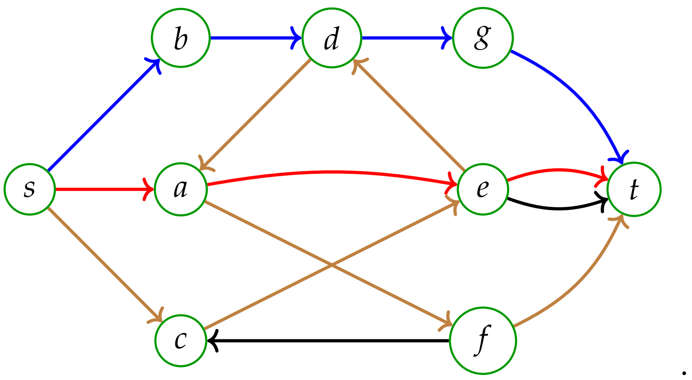

More than 3 pairwise arc-disjoint paths from s to t cannot exist in D, since (e.g.) there are only 3 arcs outgoing from s.

By Theorem 1.1.10, this shows that the minimum size of an s-t-cut in D is 3 as well. There are many s-t-cuts of size 3 (for instance, the "obvious" cut $\left[\{s\}, \overline{\{s\}}\right]$ has this property, as does the s-t-cut $\left[\{s,a,f\}, \overline{\{s,a,f\}}\right]$).

Let us now reverse of the direction of the arc from c to e in D (thus destroying the brown path). The resulting multidigraph D' looks as follows:

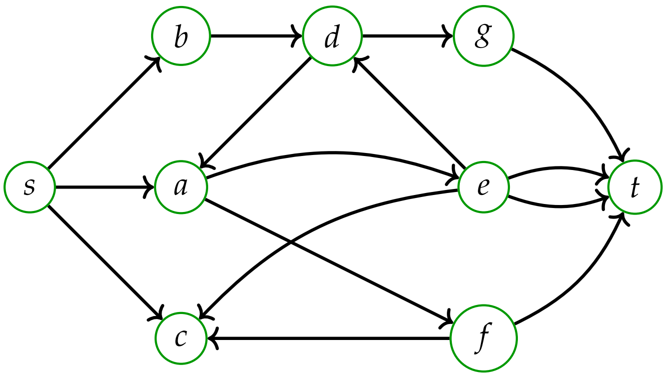

This digraph D' has no more than 2 pairwise arc-disjoint paths from s to t. This can be seen by observing that the s-t-cut $\left[\{s,c\}, \overline{\{s,c\}}\right]$ has size 2 (it consists of the arc from s to a and the arc from s to b), so that the minimum size of an s-t-cut is at most 2, and therefore (by Theorem 1.1.10) the maximum number of pairwise arc-disjoint paths from s to t is at most 2 as well. It is easy to see that the latter number is exactly 2 (since our red and blue paths still exist in D').

To prove the above two arc-Menger theorems, we need one more lemma about networks. We recall the notations from Lecture 26, and introduce a couple more:

Definition 1.1.12. Let $D = (V, A, \psi)$ be a multidigraph. Let $f, g : A \to \mathbb{N}$ be two maps. Then:

(a) We let $f + g$ denote the map from A to $\mathbb{N}$ that sends each arc $a \in A$ to $f(a) + g(a)$. (This is the pointwise sum of f and g.)

(b) We write $g \leq f$ if and only if each arc $a \in A$ satisfies $g(a) \leq f(a)$.

(c) If $g \leq f$, then we let $f - g$ denote the map from A to $\mathbb{N}$ that sends each arc $a \in A$ to $f(a) - g(a)$. (This is really a map to $\mathbb{N}$, since $g \leq f$ entails $g(a) \leq f(a)$.)

These notations satisfy the properties that you'd expect: e.g., the pointwise sum of maps from A to $\mathbb{N}$ is associative (meaning that $(f+g)+h = f+(g+h)$, so that you can write $f+g+h$ for both sides); inequalities can be manipulated in the usual way (e.g., we have $f - g \leq h$ if and only if $f \leq g + h$). Verifying this all is straightforward.

The following definition codifies the flows that we constructed in Remark 1.1.9 in Lecture 26:

Definition 1.1.13. Let N be a network consisting of a multidigraph $D = (V, A, \psi)$, a source $s \in V$, a sink $t \in V$ and a capacity function $c : A \to \mathbb{N}$. Let $\mathbf{p}$ be a path from s to t in D. Then, we define a map $f_{\mathbf{p}} : A \to \mathbb{N}$ by setting

$$f_{\mathbf{p}}(a) = \begin{cases} 1, & \text{if } a \text{ is an arc of } \mathbf{p}; \\ 0, & \text{otherwise} \end{cases} \qquad \text{for each } a \in A.$$

We call this map $f_{\mathbf{p}}$ the **path flow** of $\mathbf{p}$. It is an actual flow of value 1 if all the arcs of $\mathbf{p}$ have capacity ≥ 1.

Example 1.1.14. Consider the following network:

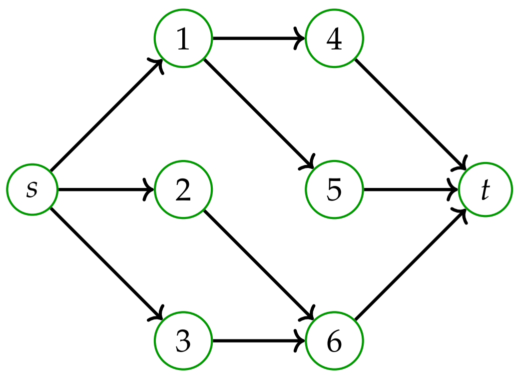

,

where each arc has capacity 1. Then, the path $\mathbf{p} = (s, 2, 6, t)$ leads to the following path flow $f_{\mathbf{p}}$:

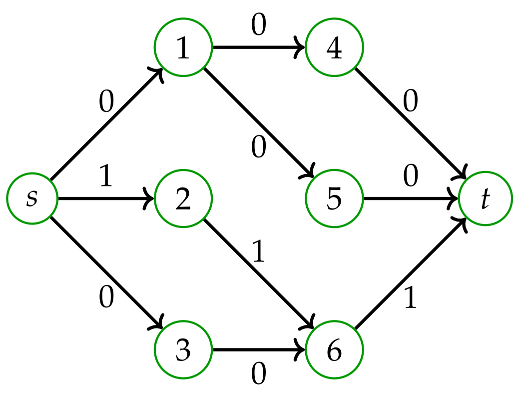

.

Here, in order not to crowd the picture, we have left out the "of 1" part of the label of each arc (so you should read the "0"s and the "1"s atop the arcs as "0 of 1" and "1 of 1", respectively).

The path flow thus turns any path from s to t in a network into a flow, provided that the arcs have enough capacity to carry this flow. If we have m paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ from s to t, then we can add their path flows together, and obtain a flow $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}$ of value m, provided (again) that the arcs have enough capacity for it. (In general, we cannot uniquely reconstruct $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ back from this latter flow, as they might have gotten "mixed together".)

Our next lemma can be viewed as a (partial) converse of this observation: Any flow f of value m "contains" a sum $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}$ of m path flows $f_{\mathbf{p}_1}, f_{\mathbf{p}_2}, \ldots, f_{\mathbf{p}_m}$ corresponding to m (not necessarily distinct) paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ from s to t. Here, the word "contains" signals that f is not necessarily equal to $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}$, but only satisfies $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m} \leq f$ in general. So here is the lemma:

Lemma 1.1.15 (flow path decomposition lemma)**.** Let N be a network consisting of a multidigraph $D = (V, A, \psi)$, a source $s \in V$, a sink $t \in V$ and a capacity function $c : A \to \mathbb{N}$. Let f be a flow on N that has value m. Then, there exist m paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ from s to t in D such that

$$f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m} \leq f.$$

Proof. We induct on m.

The *base case* ($m = 0$) is obvious (since the empty sum $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}$ is the zero flow, and thus is $\leq f$ because of the capacity constraints).

Induction step: Let m be a positive integer. Assume (as the induction hypothesis) that the lemma holds for $m - 1$. We must prove the lemma for m.

So we consider a flow f on N that has value m. We need to show that there exist m paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ from s to t in D such that $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m} \leq f$.

We shall first find some path $\mathbf{p}$ from s to t such that $f_{\mathbf{p}} \leq f$.

We shall refer to the arcs $a \in A$ satisfying $f(a) > 0$ as the **active** arcs. Let $A' := \{a \in A \mid f(a) > 0\}$ be the set of these active arcs. Consider the spanning subdigraph $D' := (V, A', \psi\mid_{A'})$ of D.

Let S be the set of all vertices $v \in V$ such that D' has a path from s to v. Then, $s \in S$ (since the trivial path (s) is a path of D').

We next claim that each arc $b \in [S, \overline{S}]$ satisfies $f(b) = 0$.

[*Proof:* Assume the contrary. Thus, some arc $b \in [S, \overline{S}]$ satisfies $f(b) \neq 0$. Consider this b. From $f(b) \neq 0$, we obtain $f(b) > 0$ (since f is a flow), thus $b \in A'$ (by the definition of A'). Hence, b is an arc of D' (by the definition of D').

Let u be the source of the arc b, and v its target. Since $b \in [S, \overline{S}]$, we therefore have $u \in S$ and $v \in \overline{S}$. Since $u \in S$, the digraph D' has a path $\mathbf{p}$ from s to u (by the definition of S). Consider this path $\mathbf{p}$. Appending the arc b and the vertex v at the end of this path $\mathbf{p}$, we obtain a walk from s to v in D' (since b is an arc of D' with source u and target v). Hence, the digraph D' has a walk from s to v, thus also a path from s to v (by Corollary 1.2.8 in Lecture 10). This means that $v \in S$ (by the definition of S). But this contradicts $v \in \overline{S} = V \setminus S$. This contradiction shows that our assumption was wrong, qed.]

We thus have proved that each $b \in \left[S, \overline{S}\right]$ satisfies $f(b) = 0$. Therefore, $f\left(S, \overline{S}\right) = 0$ (using the notations of Definition 1.1.5 **(d)** from Lecture 26). However, recall that $s \in S$. Thus, if we had $t \notin S$, then Proposition 1.2.2 **(b)** in Lecture 26 would yield

$$|f| = \underbrace{f\left(S, \overline{S}\right)}_{=0} - \underbrace{f\left(\overline{S}, S\right)}_{\geq 0} \leq 0 - 0 = 0,$$

which would contradict $|f| = m > 0$. Hence, we must have $t \in S$. In other words, the digraph D' has a path from s to t (by the definition of S). Let $\mathbf{p}$ be this path. Then, $\mathbf{p}$ is also a path in D and satisfies $f_{\mathbf{p}} \leq f$ [3]. Therefore, $f - f_{\mathbf{p}}$ is a map from A to $\mathbb{N}$. Moreover, $f - f_{\mathbf{p}}$ is again a flow[4], and has value $\left|f - f_{\mathbf{p}}\right| = m - 1$ [5]. Thus, by the induction hypothesis, we can apply Lemma 1.1.15 to $m-1$ and $f - f_{\mathbf{p}}$ instead of m and f. As a result, we conclude that there exist $m-1$ paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_{m-1}$ from s to t in D such that $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_{m-1}} \leq f - f_{\mathbf{p}}$. Consider these $m-1$ paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_{m-1}$, and set $\mathbf{p}_m := \mathbf{p}$. Then, $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_{m-1}} \leq f - f_{\mathbf{p}} = f - f_{\mathbf{p}_m}$ (since $\mathbf{p} = \mathbf{p}_m$), so that $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m} \leq f$.

