Math 4990 Fall 2017 (Darij Grinberg): homework set 9 solutions

The solutions below were written by GPT-5.5 and edited by myself (DG). They are less detailed than what I usually write.

One last binomial sum

> **Exercise 1.** Let $n \in \mathbb{N}$. Prove that
> $$\sum_{k=0}^{n} \binom{-2}{k} = (-1)^n \left\lfloor \frac{n+2}{2} \right\rfloor .$$

Solution to Exercise 1. For every $k \in \mathbb{N}$, we have

$$\binom{-2}{k} = \frac{(-2)(-3)\cdots(-k-1)}{k!} = (-1)^k \frac{2 \cdot 3 \cdots (k+1)}{k!} = (-1)^k (k+1) .$$

Hence,

$$\sum_{k=0}^{n} \binom{-2}{k} = \sum_{k=0}^{n} (-1)^k (k+1) . \qquad (1)$$

Now, we shall prove that

$$\sum_{k=0}^{n} (-1)^k (k+1) = \begin{cases} m+1, & \text{if } n = 2m \text{ for some } m \in \mathbb{N}; \\ -(m+1), & \text{if } n = 2m+1 \text{ for some } m \in \mathbb{N}. \end{cases} \qquad (2)$$

Proof of (2). A detailed proof of (2) can be found in [Grinbe15, Exercise 2.9], so we will only give a sketch.

We distinguish two cases.

Case 1: We have $n = 2m$ for some $m \in \mathbb{N}$.

In this case,

$$\begin{aligned} \sum_{k=0}^{n} (-1)^k (k+1) &= \sum_{k=0}^{2m} (-1)^k (k+1) \\ &= \sum_{i=0}^{m-1} ((2i+1) - (2i+2)) + (2m+1) \\ &= \sum_{i=0}^{m-1} (-1) + (2m+1) = -m + 2m + 1 = m+1. \end{aligned}$$

This proves (2) in Case 1.

Case 2: We have $n = 2m+1$ for some $m \in \mathbb{N}$.

In this case,

$$\begin{aligned}\sum_{k=0}^{n}(-1)^{k}(k+1) &= \sum_{k=0}^{2m+1}(-1)^{k}(k+1) \\ &= \sum_{i=0}^{m}((2i+1)-(2i+2)) = \sum_{i=0}^{m}(-1) \\ &= -(m+1).\end{aligned}$$

This proves (2) in Case 2.

So (2) is proved in both cases. □

Now, (1) becomes

$$\begin{aligned}\sum_{k=0}^{n}\binom{-2}{k} &= \sum_{k=0}^{n}(-1)^{k}(k+1) \\ &= \begin{cases} m+1, & \text{if } n=2m \text{ for some } m\in\mathbb{N}; \\ -(m+1), & \text{if } n=2m+1 \text{ for some } m\in\mathbb{N} \end{cases} \qquad \text{(by (2))} \\ &= (-1)^{n}\left\lfloor\frac{n+2}{2}\right\rfloor\end{aligned}$$

(the last equality sign can be verified by straightforward distinction of cases: if $n=2m$ for some $m\in\mathbb{N}$, then $(-1)^{n}=1$ and $\left\lfloor\frac{n+2}{2}\right\rfloor=m+1$; on the other hand, if $n=2m+1$ for some $m\in\mathbb{N}$, then $(-1)^{n}=-1$ and $\left\lfloor\frac{n+2}{2}\right\rfloor=m+1$).

We have thus proved Exercise 1. □

The Cartesian product of two permutations

We recall the following definition.

Definition 0.1. Let X be a finite set. Choose a bijection $\phi : X \to [n]$ for some $n\in\mathbb{N}$. For every permutation σ of X, define

$$(-1)^{\sigma}_{\phi} = (-1)^{\phi\circ\sigma\circ\phi^{-1}}.$$

It is known that this number depends only on σ, and not on the choice of ϕ. This number is called the *sign* of σ, and is denoted by $(-1)^{\sigma}$.

Exercise 2. Let U and V be two finite sets. Let σ be a permutation of U. Let τ be a permutation of V. We define a map $\sigma\times\tau : U\times V\to U\times V$ by

$$(\sigma\times\tau)(a,b) = (\sigma(a),\tau(b)) \qquad \text{for every } (a,b)\in U\times V.$$

(a) Prove that $\sigma \times \tau$ is a well-defined permutation.
(b) Prove that

$$\sigma \times \tau = (\sigma \times \operatorname{id}) \circ (\operatorname{id} \times \tau).$$

(c) Prove that

$$(-1)^{\sigma \times \tau} = \left((-1)^{\sigma}\right)^{|V|} \left((-1)^{\tau}\right)^{|U|}.$$

Solution to Exercise 2. A detailed solution can be found in [Grinbe15, Exercise 5.27]. What follows is an outline.

(a) The map $\sigma \times \tau$ is well-defined, because for every $(a,b) \in U \times V$, we have $\sigma(a) \in U$ and $\tau(b) \in V$, and thus $(\sigma(a), \tau(b)) \in U \times V$.

Moreover, the inverse map of $\sigma \times \tau$ is $\sigma^{-1} \times \tau^{-1}$. Indeed, for every $(a,b) \in U \times V$, we have

$$\begin{aligned}\left(\left(\sigma^{-1} \times \tau^{-1}\right) \circ (\sigma \times \tau)\right)(a,b) &= \left(\sigma^{-1} \times \tau^{-1}\right)(\sigma(a), \tau(b)) \\ &= \left(\sigma^{-1}(\sigma(a)), \tau^{-1}(\tau(b))\right) = (a,b)\end{aligned}$$

and

$$\begin{aligned}\left((\sigma \times \tau) \circ \left(\sigma^{-1} \times \tau^{-1}\right)\right)(a,b) &= (\sigma \times \tau)\left(\sigma^{-1}(a), \tau^{-1}(b)\right) \\ &= \left(\sigma\left(\sigma^{-1}(a)\right), \tau\left(\tau^{-1}(b)\right)\right) = (a,b).\end{aligned}$$

Thus,

$$\left(\sigma^{-1} \times \tau^{-1}\right) \circ (\sigma \times \tau) = \operatorname{id} \qquad \text{and} \qquad (\sigma \times \tau) \circ \left(\sigma^{-1} \times \tau^{-1}\right) = \operatorname{id}.$$

Hence $\sigma \times \tau$ is a permutation.

(b) For every $(a,b) \in U \times V$, we have

$$\begin{aligned}((\sigma \times \operatorname{id}) \circ (\operatorname{id} \times \tau))(a,b) &= (\sigma \times \operatorname{id})(a, \tau(b)) \\ &= (\sigma(a), \tau(b)) = (\sigma \times \tau)(a,b).\end{aligned}$$

Thus, $(\sigma \times \operatorname{id}) \circ (\operatorname{id} \times \tau) = \sigma \times \tau$. This proves part **(b)**.

(c) Set $m = |U|$ and $n = |V|$. Choose bijections $\alpha : U \to [m]$ and $\beta : V \to [n]$.

Define a map $\chi : U \times V \to [mn]$ by

$$\chi(u,v) = \alpha(u) + m(\beta(v) - 1).$$

This map is a bijection: indeed, each integer $t \in [mn]$ can be written in one and only one way in the form

$$t = i + m(j-1) \qquad \text{with } i \in [m] \text{ and } j \in [n],$$

namely with $j = \lceil t/m \rceil$ and $i = t - m(j-1)$. Thus $t = \chi\left(\alpha^{-1}(i), \beta^{-1}(j)\right)$. Hence χ is indeed a bijection. Also, for each fixed $v \in V$, the subset $U \times \{v\}$ is sent by χ onto one consecutive block of m integers[1].

Now consider the permutation $\sigma \times \operatorname{id}$ of $U \times V$. For each fixed $v \in V$, this permutation leaves the subset $U \times \{v\}$ stable, since

$$(\sigma \times \operatorname{id})(u,v) = (\sigma(u), v) \qquad \text{for every } (u,v) \in U \times \{v\}.$$

Moreover, under the identification[2]

$$U \times \{v\} \longrightarrow U, \qquad (u,v) \mapsto u,$$

the restriction of $\sigma \times \operatorname{id}$ to $U \times \{v\}$ becomes exactly the permutation σ of U. In this sense, the restriction of $\sigma \times \operatorname{id}$ to $U \times \{v\}$ acts just like σ.