Thus, we have found m paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ from s to t in D such that $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m} \leq f$. But this is precisely what we wanted. Thus, the induction step is complete, and Lemma 1.1.15 is proved. □

Remark 1.1.16. There exists an alternative proof of Lemma 1.1.15, which is too nice to leave unmentioned. Here is a quick outline: Consider a new multidigraph that is obtained from D by replacing each arc a by $f(a)$ many parallel arcs (if $f(a) = 0$, this means that a is simply removed). Add m many arcs from t to s to this new multidigraph. The resulting digraph is balanced (because of the conservation constraints for f). It may fail to be

[3] *Proof.* We need to prove that each arc $a \in A$ satisfies $f_{\mathbf{p}}(a) \leq f(a)$.

So let $a \in A$ be an arc. If a is not an arc of $\mathbf{p}$, then the definition of $f_{\mathbf{p}}$ yields $f_{\mathbf{p}}(a) = 0 \leq f(a)$ (since f is a flow), so we are done in this case. Hence, assume WLOG that a is an arc of $\mathbf{p}$. Thus, a is an arc of D' (since $\mathbf{p}$ is a path of D'). In other words, $a \in A'$. By the definition of A', this means that $f(a) > 0$. Since $f(a)$ is an integer, we thus have $f(a) \geq 1 = f_{\mathbf{p}}(a)$ (since a is an arc of $\mathbf{p}$). In other words, $f_{\mathbf{p}}(a) \leq f(a)$. This is precisely what we wanted to prove.

[4] Here, we are using the fact (which is straightforward to prove) that if g and h are two flows with $h \leq g$, then $g - h$ is again a flow.

[5] Here, we are using the fact (which is straightforward to prove) that if g and h are two flows satisfying $h \leq g$, then $|g - h| = |g| - |h|$. Applying this fact to $g = f$ and $h = f_{\mathbf{p}}$, we obtain $\left|f - f_{\mathbf{p}}\right| = \underbrace{|f|}_{=m} - \underbrace{\left|f_{\mathbf{p}}\right|}_{=1} = m - 1$.

weakly connected; however, the vertices s and t belong to the same weak component of it (as long as $m > 0$). Hence, applying the directed Euler-Hierholzer theorem (Theorem 1.4.2 **(a)** in Lecture 10) to this component, we see that this component has a Eulerian circuit. Cutting the m arcs from t to s out of this circuit, we obtain m arc-disjoint walks from s to t. Each of these m walks contains some path from s to t, and thus we obtain m paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ from s to t in D such that $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m} \leq f$.

Remark 1.1.17. Let N be a network consisting of a multidigraph $D = (V, A, \psi)$, a source $s \in V$, a sink $t \in V$ and a capacity function $c : A \to \mathbb{N}$. If $\mathbf{c}$ is a cycle of D, then we can define a map $f_{\mathbf{c}} : A \to \mathbb{N}$ by setting

$$f_{\mathbf{c}}(a) = \begin{cases} 1, & \text{if } a \text{ is an arc of } \mathbf{c}; \\ 0, & \text{otherwise} \end{cases} \qquad \text{for each } a \in A.$$

We call this map $f_{\mathbf{c}}$ the **cycle flow** of $\mathbf{c}$. It is an actual flow of value 0 if all the arcs of $\mathbf{c}$ have capacity ≥ 1.

Now, the conclusion of Lemma 1.1.15 can be improved as follows: There exist m paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ from s to t in D as well as a (possibly empty) collection of cycles $\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_k$ of D such that

$$f = \left(f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}\right) + \left(f_{\mathbf{c}_1} + f_{\mathbf{c}_2} + \cdots + f_{\mathbf{c}_k}\right).$$

Proving this improved claim is a bit harder than proving Lemma 1.1.15, but not by too much (in particular, the argument in Remark 1.1.16 can be adapted, since a walk becomes a path if we successively remove all cycles from it).

Proof of Theorem 1.1.10. We make D into a network N (with source s and sink t) by assigning the capacity 1 to each arc $a \in A$. Clearly, a cut of this network is the same as what we call an s-t-cut. Moreover, the capacity $c\left(S, \overline{S}\right)$ of a cut $\left[S, \overline{S}\right]$ is simply the size of this cut (since each arc has capacity 1).

The max-flow-min-cut theorem (Theorem 1.3.3 in Lecture 26) tells us that the maximum value of a flow equals the minimum capacity of a cut, i.e., the minimum size of an s-t-cut (because, as we just explained, a cut is the same as an s-t-cut, and its capacity is simply its size). It thus remains to show that the maximum value of a flow is the maximum number of pairwise arc-disjoint paths from s to t. But this is easy by now:

- If you have a flow f of value m, then you can find m pairwise arc-disjoint paths from s to t (because Lemma 1.1.15 gives you m paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ such that $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m} \leq f$, and the latter inequality tells you that these m paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ are arc-disjoint[6]). Thus,

$$\begin{aligned} & (\text{the maximum number of pairwise arc-disjoint paths from } s \text{ to } t) \\ & \geq (\text{the maximum value of a flow}) . \end{aligned} \tag{1}$$

- Conversely, if you have m pairwise arc-disjoint paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ from s to t, then you obtain a flow of value m (namely, $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}$ is

---

[6] *Proof.* Assume the contrary. Thus, these m paths are not arc-disjoint. In other words, there exists an arc a that is used by two paths $\mathbf{p}_i$ and $\mathbf{p}_j$ with $i \neq j$. Consider this arc a and the corresponding indices i and j. Since a is used by $\mathbf{p}_i$, we have $f_{\mathbf{p}_i}(a) = 1$. Likewise, $f_{\mathbf{p}_j}(a) = 1$. However, $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m} \leq f$, so that

$$\begin{aligned} \left(f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}\right)(a) \leq f(a) \leq c(a) \qquad & (\text{by the capacity constraints}) \\ = 1 \qquad & (\text{since each arc has capacity } 1) . \end{aligned}$$

Thus,

$$\begin{aligned} 1 & \geq \left(f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}\right)(a) = f_{\mathbf{p}_1}(a) + f_{\mathbf{p}_2}(a) + \cdots + f_{\mathbf{p}_m}(a) \\ & \geq \underbrace{f_{\mathbf{p}_i}(a)}_{=1} + \underbrace{f_{\mathbf{p}_j}(a)}_{=1} \qquad \left(\begin{array}{c} \text{since } f_{\mathbf{p}_i}(a) \text{ and } f_{\mathbf{p}_j}(a) \text{ are two distinct addends} \\ \text{of the sum } f_{\mathbf{p}_1}(a) + f_{\mathbf{p}_2}(a) + \cdots + f_{\mathbf{p}_m}(a) \\ \text{(because } i \neq j\text{), and since all the remaining} \\ \text{addends are } \geq 0 \text{ (since } f_{\mathbf{p}}(a) \geq 0 \text{ for each path } \mathbf{p}\text{)} \end{array}\right) \\ & = 1 + 1 > 1, \end{aligned}$$

which is absurd. This contradiction shows that our assumption was false, qed.

such a flow[7]). Thus,

$$\begin{aligned}&(\text{the maximum value of a flow})\\&\geq (\text{the maximum number of pairwise arc-disjoint paths from } s \text{ to } t)\,.\end{aligned}$$

Combining this last inequality with (1), we obtain

$$\begin{aligned}&(\text{the maximum number of pairwise arc-disjoint paths from } s \text{ to } t)\\&= (\text{the maximum value of a flow})\\&= (\text{the minimum size of an } s\text{-}t\text{-cut}) \qquad (\text{as we have proved before})\,.\end{aligned}$$

Thus, Theorem 1.1.10 is proved. □

Theorem 1.1.10 can also be proved without using network flows (see, e.g., [Schrij17, Corollary 4.1b] for such a proof).

Proof of Theorem 1.1.6. Let x denote the maximum number of pairwise arc-disjoint paths from s to t.

Let n_c denote the minimum size of an s-t-cut.

Let n_s denote the minimum size of an s-t-arc-separator.[8]

Theorem 1.1.10 says that $x = n_c$. Our goal is to prove that $x = n_s$.

---

[7]*Proof.* First, we observe that the map $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}$ satisfies the conservation constraints (because it is the sum of the functions $f_{\mathbf{p}_1}, f_{\mathbf{p}_2}, \ldots, f_{\mathbf{p}_m}$, each of which satisfies the conservation constraints). Let us now check that it satisfies the capacity constraints.

Indeed, let $a \in A$ be an arc. Then, a belongs to at most one of the m paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ (since these m paths are arc-disjoint). In other words, at most one of the m numbers $f_{\mathbf{p}_1}(a)$, $f_{\mathbf{p}_2}(a)$, $\ldots$, $f_{\mathbf{p}_m}(a)$ equals 1; all the remaining numbers equal 0. Hence, the sum $f_{\mathbf{p}_1}(a) + f_{\mathbf{p}_2}(a) + \cdots + f_{\mathbf{p}_m}(a)$ of these m numbers equals either 1 or 0; in either case, we thus have $f_{\mathbf{p}_1}(a) + f_{\mathbf{p}_2}(a) + \cdots + f_{\mathbf{p}_m}(a) \in \{0,1\}$. Now,

$$\left(f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}\right)(a) = f_{\mathbf{p}_1}(a) + f_{\mathbf{p}_2}(a) + \cdots + f_{\mathbf{p}_m}(a) \in \{0,1\}\,,$$

so that

$$0 \leq \left(f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}\right)(a) \leq 1 = c(a)$$

(since each arc has capacity 1). Since we have proved this for each arc $a \in A$, we thus have shown that the map $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}$ satisfies the capacity constraints. Hence, this map is a flow (since it also satisfies the conservation constraints).

It remains to show that the value of this flow is m. But this is easy: For any flows $g_1, g_2, \ldots, g_k$, we have $|g_1 + g_2 + \cdots + g_k| = |g_1| + |g_2| + \cdots + |g_k|$ (this is straightforward to see from the definition of value). Thus,

$$\left|f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}\right| = \left|f_{\mathbf{p}_1}\right| + \left|f_{\mathbf{p}_2}\right| + \cdots + \left|f_{\mathbf{p}_m}\right| = \sum_{k=1}^{m} \underbrace{\left|f_{\mathbf{p}_k}\right|}_{=1} = \sum_{k=1}^{m} 1 = m.$$

In other words, the value of the flow $f_{\mathbf{p}_1} + f_{\mathbf{p}_2} + \cdots + f_{\mathbf{p}_m}$ is m.

[8]If you are wondering why we chose the baroque notations "x", "n_c" and "n_s" for these three numbers: The letter "x" appears in "maximum", whereas the letter "n" appears in "minimum". The subscripts "c" and "s" should be reasonably clear.

Remark 1.1.9 shows that any s-t-cut is an s-t-arc-separator. Thus, $n_s \leq n_c$.

The inequality $x \leq n_s$ follows easily from the pigeonhole principle[9]. Combining this with $n_s \leq n_c = x$ (since $x = n_c$), we obtain $x = n_s$. Thus, Theorem 1.1.6 is proved. □

We can also extend the arc-Menger theorem to paths between different pairs of vertices:

Theorem 1.1.18 (arc-Menger theorem for directed graphs, multi-terminal version)**.** Let $D = (V, A, \psi)$ be a multidigraph, and let X and Y be two disjoint subsets of V.

A **path from** X **to** Y shall mean a path whose starting point belongs to X and whose ending point belongs to Y.

An X**-**Y**-cut** shall mean a subset of A that has the form $[S, \overline{S}]$, where S is a subset of V that satisfies $X \subseteq S$ and $Y \subseteq \overline{S}$.

Then, the maximum number of pairwise arc-disjoint paths from X to Y equals the minimum size of an X-Y-cut.

Example 1.1.19. Here is an example of a digraph $D = (V, A, \psi)$, with two disjoint subsets X and Y of V drawn as ovals:

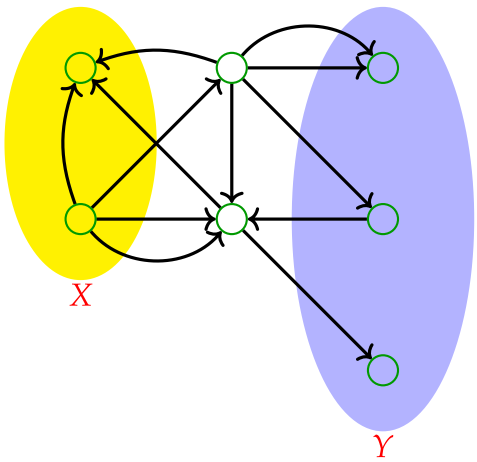

.

[9]*Proof.* We know that there exist x pairwise arc-disjoint paths from s to t (by the definition of x). Let $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_x$ be these x paths.

We know that there exists an s-t-arc-separator of size n_s (by the definition of n_s). Let B be this s-t-arc-separator. Thus, each path from s to t contains at least one arc from B (by the definition of an s-t-arc-separator). Hence, in particular, each of the x paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_x$ contains at least one arc from B. These altogether x arcs must be distinct (since the x paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_x$ are arc-disjoint); thus, we have found at least x arcs that belong to B. This shows that $|B| \geq x$. However, B has size n_s; in other words, we have $|B| = n_s$. Thus, $n_s = |B| \geq x$, so that $x \leq n_s$.

In this digraph D, the maximum number of pairwise arc-disjoint paths from X to Y is 2; here are two such paths (marked in red and blue):

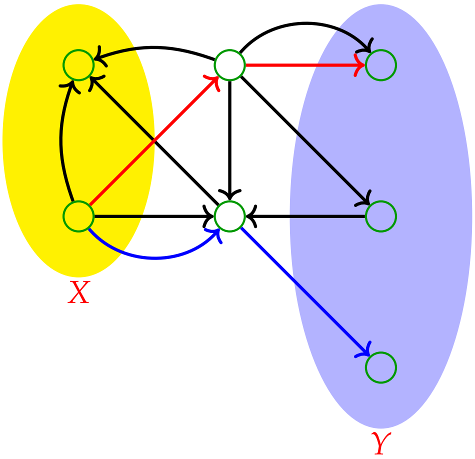

.

According to Theorem 1.1.18, the minimum size of an X-Y-cut must thus also be 2. And indeed, here is such an X-Y-cut:

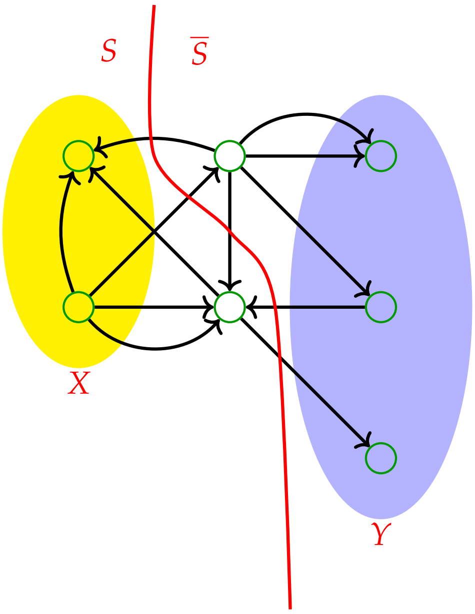

.

Proof of Theorem 1.1.18. We transform our digraph $D = (V, A, \psi)$ into a new multidigraph $D' = (V', A', \psi')$ as follows:

- We replace all the vertices in X by a single (new) vertex s, and replace all the vertices in Y by a single (new) vertex t. (Thus, formally speaking, we set $V' = (V \setminus (X \cup Y)) \cup \{s, t\}$, where s and t are two objects not in V.)

For any vertex $p \in V$, we define a vertex $p' \in V'$ by

$$p' = \begin{cases} s, & \text{if } p \in X; \\ t, & \text{if } p \in Y; \\ p, & \text{otherwise.} \end{cases}$$

We refer to this vertex p' as the **projection** of p.

- We keep all the arcs of D around, but we replace all their endpoints (i.e., sources and targets) by their projections (thus, any endpoint in X gets replaced by s, and any endpoint in Y gets replaced by t, while an endpoint that belongs neither to X nor to Y stays unchanged). For example, an arc with source in X becomes an arc with source in s. (Formally speaking, this means the following: We set $A' = A$ and we define the map $\psi' : A' \to V' \times V'$ as follows: For any $a \in A' = A$, we set $\psi'(a) = (u', v')$, where $(u, v) = \psi(a)$.)

For instance, if D is the digraph from Example 1.1.19, then D' looks as follows:

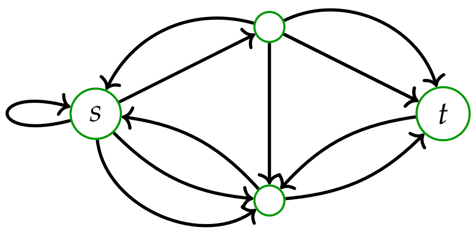

.

Now, Theorem 1.1.10 (applied to $D' = (V', A', \psi')$ instead of $D = (V, A, \psi)$) shows that the maximum number of pairwise arc-disjoint paths from s to t in D' equals the minimum size of an s-t-cut in D'.

Let us now connect this with the claim that we want to prove. It is easy to see that the minimum size of an s-t-cut in D' equals the minimum size of an X-Y-cut in D (indeed, the s-t-cuts in D' are precisely the X-Y-cuts in D [10]). If we can also show that the maximum number of pairwise arc-disjoint paths from s to t in D' equals the maximum number of pairwise arc-disjoint paths

[10] In more detail:

- Any s-t-cut in D' has the form $\left[S, \overline{S}\right]$ for some subset S of V' satisfying $s \in S$ and $t \notin S$; it is therefore equal to the set $\left[S', \overline{S'}\right]$, where S' is the subset of V given by $S' := (S \setminus \{s\}) \cup X$. Therefore, it is an X-Y-cut in D.
- Conversely, any X-Y-cut in D has the form $\left[S, \overline{S}\right]$ for some subset S of V satisfying $X \subseteq S$ and $Y \subseteq \overline{S}$; it is therefore equal to the set $\left[S', \overline{S'}\right]$, where S' is the subset of V' given by $S' := (S \setminus X) \cup \{s\}$. Therefore, it is an s-t-cut in D'.

from X to Y in D, then the result of the preceding paragraph will thus become the claim of Theorem 1.1.18, so we will be done.

So how can we show that the maximum number of pairwise arc-disjoint paths from s to t in D' equals the maximum number of pairwise arc-disjoint paths from X to Y in D ? It would be easy if there was a well-behaved bijection between the former paths and the latter paths that preserves the arcs of any path, but this is not quite the case. Each path from X to Y in D becomes a **walk** from s to t in D' if we replace each of its vertices by its projection. However, the latter walk is not necessarily a path, since different vertices can have the same projection.

Fortunately, this is easy to fix. If we have k pairwise arc-disjoint paths from X to Y in D, then we can turn them into k pairwise arc-disjoint **walks** from s to t in D', and then we also obtain k pairwise arc-disjoint **paths** from s to t in D' (since any walk from s to t contains a path from s to t). Thus,

$$\begin{aligned}&\left(\text{the maximum number of pairwise arc-disjoint paths from } s \text{ to } t \text{ in } D'\right)\\&\geq \left(\text{the maximum number of pairwise arc-disjoint paths from } X \text{ to } Y \text{ in } D\right).\end{aligned}$$

Conversely, if we have k pairwise arc-disjoint paths from s to t in D', then we can "lift" these k paths back to the digraph D (preserving the arcs, and replacing the vertices s and t by appropriate vertices in X and Y to make them belong to the right arcs), and thus obtain k pairwise arc-disjoint paths from X to Y in D. Thus,

$$\begin{aligned}&\left(\text{the maximum number of pairwise arc-disjoint paths from } X \text{ to } Y \text{ in } D\right)\\&\geq \left(\text{the maximum number of pairwise arc-disjoint paths from } s \text{ to } t \text{ in } D'\right).\end{aligned}$$

Combining these two inequalities, we obtain

$$\begin{aligned}&\left(\text{the maximum number of pairwise arc-disjoint paths from } s \text{ to } t \text{ in } D'\right)\\&= \left(\text{the maximum number of pairwise arc-disjoint paths from } X \text{ to } Y \text{ in } D\right).\end{aligned}$$

As explained above, this completes the proof of Theorem 1.1.18. $\square$

1.1.2. The edge-Menger theorem for undirected graphs

We shall now state analogues of Theorem 1.1.6 and Theorem 1.1.10 for undirected graphs. First, the unsurprising definitions:

Definition 1.1.20. Two walks $\mathbf{p}$ and $\mathbf{q}$ in a graph are said to be **edge-disjoint** if they have no edge in common.

Definition 1.1.21. Let $G = (V, E, \varphi)$ be a multigraph, and let s and t be two vertices of G. A subset B of E is said to be an **s-t-edge-separator** if each path from s to t contains at least one edge from B. Equivalently, a subset B of E is said to be an s-t-**edge-separator** if the multigraph $\left(V,\ E \setminus B,\ \varphi\mid_{E \setminus B}\right)$ has no path from s to t (in other words, removing from G all edges contained in B destroys all paths from s to t).

Now comes the analogue of Theorem 1.1.6:

Theorem 1.1.22 (edge-Menger theorem for undirected graphs, version 1)**.** Let $G = (V, E, \varphi)$ be a multigraph, and let s and t be two distinct vertices of G. Then, the maximum number of pairwise edge-disjoint paths from s to t equals the minimum size of an s-t-edge-separator.

To state the analogue of Theorem 1.1.10, we need to first adopt Definition 1.1.8 to undirected graphs:

Definition 1.1.23. Let $G = (V, E, \varphi)$ be a multigraph, and let s and t be two distinct vertices of G.

(a) For each subset S of V, we set $\overline{S} := V \setminus S$ and

$$\left[S, \overline{S}\right]_{\text{und}} := \{e \in E \mid \text{one endpoint of } e \text{ belongs to } S, \\ \text{while the other belongs to } \overline{S}\}.$$

(b) An **(undirected)** s-t-**cut** means a subset of E that has the form $\left[S, \overline{S}\right]_{\text{und}}$, where S is a subset of V that satisfies $s \in S$ and $t \notin S$.

The following remark is an analogue of Remark 1.1.9:

Remark 1.1.24. Let $G = (V, E, \varphi)$ be a multigraph, and let s and t be two distinct vertices of G. Then, any (undirected) s-t-cut is an s-t-edge-separator.

Proof. Analogous to the proof of Remark 1.1.9. □

And here is the analogue of Theorem 1.1.10:

Theorem 1.1.25 (edge-Menger theorem for undirected graphs, version 2)**.** Let $G = (V, E, \varphi)$ be a multigraph, and let s and t be two distinct vertices of G. Then, the maximum number of pairwise edge-disjoint paths from s to t equals the minimum size of an (undirected) s-t-cut.

Proof of Theorem 1.1.25. We shall not prove this from scratch, but rather derive this from the directed version (Theorem 1.1.10).

Namely, we apply Theorem 1.1.10 to[11] $D = G^{\text{bidir}}$. We thus see that the maximum number of pairwise arc-disjoint paths from s to t (in G^{bidir}) equals the minimum size of an s-t-cut (in G^{bidir}). This is similar to the claim that we want to prove, but not quite the same statement, because G^{bidir} is not G. To obtain the claim that we want to prove, we must prove the following two claims:

> *Claim 1:* The maximum number of pairwise arc-disjoint paths from s to t (in G^{bidir}) equals the maximum number of pairwise edge-disjoint paths from s to t (in G).
>
> *Claim 2:* The minimum size of a directed s-t-cut[12] (in G^{bidir}) equals the minimum size of an (undirected) s-t-cut (in G).

Claim 2 is very easy to verify, since the directed s-t-cuts in G^{bidir} are essentially the same as the undirected s-t-cuts in G [13].

It remains to verify Claim 1. The simplest approach is to argue that each path from s to t in G^{bidir} becomes a path from s to t in G (just replace each arc of the path by the corresponding undirected edge). Unfortunately, this alone does not suffice, since two arc-disjoint paths in G^{bidir} won't necessarily become edge-disjoint paths in G. Here is an example of how this can go wrong (imagine that the two arcs between u and v come from the same edge of G, and the two paths are marked red and blue):

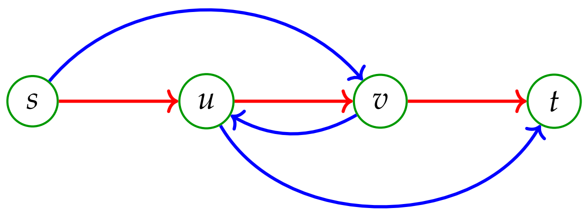

. (2)

If we replace each arc by the corresponding edge here, then the two paths will no longer be edge-disjoint (since the edge between u and v will be used by both paths).

---

[11] Recall that G^{bidir} is the multidigraph obtained from G by replacing each edge by two arcs in opposite directions. (If the edge has endpoints u and v, then one of the two arcs has source u and target v, while the other has source v and target u.) See Definition 1.1.1 in Lecture 10 for a formal definition.

[12] A "directed s-t-cut" here simply means an s-t-cut in a digraph.

[13] In more detail: If S is a subset of V that satisfies $s \in S$ and $t \notin S$, then the directed s-t-cut $[S, \overline{S}]$ in G^{bidir} and the undirected s-t-cut $[S, \overline{S}]_{\text{und}}$ in G have the same size (because each edge in $[S, \overline{S}]_{\text{und}}$ corresponds to exactly one arc in $[S, \overline{S}]$). Thus, the sizes of the directed s-t-cuts in G^{bidir} are exactly the sizes of the undirected s-t-cuts in G. In particular, the minimum size of a former cut equals the minimum size of a latter cut. This proves Claim 2.

However, this kind of situation can be averted. To do this, we let k be the maximum number of pairwise arc-disjoint paths from s to t in G^{bidir}. We now choose k pairwise arc-disjoint paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$ from s to t in G^{bidir} in such a way that their **total length** (i.e., the sum of the lengths of $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$) is **as small as possible**. Then, it is easy to see that these paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$ become edge-disjoint paths in G when we replace each arc by the corresponding edge.

[*Proof:* Assume the contrary. Thus, two of these paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$ end up sharing an edge when we replace each arc by the corresponding edge. Let $\mathbf{p}_i$ and $\mathbf{p}_j$ be these two paths (where $i \neq j$, of course). Let e be the edge that they end up sharing, and let u and v be the two endpoints of e, in the order in which they appear on $\mathbf{p}_i$. Hence, the path $\mathbf{p}_i$ uses the edge e (or, more precisely, one of the two arcs of G^{bidir} corresponding to e) to get from u to v.

Since the paths $\mathbf{p}_i$ and $\mathbf{p}_j$ are arc-disjoint, they cannot both use the edge e in the same direction (because this would mean that $\mathbf{p}_i$ and $\mathbf{p}_j$ share the same arc of G^{bidir}). Hence, the path $\mathbf{p}_j$ uses the edge e to get from v to u (since the path $\mathbf{p}_i$ uses the edge e to get from u to v). Hence, the paths $\mathbf{p}_i$ and $\mathbf{p}_j$ have the following forms:

$$\mathbf{p}_i = (\ldots, u, e_1, v, \ldots);$$
$$\mathbf{p}_j = (\ldots, v, e_2, u, \ldots),$$

where e_1 and e_2 are the two arcs of G^{bidir} that correspond to the edge e. Now, let us replace the two paths $\mathbf{p}_i$ and $\mathbf{p}_j$ by two new walks[14]

$$\mathbf{p}_i' = \left(\underbrace{\ldots}_{\substack{\text{the part of } \mathbf{p}_i \\ \text{before } u}}, u, \underbrace{\ldots}_{\substack{\text{the part of } \mathbf{p}_j \\ \text{after } u}} \right) \qquad \text{and}$$

$$\mathbf{p}_j' = \left(\underbrace{\ldots}_{\substack{\text{the part of } \mathbf{p}_j \\ \text{before } v}}, v, \underbrace{\ldots}_{\substack{\text{the part of } \mathbf{p}_i \\ \text{after } v}} \right).$$

These walks $\mathbf{p}_i'$ and $\mathbf{p}_j'$ are two walks from s to t, and they don't use any arcs that were not already used by $\mathbf{p}_i$ or $\mathbf{p}_j$. Thus, they are arc-disjoint from all of the paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$ except for $\mathbf{p}_i$ and $\mathbf{p}_j$. Moreover, they are arc-disjoint from each other (since $\mathbf{p}_i$ and $\mathbf{p}_j$ were arc-disjoint, and since the arcs of any path are distinct). Furthermore, their total length is smaller by 2 than the total length of $\mathbf{p}_i$ and $\mathbf{p}_j$ (since they use all the arcs of $\mathbf{p}_i$ and $\mathbf{p}_j$ except for e_1 and e_2). They are not necessarily paths, but we can turn them into paths from s to t by successively removing cycles (as in the proof of Corollary 1.2.8 in Lecture 10). If we do this, we end up with two paths $\mathbf{p}_i''$ and $\mathbf{p}_j''$ from s to t that are arc-disjoint from each other and from all of the paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$ except for $\mathbf{p}_i$ and $\mathbf{p}_j$, and whose total length is smaller by at least 2 than the total length of $\mathbf{p}_i$ and $\mathbf{p}_j$.

Thus, if we replace $\mathbf{p}_i$ and $\mathbf{p}_j$ by these two paths $\mathbf{p}_i''$ and $\mathbf{p}_j''$ (while leaving the remaining $k-2$ of our k paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$ unchanged), then we obtain k mutually

[14]Here is an illustration:

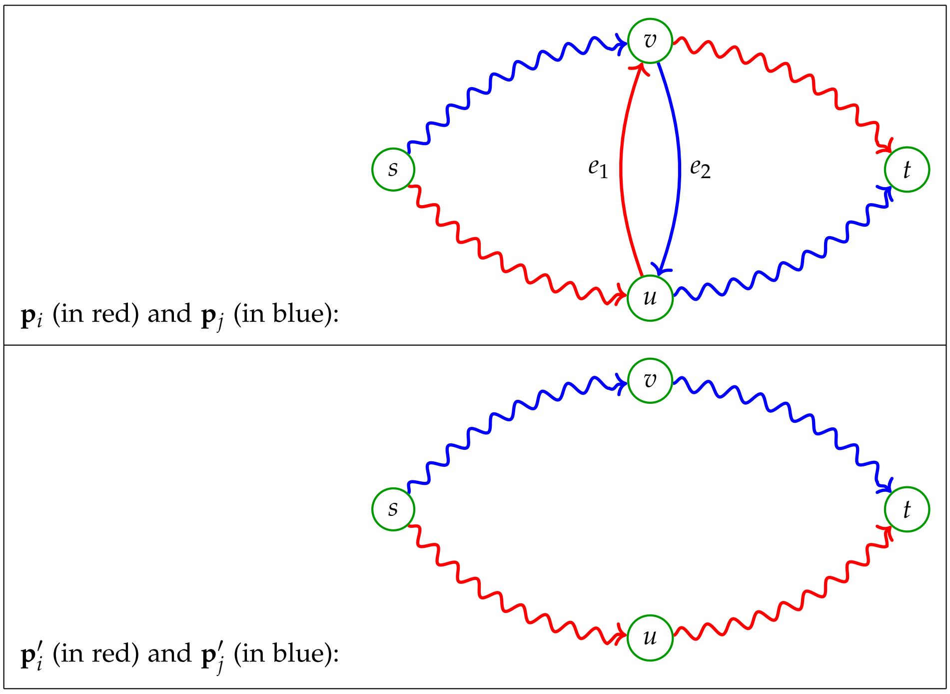

.

(The wavy arrows stand not for single arcs, but for sequences of multiple arcs.)

arc-disjoint paths from s to t whose total length is smaller than the total length of our original k paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$. However, this is absurd, because we chose our original k pairwise arc-disjoint paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$ from s to t in such a way that their total length is as small as possible. This contradiction shows that our assumption was wrong. Thus, we have proved that the paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$ become edge-disjoint paths in G when we replace each arc by the corresponding edge.]

Hence, we have found k pairwise edge-disjoint paths from s to t in G (namely, the k paths that are obtained from the paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k$ when we replace each arc by the corresponding edge). This shows that

$$\begin{aligned}
&(\text{the maximum number of pairwise edge-disjoint paths from } s \text{ to } t \text{ in } G) \\
&\geq k \\
&= \left(\text{the maximum number of pairwise arc-disjoint paths from } s \text{ to } t \text{ in } G^{\text{bidir}}\right)
\end{aligned}$$

(by the definition of k). Conversely, we can easily see that

$$\begin{aligned}
&\left(\text{the maximum number of pairwise arc-disjoint paths from } s \text{ to } t \text{ in } G^{\text{bidir}}\right) \\
&\geq (\text{the maximum number of pairwise edge-disjoint paths from } s \text{ to } t \text{ in } G)
\end{aligned}$$

(since there is an obvious way to transform paths in G into paths in G^{bidir} (just replace each edge by one of the two corresponding arcs of G^{bidir}), and applying this transformation to edge-disjoint paths of G yields arc-disjoint paths of G^{bidir}). Combining these two inequalities, we obtain

$$\begin{aligned}
&\left(\text{the maximum number of pairwise arc-disjoint paths from } s \text{ to } t \text{ in } G^{\text{bidir}}\right) \\
&= (\text{the maximum number of pairwise edge-disjoint paths from } s \text{ to } t \text{ in } G).
\end{aligned}$$

This proves Claim 1. As we explained, this concludes the proof of Theorem 1.1.25. $\square$

Proof of Theorem 1.1.22. This can be derived from Theorem 1.1.25 and Remark 1.1.24 in the same way as we derived Theorem 1.1.6 from Theorem 1.1.10 and Remark 1.1.9. $\square$

1.1.3. The vertex-Menger theorem for directed graphs

The Menger theorems we have seen so far have been concerned with paths not having arcs in common. What if we want to avoid common vertices too?

Definition 1.1.26. Let $\mathbf{p}$ be a path of some graph or digraph. Then, an **intermediate vertex** of $\mathbf{p}$ shall mean a vertex of $\mathbf{p}$ that is neither the starting point nor the ending point of $\mathbf{p}$.

Definition 1.1.27. Two paths $\mathbf{p}$ and $\mathbf{q}$ in a graph or digraph are said to be **internally vertex-disjoint** if they have no common intermediate vertices.

Example 1.1.28. The two paths $\mathbf{p}$ and $\mathbf{q}$ in Example 1.1.2 are arc-disjoint, but not internally vertex-disjoint.

Here are two internally vertex-disjoint paths $\mathbf{p}$ and $\mathbf{q}$:

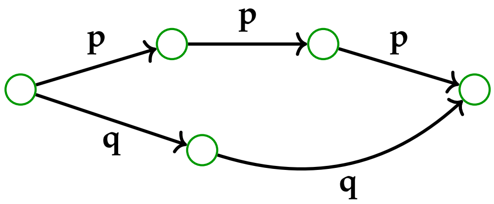

.

One trivial case of internally vertex-disjoint paths is a path of length ≤ 1: Namely, a path of length ≤ 1 is internally vertex-disjoint from any path, including itself (since it has no intermediate vertices).

Definition 1.1.29. Let $D = (V, A, \psi)$ be a multidigraph, and let s and t be two vertices of D. A subset W of $V \setminus \{s, t\}$ is said to be an **internal s-t-vertex-separator** if each path from s to t contains at least one vertex from W. Equivalently, a subset W of $V \setminus \{s, t\}$ is said to be an **internal s-t-vertex-separator** if the induced subdigraph of D on the set $V \setminus W$ has no path from s to t (in other words, removing from D all vertices contained in W destroys all paths from s to t).

Example 1.1.30. Let $D = (V, A, \psi)$ be the following multidigraph:

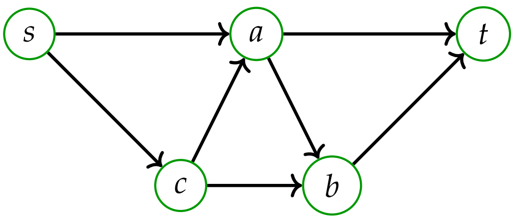

.

Then, the sets $\{a, b\}$ and $\{a, c\}$ are internal s-t-vertex-separators (indeed, removing the vertices a and b cuts off s from the rest of the digraph, whereas removing the vertices a and c does the same to t), but the sets $\{a\}$ and $\{b, c\}$ are not (since the path from s to t via c and b avoids a, whereas the path from s to t via a avoids b and c).

Example 1.1.31. Let $D = (V, A, \psi)$ be a multidigraph. Let s and t be two distinct vertices of D. Then:

(a) The empty set $\varnothing$ is an internal s-t-vertex-separator if and only if D has no path from s to t.

(b) If D has no arc with source s and target t, then the set $V \setminus \{s, t\}$ is an internal s-t-vertex-separator (since any path from s to t contains at least one intermediate vertex, and such a vertex must belong to $V \setminus \{s, t\}$).

(c) If D has an arc with source s and target t, then there exists no internal s-t-vertex-separator (since the "direct" length-1 path from s to t contains no vertices besides s and t).

Now, we state the analogue of Theorem 1.1.6 and Theorem 1.1.10 for internally vertex-disjoint paths:

Theorem 1.1.32 (vertex-Menger theorem for directed graphs)**.** Let $D = (V, A, \psi)$ be a multidigraph, and let s and t be two distinct vertices of D. Assume that D has no arc with source s and target t. Then, the maximum number of pairwise internally vertex-disjoint paths from s to t equals the minimum size of an internal s-t-vertex-separator.

Example 1.1.33. Let D be the following multidigraph:

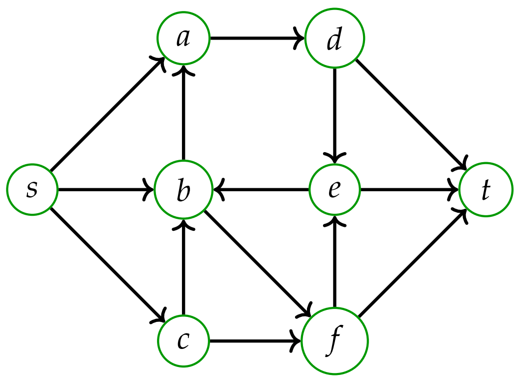

.

Then, the maximum number of pairwise internally vertex-disjoint paths from s to t is 2. Indeed, the following picture shows 2 such paths in red and blue, respectively:

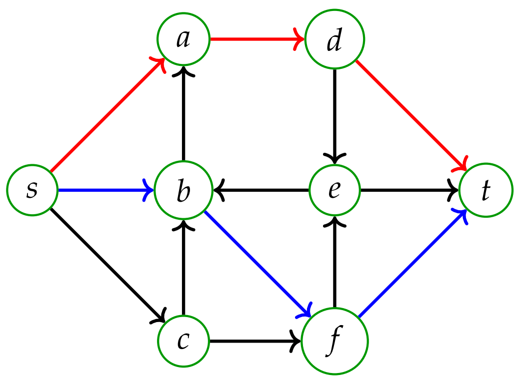

.

Why can there be no 3 such paths? This is not obvious from a quick look, but can be easily derived from Theorem 1.1.32. Indeed, Theorem 1.1.32 yields that the maximum number of pairwise internally vertex-disjoint paths from s to t equals the minimum size of an internal s-t-vertex-separator. Thus, if the former number was larger than 2, then so would be the latter number. But this cannot be the case, since the 2-element set $\{a, f\}$ is easily checked to be an internal s-t-vertex-separator. Hence, we see that both of these numbers are 2.

Example 1.1.34. Consider again the digraph D from Example 1.1.11. In that example, we found 3 pairwise arc-disjoint paths from s to t. These 3 paths are not internally vertex-disjoint (in fact, the brown path has non-starting and non-ending vertices in common with both the red and the blue path). However, there do exist 3 pairwise internally vertex-disjoint paths from s to t. Can you find them?

Proof of Theorem 1.1.32. We will again derive this from the arc-Menger theorem (Theorem 1.1.10), applied to an appropriate multidigraph $D' = (V', A', \psi')$.

What is this multidigraph D' ? The idea is to modify the digraph D in such a way that paths having a common vertex become paths having a common arc. The most natural way to achieve this is to "stretch out" each vertex v of D into a little arc. In order to do this in a systematic manner, we replace each vertex v of D by two distinct vertices v^i and v^o (the notations stand for "v-in" and "v-out", and we can think of v^i as the "entrance" to v while v^o is the "exit" from v) and an arc v^{io} that goes from v^i to v^o. Any existing arc a of D becomes a new arc a^{oi} of D, whose source and target are specified as follows: If a has source u and target v, then a^{oi} will have source u^o and target v^i.

Here is an example: If

$D =$ 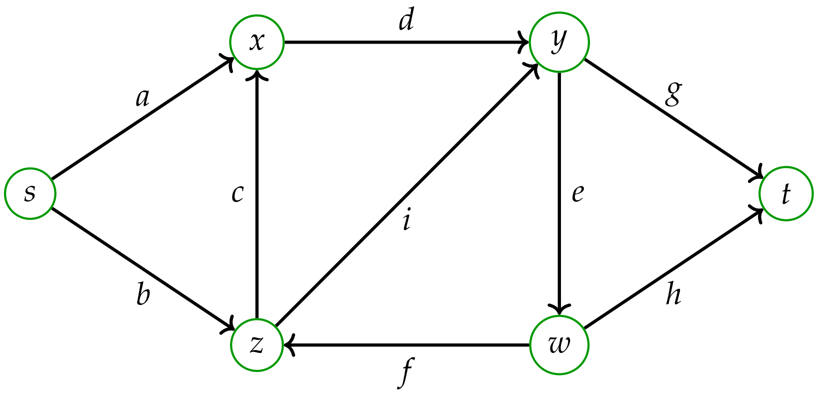

,

then

$D' =$ 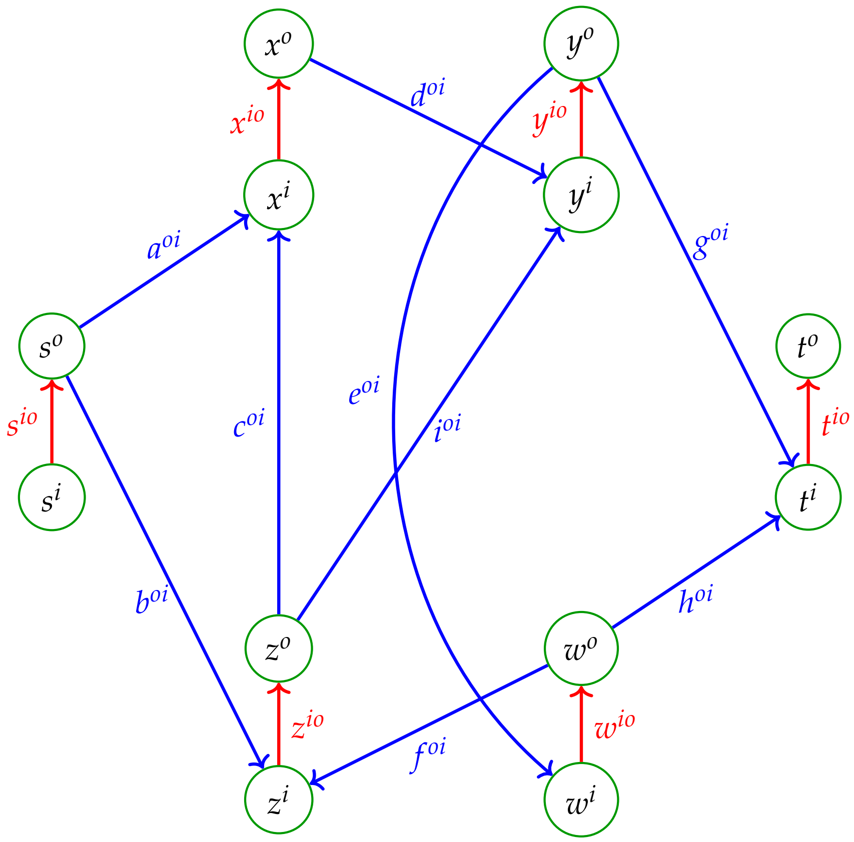

(where all arcs of the form a^{oi} for $a \in A$ have been colored blue, whereas all arcs of the form v^{io} for $v \in V$ have been colored red). This D' satisfies the property that we want it to satisfy: For instance, the two paths

$$(s, a, x, d, y, g, t) \qquad \text{and}$$
$$(s, b, z, i, y, e, w, h, t)$$

of D have the vertex y in common, so the corresponding two paths

$$\left(s^o, a^{oi}, x^i, x^{io}, x^o, d^{oi}, y^i, y^{io}, y^o, g^{oi}, t^i\right) \qquad \text{and}$$
$$\left(s^o, b^{oi}, z^i, z^{io}, z^o, i^{oi}, y^i, y^{io}, y^o, e^{oi}, w^i, w^{io}, w^o, h^{oi}, t^i\right)$$

of D' have the arc y^{io} in common. If you think of D as a railway network with the vertices being train stations and the arcs being train rides, then D' is a more detailed version of this network that records a change of platforms as an arc as well.

Here is a formal definition of the multidigraph $D' = (V', A', \psi')$ in full generality:

- We replace each vertex v of D by two new vertices v^i and v^o. We call v^i an "**in-vertex**" and v^o an "**out-vertex**". The vertex set of D' will be the set
$$V' := \underbrace{\left\{v^i \mid v \in V\right\}}_{\text{in-vertices}} \cup \underbrace{\left\{v^o \mid v \in V\right\}}_{\text{out-vertices}}.$$

- Each arc $a \in A$ is replaced by a new arc a^{oi}, which is defined as follows: If the arc $a \in A$ has source u and target v, then we replace it by a new arc a^{oi}, which has source u^o and target v^i. This arc a^{oi} will be called an "**arc-arc**" of D' (since it originates from an arc of D).

- For any vertex $v \in V$ of D, we introduce a new arc v^{io}, which has source v^i and target v^o. This arc v^{io} will be called a "**vertex-arc**" of D' (since it originates from a vertex of D).

- The arc set of D' will be the set
$$A' := \underbrace{\left\{a^{oi} \mid a \in A\right\}}_{\text{the arc-arcs}} \cup \underbrace{\left\{v^{io} \mid v \in V\right\}}_{\text{the vertex-arcs}}.$$

  The map $\psi' : A' \to V' \times V'$ is defined as we already explained:

  - For any arc-arc $a^{oi} \in A$, we let $\psi'\left(a^{oi}\right) := \left(u^o, v^i\right)$, where $(u, v) = \psi(a)$.
  - For any vertex-arc $v^{io} \in V$, we let $\psi'\left(v^{io}\right) := \left(v^i, v^o\right)$.

Note that D' is something like a "bipartite digraph": Each of its arcs goes either from an out-vertex to an in-vertex or vice versa. Namely, each arc-arc goes from an out-vertex to an in-vertex, whereas each vertex-arc goes from an in-vertex to an out-vertex. Thus, on any walk of D', the arc-arcs and the vertex-arcs have to alternate.

If $\mathbf{p} = (v_0, a_1, v_1, a_2, v_2, \ldots, a_k, v_k)$ is any nontrivial[15] path of D, then we can define a corresponding path $\mathbf{p}^{oi}$ of D' by
$$\mathbf{p}^{oi} := \left(v_0^o, a_1^{oi}, v_1^i, v_1^{io}, v_1^o, a_2^{oi}, v_2^i, v_2^{io}, v_2^o, \ldots, a_k^{oi}, v_k^i\right).$$

This path $\mathbf{p}^{oi}$ is obtained from $\mathbf{p}$ by

- replacing the starting point v_0 by v_0^o;
- replacing the ending point v_k by v_k^i;
- replacing each other vertex v_j by the sequence v_j^i, v_j^{io}, v_j^o;

[15] We say that a path is **nontrivial** if it has length > 0.

- replacing each arc a_j by a_j^{oi}.

Informally speaking, this simply means that we stretch out each intermediate vertex of $\mathbf{p}$ to the corresponding arc.

If $\mathbf{p}$ is a path from s to t in D, then $\mathbf{p}^{oi}$ is a path from s^o to t^i in D'. Conversely, any path from s^o to t^i in D' must have the form $\mathbf{p}^{oi}$, where $\mathbf{p}$ is some path from s to t in D (because on any walk of D', the arc-arcs and the vertex-arcs have to alternate). Therefore, the map

$$\{\text{paths from } s \text{ to } t \text{ in } D\} \to \left\{\text{paths from } s^o \text{ to } t^i \text{ in } D'\right\},$$
$$\mathbf{p} \mapsto \mathbf{p}^{oi} \tag{3}$$

is a bijection. Moreover, two paths $\mathbf{p}$ and $\mathbf{q}$ of D are internally vertex-disjoint if and only if the paths $\mathbf{p}^{oi}$ and $\mathbf{q}^{oi}$ are arc-disjoint (since each vertex of a path $\mathbf{p}$ except for its starting and ending points is represented by an arc in $\mathbf{p}^{oi}$).

Now, let k be the maximum number of pairwise arc-disjoint paths from s^o to t^i in D'. Thus, D' has k pairwise arc-disjoint paths from s^o to t^i. Applying the inverse of the bijection (3) to these k paths, we obtain k pairwise internally vertex-disjoint paths from s to t in D (because two paths $\mathbf{p}$ and $\mathbf{q}$ of D are internally vertex-disjoint if and only if the paths $\mathbf{p}^{oi}$ and $\mathbf{q}^{oi}$ are arc-disjoint). Hence,

$$\begin{aligned}&(\text{the maximum number of pairwise internally vertex-disjoint}\\ &\qquad\qquad \text{paths from } s \text{ to } t \text{ in } D)\\ &\geq k. \end{aligned}\tag{4}$$

Our next goal is to find an internal s-t-vertex-separator $W \subseteq V \setminus \{s, t\}$ of size $|W| \leq k$.

First, we simplify our setting a bit.

A path from s to t cannot contain a loop; nor can it contain an arc with source t and target s (since the vertices of a path must be distinct). Hence, we can remove such arcs (i.e., loops as well as arcs with source t and target s) from D without affecting the meaning of the claim we are proving. Thus, we WLOG assume that the digraph D has no such arcs. Since we also know (by assumption) that D has no arc with source s and target t, we thus conclude that D has no arc with source $\in \{s, t\}$ and target $\in \{s, t\}$ (because each such arc would either have source s and target t, or have source t and target s, or be a loop). In other words, each arc of D has at least one endpoint[16] distinct from both s and t.

However, k is the maximum number of pairwise arc-disjoint paths from s^o to t^i in D'. Therefore, by Theorem 1.1.10 (applied to $D' = (V', A', \psi')$, s^o and t^i instead of $D = (V, A, \psi)$, s and t), this number k equals the minimum size of an

[16] An **endpoint** of an arc means a vertex that is either the source or the target of this arc.

s^o-t^i-cut in D'. Hence, there exists an s^o-t^i-cut $\left[S, \overline{S}\right]$ in D' such that $\left|\left[S, \overline{S}\right]\right| = k$. Consider this s^o-t^i-cut $\left[S, \overline{S}\right]$. Since $\left[S, \overline{S}\right]$ is an s^o-t^i-cut, we have $S \subseteq V'$ and $s^o \in S$ and $t^i \notin S$.

Let $B := \left[S, \overline{S}\right]$. Then, $|B| = \left|\left[S, \overline{S}\right]\right| = k$. Moreover, it is easy to see that $s^{io} \notin B$ [17] and $t^{io} \notin B$ [18].

To each vertex $w \in V'$ of D', we assign a vertex $\beta(w) \in V$ of D as follows: If $w = v^i$ or $w = v^o$ for some $v \in V$, then we set $\beta(w) := v$. In other words, $\beta(w)$ is the vertex v such that $w \in \left\{v^i, v^o\right\}$. We shall call $\beta(w)$ the **base** of the vertex w.

For each arc $b \in B$, there exists at least one endpoint w of b such that $\beta(w) \in V \setminus \{s, t\}$ [19]. We choose such an endpoint w arbitrarily, and we denote its base $\beta(w)$ by $\beta(b)$. We shall call $\beta(b)$ the **basepoint** of the arc b. Thus, by definition, we have

$$\beta(b) \in V \setminus \{s, t\} \qquad \text{for each } b \in B. \tag{5}$$

We let $\beta(B)$ denote the set $\{\beta(b) \mid b \in B\}$. Clearly, $|\beta(B)| \leq |B| = k$.

[17] *Proof.* If we had $s^{io} \in \left[S, \overline{S}\right]$, then we would have $s^i \in S$ and $s^o \in \overline{S}$; however, $s^o \in \overline{S}$ would contradict $s^o \in S$. Thus, we cannot have $s^{io} \in \left[S, \overline{S}\right]$. In other words, we cannot have $s^{io} \in B$ (since $B = \left[S, \overline{S}\right]$). Hence, $s^{io} \notin B$.

[18] *Proof.* If we had $t^{io} \in \left[S, \overline{S}\right]$, then we would have $t^i \in S$ and $t^o \in \overline{S}$; however, $t^i \in S$ would contradict $t^i \in \overline{S}$. Thus, we cannot have $t^{io} \in \left[S, \overline{S}\right]$. In other words, we cannot have $t^{io} \in B$ (since $B = \left[S, \overline{S}\right]$). Hence, $t^{io} \notin B$.

[19] *Proof:* Let $b \in B$ be an arc. We must prove that there exists at least one endpoint w of b such that $\beta(w) \in V \setminus \{s, t\}$.

The arc b is either a vertex-arc or an arc-arc. Thus, we are in one of the following two cases:

Case 1: The arc b is a vertex-arc.

Case 2: The arc b is an arc-arc.

Let us first consider Case 1. In this case, the arc b is a vertex-arc. In other words, $b = v^{io}$ for some $v \in V$. Consider this v. Then, $v^{io} = b \in B$. Hence, $v \neq s$ (since $v = s$ would entail $v^{io} = s^{io} \notin B$, which would contradict $v^{io} \in B$) and $v \neq t$ (since $v = t$ would entail $v^{io} = t^{io} \notin B$, which would contradict $v^{io} \in B$). Therefore, $v \in V \setminus \{s, t\}$. Also, clearly, v^i is an endpoint of b and satisfies $\beta\left(v^i\right) = v \in V \setminus \{s, t\}$. Hence, there exists at least one endpoint w of b such that $\beta(w) \in V \setminus \{s, t\}$ (namely, $w = v^i$). Thus, our proof is complete in Case 1.

Let us now consider Case 2. In this case, the arc b is an arc-arc. In other words, $b = a^{oi}$ for some $a \in A$. Consider this a. Now, a is an arc of D (since $a \in A$), and thus has at least one endpoint distinct from both s and t (since we have shown above that each arc of D has at least one endpoint distinct from both s and t). Let x be this endpoint. Then, $x \in V \setminus \{s, t\}$ (since x is distinct from both s and t).

But x is an endpoint of a. In other words, x is either the source or the target of a. Hence, the arc a^{oi} of D' either has source x^o or has target x^i (by the definition of a^{oi}). In other words, the arc b of D' either has source x^o or has target x^i (since $b = a^{oi}$). Since $\beta(x^o) = x \in V \setminus \{s, t\}$ and $\beta\left(x^i\right) = x \in V \setminus \{s, t\}$, we thus conclude that the arc b of D' has at least one endpoint w such that $\beta(w) \in V \setminus \{s, t\}$ (namely, $w = x^o$ if b has source x^o, and $w = x^i$ if b has target x^i). This completes our proof in Case 2.

Thus, we are done in both Cases 1 and 2, so that our proof is complete.

Now, we claim that

$$\text{every path from } s \text{ to } t \text{ (in } D\text{) contains a vertex in } \beta(B). \tag{6}$$

[*Proof of (6):* Let $\mathbf{p}$ be a path from s to t (in D). We must prove that $\mathbf{p}$ contains a vertex in $\beta(B)$.

Recall that we have assigned a path $\mathbf{p}^{oi}$ of D' to the path $\mathbf{p}$ of D. The definition of $\mathbf{p}^{oi}$ shows that the base of any vertex of $\mathbf{p}^{oi}$ is a vertex of $\mathbf{p}$ (indeed, if $v_0, v_1, \ldots, v_k$ are the vertices of $\mathbf{p}$, then the vertices of $\mathbf{p}^{oi}$ are $v_0^o, v_1^i, v_1^o, v_2^i, v_2^o, \ldots, v_{k-1}^i, v_{k-1}^o, v_k^i$, and their respective bases are $v_0, v_1, v_1, v_2, v_2, \ldots, v_{k-1}, v_{k-1}, v_k$).

The path $\mathbf{p}^{oi}$ is a path from s^o to t^i (since $\mathbf{p}$ is a path from s to t). Hence, it starts at a vertex in S (since $s^o \in S$) and ends at a vertex that is not in S (since $t^i \notin S$). Thus, this path $\mathbf{p}^{oi}$ must cross from S into $\overline{S}$ somewhere. In other words, there exists an arc b of $\mathbf{p}^{oi}$ such that the source of b belongs to S but the target of b belongs to $\overline{S}$. Consider this arc b. Thus, $b \in [S, \overline{S}] = B$, so that $\beta(b) \in \beta(B)$ (by the definition of $\beta(B)$). Both endpoints of b are vertices of $\mathbf{p}^{oi}$ (since b is an arc of $\mathbf{p}^{oi}$).

Now, consider the basepoint $\beta(b)$ of this arc b. This basepoint $\beta(b)$ is the base of an endpoint of b (by the definition of $\beta(b)$). Thus, $\beta(b)$ is the base of a vertex of $\mathbf{p}^{oi}$ (since both endpoints of b are vertices of $\mathbf{p}^{oi}$). Hence, $\beta(b)$ is a vertex of $\mathbf{p}$ (since the base of any vertex of $\mathbf{p}^{oi}$ is a vertex of $\mathbf{p}$). In other words, the path $\mathbf{p}$ contains the vertex $\beta(b)$. Since $\beta(b) \in \beta(B)$, we thus conclude that $\mathbf{p}$ contains a vertex in $\beta(B)$. This proves (6).]

Now, the set $\beta(B)$ is a subset of $V \setminus \{s, t\}$ (since $\beta(b) \in V \setminus \{s, t\}$ for each $b \in B$) and has the property that every path from s to t contains a vertex in $\beta(B)$ (by (6)). In other words, $\beta(B)$ is a subset $W \subseteq V \setminus \{s, t\}$ such that every path from s to t contains a vertex in W. In other words, $\beta(B)$ is an internal s-t-vertex-separator (by the definition of an "internal s-t-vertex-separator"). Thus,

$$\begin{aligned} &(\text{the minimum size of an internal } s\text{-}t\text{-vertex-separator}) \\ &\leq |\beta(B)| = k \\ &\leq (\text{the maximum number of pairwise internally vertex-disjoint} \\ &\qquad\qquad \text{paths from } s \text{ to } t \text{ in } D) \end{aligned}$$

(by (4)).

On the other hand, we have

$$\begin{aligned} &(\text{the minimum size of an internal } s\text{-}t\text{-vertex-separator}) \\ &\geq (\text{the maximum number of pairwise internally vertex-disjoint} \\ &\qquad\qquad \text{paths from } s \text{ to } t \text{ in } D) \end{aligned}$$

(by the pigeonhole principle[20]). Combining this inequality with the preceding

[20] *Proof* in more detail*:* Let n be the minimum size of an internal s-t-vertex-separator. Let x be the maximum number of pairwise internally vertex-disjoint paths from s to t in D. We must show that $n \geq x$.

one, we obtain

$$\begin{aligned}
&(\text{the minimum size of an internal } s\text{-}t\text{-vertex-separator}) \\
&= (\text{the maximum number of pairwise internally vertex-disjoint} \\
&\qquad\qquad \text{paths from } s \text{ to } t \text{ in } D) .
\end{aligned}$$

This proves Theorem 1.1.32. □

There is also a variant of the vertex-Menger theorem similar to what Theorem 1.1.18 did for the arc-Menger theorem. Again, we need some notations first:

Definition 1.1.35. Let $D = (V, A, \psi)$ be a multidigraph, and let X and Y be two subsets of V.

(a) A **path from** X **to** Y shall mean a path whose starting point belongs to X and whose ending point belongs to Y.

(b) A subset W of V is said to be an X**-**Y**-vertex-separator** if each path from X to Y contains at least one vertex from W. Equivalently, a subset W of V is said to be an X**-**Y**-vertex-separator** if the induced subdigraph of D on the set $V \setminus W$ has no path from X to Y (in other words, removing from D all vertices contained in W destroys all paths from X to Y).

(c) An X-Y-vertex-separator W is said to be **internal** if it is a subset of $V \setminus (X \cup Y)$ (that is, if it is disjoint from X and from Y).

Assume the contrary. Thus, $n < x$.

The definition of n shows that there exists an internal s-t-vertex-separator W that has size n.

The set W is an internal s-t-vertex-separator. In other words, W is a subset of $V \setminus \{s, t\}$ such that every path from s to t contains a vertex in W. Moreover, W has size n; thus, $|W| = n < x$.

The definition of x shows that there exist x pairwise internally vertex-disjoint paths from s to t in D. Let $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_x$ be these x paths. Each of these x paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_x$ must contain at least one vertex in W (since every path from s to t contains a vertex in W). Since $|W| < x$, we thus conclude by the pigeonhole principle that at least two of the x paths $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_x$ must contain the same vertex in W. In other words, there exist two distinct elements $i, j \in \{1, 2, \ldots, x\}$ such that $\mathbf{p}_i$ and $\mathbf{p}_j$ contain the same vertex in W. Let w be the latter vertex. Thus, $w \in W \subseteq V \setminus \{s, t\}$. Hence, w is distinct from both s and t. Therefore, w is an intermediate vertex of $\mathbf{p}_i$ (since the path $\mathbf{p}_i$ has starting point s and ending point t). Likewise, w is an intermediate vertex of $\mathbf{p}_j$.

However, the paths $\mathbf{p}_i$ and $\mathbf{p}_j$ are internally vertex-disjoint, and thus have no common intermediate vertex. This contradicts the fact that w is an intermediate vertex of both paths $\mathbf{p}_i$ and $\mathbf{p}_j$. This contradiction shows that our assumption was false. Hence, $n \geq x$ is proved, qed.

Theorem 1.1.36 (vertex-Menger theorem for directed graphs, multi-terminal version 1)**.** Let $D = (V, A, \psi)$ be a multidigraph, and let X and Y be two disjoint subsets of V. Assume that D has no arc with source in X and target in Y.

Then, the maximum number of pairwise internally vertex-disjoint paths from X to Y equals the minimum size of an internal X-Y-vertex-separator.

Example 1.1.37. Let D be the following multidigraph:

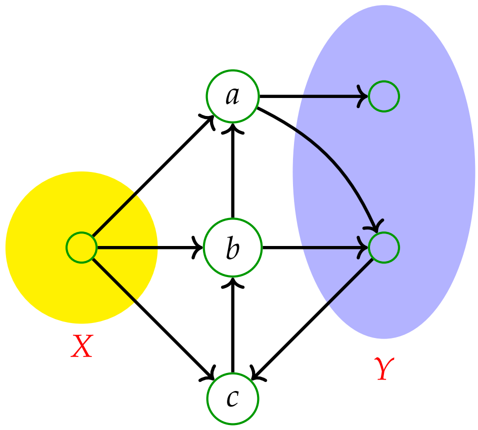

.

Then, the maximum number of pairwise internally vertex-disjoint paths from X to Y is 2; here are two such paths (drawn in red and blue):

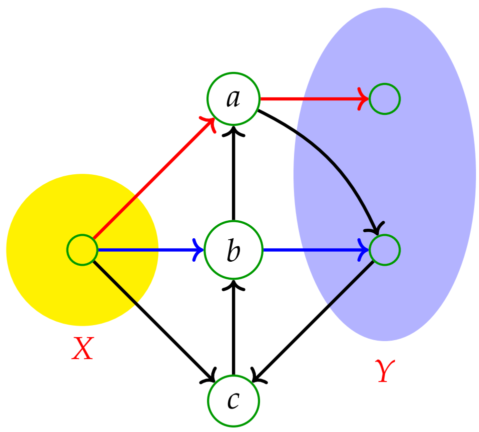

(there are other choices, of course). The minimum size of an internal X-Y-vertex-separator is 2 as well; indeed, $\{a, b\}$ is such an internal X-Y-vertex-separator. These two numbers are equal, just as Theorem 1.1.36 predicts.

Proof of Theorem 1.1.36. We define a new multidigraph $D' = (V', A', \psi')$ as in the proof of Theorem 1.1.18. Then, D' has no arc with source s and target t (since D has no arc with source in X and target in Y).

Hence, Theorem 1.1.32 (applied to $D' = (V', A', \psi')$ instead of $D = (V, A, \psi)$) shows that the maximum number of pairwise internally vertex-disjoint paths from s to t in D' equals the minimum size of an internal s-t-vertex-separator in D'.

Let us now see what this result means for our original digraph D. Indeed:

- The minimum size of an internal s-t-vertex-separator in D' equals the minimum size of an internal X-Y-vertex-separator in D (indeed, the internal s-t-vertex-separators in D' are precisely the internal X-Y-vertex-separators in D [21]).
- The maximum number of pairwise internally vertex-disjoint paths from s to t in D' equals the maximum number of pairwise internally vertex-disjoint paths from X to Y in D [22].

[21] *Proof.* We recall the definitions of internal s-t-vertex-separators in D' and of internal X-Y-vertex-separators in D:

- An internal s-t-vertex-separator in D' is a subset W of $V' \setminus \{s,t\}$ such that each path from s to t contains at least one vertex from W.
- An internal X-Y-vertex-separator in D is a subset W of $V \setminus (X \cup Y)$ such that each path from X to Y contains at least one vertex from W.

These two definitions describe the same object, because of the following two reasons:

- We have $V' \setminus \{s,t\} = V \setminus (X \cup Y)$.
- The paths from s to t are in bijection with the paths from X to Y (indeed, any path of the latter kind can be transformed into a path of the former kind by replacing the starting point by s and replacing the ending point by t). This bijection preserves the intermediate vertices (i.e., the vertices other than the starting point and the ending point). Thus, a path $\mathbf{p}$ from s to t contains at least one vertex from W if and only if the corresponding path from X to Y (that is, the image of $\mathbf{p}$ under our bijection) contains at least one vertex from W.

Thus, the internal s-t-vertex-separators in D' are precisely the internal X-Y-vertex-separators in D.

[22] *Proof.* We make the following two observations:

Observation 1: Let $k \in \mathbb{N}$. If D' has k pairwise internally vertex-disjoint paths from s to t, then D has k pairwise internally vertex-disjoint paths from X to Y.

[*Proof of Observation 1:* Assume that D' has k pairwise internally vertex-disjoint paths from s to t. We can "lift" these k paths to k paths from X to Y in D (preserving the arcs, and replacing the vertices s and t by appropriate vertices in X and Y to make them belong to the right arcs). The resulting k paths from X to Y in D are still pairwise internally vertex-disjoint (since our "lifting" operation has not changed the intermediate vertices of our paths). Thus, D has k pairwise internally vertex-disjoint paths from X to Y. This proves Observation 1.]

Observation 2: Let $k \in \mathbb{N}$. If D has k pairwise internally vertex-disjoint paths from X to Y, then D' has k pairwise internally vertex-disjoint paths from s to t.

[*Proof of Observation 2:* Assume that D has k pairwise internally vertex-disjoint paths from X to Y. We can replace these k paths by k pairwise internally vertex-disjoint **walks** from s to t in D' (by replacing their starting points with s and replacing their ending points with t). Thus, D' has k pairwise internally vertex-disjoint **walks** from s to t. Therefore, D' has

Hence, the result of the preceding paragraph is precisely the claim of Theorem 1.1.36, and our proof is thus complete. □

Another variant of this result can be stated for vertex-disjoint (as opposed to internally vertex-disjoint) paths. These are even easier to define:

Definition 1.1.38. Two paths $\mathbf{p}$ and $\mathbf{q}$ in a graph or digraph are said to be **vertex-disjoint** if they have no common vertices.

Theorem 1.1.39 (vertex-Menger theorem for directed graphs, multi-terminal version 2)**.** Let $D = (V, A, \psi)$ be a multidigraph, and let X and Y be two subsets of V.

Then, the maximum number of pairwise vertex-disjoint paths from X to Y equals the minimum size of an X-Y-vertex-separator.

Example 1.1.40. Let D be the following multidigraph:

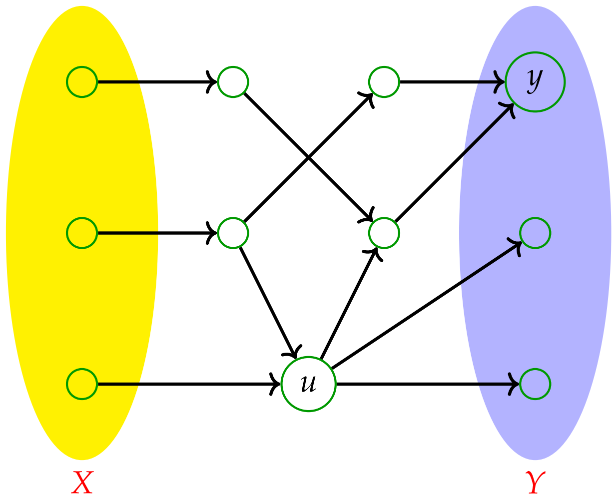

.

Then, the maximum number of pairwise vertex-disjoint paths from X to Y is

k pairwise internally vertex-disjoint **paths** from s to t as well (since each walk contains a path, and of course we don't lose internal vertex-disjointness if we restrict our walk to a path contained in it). This proves Observation 2.]

Observation 2 shows that the maximum number of pairwise internally vertex-disjoint paths from s to t in D' is $\geq$ to the maximum number of pairwise internally vertex-disjoint paths from X to Y in D. But Observation 1 shows the reverse inequality (i.e., it shows that the former number is $\leq$ to the latter number). Thus, the inequality is an equality, i.e., the two numbers are equal. Qed.

2. Here are two such paths (drawn in red and blue):

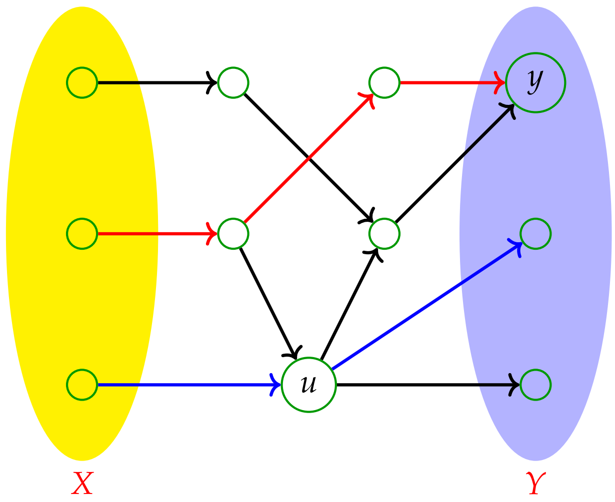

.

If we were only looking for **internally** vertex-disjoint paths, then we could add a third path to these two (namely, the path that starts at the topmost vertex of X and ends at the topmost vertex of Y). However, this path and our red paths are only **internally** vertex-disjoint, not vertex-disjoint. A little bit of thought shows that D has no more than 2 vertex-disjoint paths from X to Y.

The minimum size of an X-Y-vertex-separator is 2 as well; indeed, $\{u, y\}$ is such an X-Y-vertex-separator. This number equals the maximum number of pairwise vertex-disjoint paths from X to Y, just as Theorem 1.1.39 predicts.

Proof of Theorem 1.1.39. We will reduce this to Theorem 1.1.32, again by tweaking our digraph appropriately. This time, the tweak is pretty simple: We add two new vertices s and t to D, and we furthermore add an arc from s to each $x \in X$ and an arc from each $y \in Y$ to t (thus, we add a total of $|X| + |Y|$ new arcs). We denote the resulting digraph by D'. In more detail, the definition of D' is as follows:

- We introduce two new vertices s and t, and we set $V' := V \cup \{s, t\}$. This set V' will be the vertex set of D'.
- For each $x \in X$, we introduce a new arc a_x, which shall have source s and target x.
- For each $y \in Y$, we introduce a new arc b_y, which shall have source y and target t.
- We let $A' := A \cup \{a_x \mid x \in X\} \cup \{b_y \mid y \in Y\}$. This set A' will be the arc set of D'.

- We extend our map $\psi : A \to V \times V$ to a map $\psi' : A' \to V' \times V'$ by setting

$$\psi'(a_x) = (s, x) \qquad \text{for each } x \in X$$

and

$$\psi'(b_y) = (y, t) \qquad \text{for each } y \in Y$$

(and, of course, $\psi'(c) = \psi(c)$ for each $c \in A$).

- We define D' to be the multidigraph (V', A', ψ').

For instance, if D is the multidigraph from Example 1.1.40, then D' looks as follows:

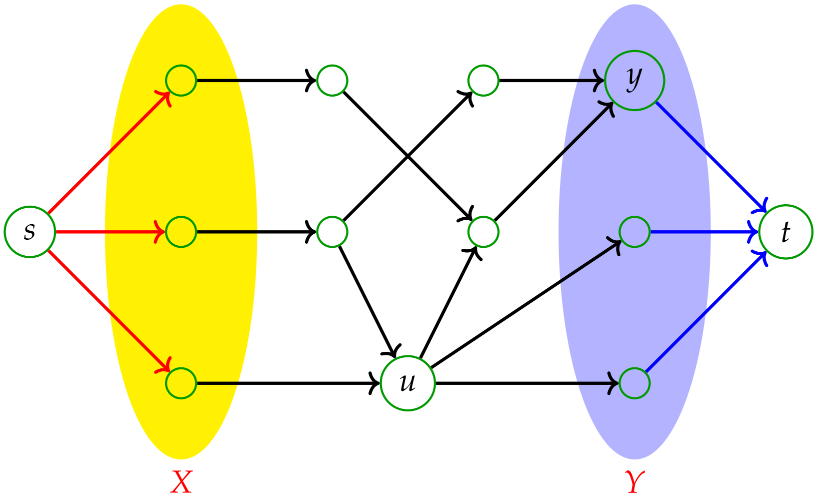

(The arcs a_x are drawn in red; the arcs b_y are drawn in blue.)

By its construction, the digraph D' has no arc with source s and target t. Hence, Theorem 1.1.32 (applied to $D' = (V', A', \psi')$ instead of $D = (V, A, \psi)$) yields that the maximum number of pairwise internally vertex-disjoint paths from s to t equals the minimum size of an internal s-t-vertex-separator. However, it is easy to see the following two claims:

> *Claim 1:* The maximum number of pairwise internally vertex-disjoint paths from s to t equals the maximum number of pairwise vertex-disjoint paths from X to Y (in D).

> *Claim 2:* The minimum size of an internal s-t-vertex-separator equals the minimum size of an X-Y-vertex-separator (in D).

[*Proof of Claim 1 (sketched):* Given any path $\mathbf{p}$ from s to t, we can remove the starting point and the ending point of this path; the result will always be a path from X to Y (in D). Let us denote the latter path by $\overline{\mathbf{p}}$. Thus, we obtain a map

$$\{\text{paths from } s \text{ to } t\} \to \{\text{paths from } X \text{ to } Y \text{ (in } D)\},$$
$$\mathbf{p} \mapsto \overline{\mathbf{p}}.$$

This map is easily seen to be a bijection (indeed, if $\mathbf{q}$ is a path from X to Y (in D), then we can easily extend it to a path $\mathbf{p}$ from s to t by inserting the appropriate arc a_x at its beginning and the appropriate arc b_y at its end; this latter path $\mathbf{p}$ will then satisfy $\overline{\mathbf{p}} = \mathbf{q}$). Moreover, two paths $\mathbf{p}$ and $\mathbf{q}$ from s to t are internally vertex-disjoint if and only if the corresponding paths $\overline{\mathbf{p}}$ and $\overline{\mathbf{q}}$ are vertex-disjoint (because the intermediate vertices of $\mathbf{p}$ are the vertices of $\overline{\mathbf{p}}$, whereas the intermediate vertices of $\mathbf{q}$ are the vertices of $\overline{\mathbf{q}}$). This proves Claim 1.]

[*Proof of Claim 2 (sketched):* It is easy to see that the internal s-t-vertex-separators are precisely the X-Y-vertex-separators (in D). (To show this, compare the definitions of these two objects using the bijection from the proof of Claim 1, and observe that $V' \setminus \{s, t\} = V$.) From this, Claim 2 follows.]

Recall that the maximum number of pairwise internally vertex-disjoint paths from s to t equals the minimum size of an internal s-t-vertex-separator. In view of Claim 1 and Claim 2, we can rewrite this as follows: The maximum number of pairwise vertex-disjoint paths from X to Y equals the minimum size of an X-Y-vertex-separator. Thus, Theorem 1.1.39 is proved. □

We note that Hall's Marriage Theorem (see Lecture 24) can be easily derived from any of the directed Menger theorems (exercise!). I have heard that this can also be done in reverse. This places the Menger theorems in the cluster of theorems equivalent to Hall's Marriage Theorem (such as König's theorem).

1.1.4. The vertex-Menger theorem for undirected graphs

Vertex-Menger theorems also exist for undirected graphs. Here are the undirected analogues of Theorem 1.1.32, Theorem 1.1.36 and Theorem 1.1.39, along with the definitions they rely on:

Definition 1.1.41. Let $G = (V, E, \varphi)$ be a multigraph, and let s and t be two vertices of G. A subset W of $V \setminus \{s, t\}$ is said to be an **internal** s-t-**vertex-separator** if each path from s to t contains at least one vertex from W. Equivalently, a subset W of $V \setminus \{s, t\}$ is said to be an **internal** s-t-**vertex-separator** if the induced subgraph of G on the set $V \setminus W$ has no path from s to t (in other words, removing from G all vertices contained in W destroys all paths from s to t).

Theorem 1.1.42 (vertex-Menger theorem for undirected graphs)**.** Let $G = (V, E, \varphi)$ be a multigraph, and let s and t be two distinct vertices of G. Assume that G has no edge with endpoints s and t. Then, the maximum number of pairwise internally vertex-disjoint paths from s to t equals the minimum size of an internal s-t-vertex-separator.

Definition 1.1.43. Let $G = (V, E, \varphi)$ be a multigraph, and let X and Y be two subsets of V.

(a) A **path from** X **to** Y shall mean a path whose starting point belongs to X and whose ending point belongs to Y.

(b) A subset W of V is said to be an X**-**Y**-vertex-separator** if each path from X to Y contains at least one vertex from W. Equivalently, a subset W of V is said to be an X**-**Y**-vertex-separator** if the induced subgraph of G on the set $V \setminus W$ has no path from X to Y (in other words, removing from G all vertices contained in W destroys all paths from X to Y).

(c) An X-Y-vertex-separator W is said to be **internal** if it is a subset of $V \setminus (X \cup Y)$ (that is, if it is disjoint from X and from Y).

Theorem 1.1.44 (vertex-Menger theorem for undirected graphs, multi-terminal version 1)**.** Let $G = (V, E, \varphi)$ be a multigraph, and let X and Y be two disjoint subsets of V. Assume that G has no edge with one endpoint in X and the other endpoint in Y.

Then, the maximum number of pairwise internally vertex-disjoint paths from X to Y equals the minimum size of an internal X-Y-vertex-separator.

Theorem 1.1.45 (vertex-Menger theorem for undirected graphs, multi-terminal version 2)**.** Let $G = (V, E, \varphi)$ be a multigraph, and let X and Y be two subsets of V.

Then, the maximum number of pairwise vertex-disjoint paths from X to Y equals the minimum size of an X-Y-vertex-separator.

Theorem 1.1.42, Theorem 1.1.44 and Theorem 1.1.45 follow immediately by applying the analogous theorems for directed graphs (i.e., Theorem 1.1.32, Theorem 1.1.36 and Theorem 1.1.39) to the digraph G^{bidir} instead of D (since the paths of G are in bijection with the paths of G^{bidir}).

References

[Schrij03] Alexander Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*, Springer 2003.
See `https://homepages.cwi.nl/~lex/co/` for errata.

[Schrij17] Alexander Schrijver, *A Course in Combinatorial Optimization*, March 23, 2017.
`https://homepages.cwi.nl/~lex/files/dict.pdf`