Therefore, the permutation

$$\chi \circ (\sigma \times \operatorname{id}) \circ \chi^{-1}$$

of $[mn]$ is a product of n permutations with pairwise disjoint supports[3], one on each of the above consecutive blocks of size m. Each of these n permutations has sign $(-1)^{\sigma}$. Therefore,

$$(-1)^{\sigma \times \operatorname{id}} = \left((-1)^{\sigma}\right)^{n}. \tag{3}$$

Likewise, define a map $\chi' : U \times V \to [mn]$ by

$$\chi'(u,v) = \beta(v) + n(\alpha(u) - 1).$$

By the same argument as for χ, this map χ' is a bijection. For each fixed $u \in U$, the subset $\{u\} \times V$ is sent by χ' onto one consecutive block of n integers. Also, the restriction of $\operatorname{id} \times \tau$ to $\{u\} \times V$ corresponds, under the identification

$$\{u\} \times V \longrightarrow V, \qquad (u,v) \mapsto v,$$

[1] namely, of the m integers $1 + m(\beta(v) - 1), 2 + m(\beta(v) - 1), \ldots, m\beta(v)$

[2] "Identifying" things in mathematics generally means finding a bijection between two sets and pretending – at least notationally – that each element of its domain is the same as its image under this bijection. Of course, this is always risky, since pretending that different things are the same is a quick way to contradiction. Make sure you know what you are doing and be ready to explain yourself without such shortcuts.

In our specific case, we are considering the bijection

$$\gamma_v : U \times \{v\} \longrightarrow U, \qquad (u,v) \mapsto u$$

for a fixed $v \in V$. Then, the restriction of the permutation $\sigma \times \operatorname{id}$ to $U \times \{v\}$ is the map $\gamma_v^{-1} \circ \sigma \circ \gamma_v$ (as you can easily check by comparing how both maps act on a given element (u,v)). Hence, the sign of the former restriction equals the sign of σ (because of the general fact that if X and Y are two finite sets and $\gamma : X \to Y$ is a bijection, then, for any permutation π of Y, the sign of $\gamma^{-1} \circ \pi \circ \gamma$ equals the sign of π).

[3] The *support* of a permutation π of a set X is the set of all elements $x \in X$ that are not fixed by π. A permutation of X can always be restricted to a permutation of its support (since it sends elements of its support to other elements of its support), and moreover, this restriction has the same sign as π.

to the permutation τ of V. Hence, by the same argument as above,

$$(-1)^{\operatorname{id}\times\tau} = \left((-1)^{\tau}\right)^{m}. \tag{4}$$

But part **(b)** yields

$$\sigma\times\tau = (\sigma\times\operatorname{id})\circ(\operatorname{id}\times\tau).$$

Since sign is multiplicative, we thus obtain

$$\begin{aligned}(-1)^{\sigma\times\tau} &= (-1)^{\sigma\times\operatorname{id}}\cdot(-1)^{\operatorname{id}\times\tau} \\ &= \left((-1)^{\sigma}\right)^{n}\left((-1)^{\tau}\right)^{m} \qquad \text{(by (3) and (4))} \\ &= \left((-1)^{\sigma}\right)^{|V|}\left((-1)^{\tau}\right)^{|U|} \qquad (\text{since } m=|U| \text{ and } n=|V|).\end{aligned}$$

This proves part **(c)**. □

Non-cut vertices I and II

We recall a few definitions.

Definition 0.2. A *multigraph* means a triple (V, E, φ), where V and E are finite sets and $\varphi : E \to \mathcal{P}_2(V)$.

Definition 0.3. The *vertex set* of a multigraph $G = (V, E, \varphi)$ means the set V. It shall be denoted by $\operatorname{V}(G)$.

Definition 0.4. Let $G = (V, E, \varphi)$ and $G' = (V', E', \varphi')$ be two multigraphs. Then, G' is said to be a *subgraph* of G if and only if $V' \subseteq V$, $E' \subseteq E$, and $\varphi'(e) = \varphi(e)$ for each $e \in E'$.

Definition 0.5. If $G = (V, E, \varphi)$ is a multigraph and $v \in V$, then $G \setminus v$ denotes the subgraph obtained from G by removing the vertex v and all edges containing v.

Definition 0.6. Let G be a multigraph. A vertex v of G is called *non-cut* if and only if $G \setminus v$ is connected or has no vertices.

Definition 0.7. If G is a connected multigraph and if v and w are two vertices of G, then $d(v, w)$ denotes the smallest length of a path from v to w.

Exercise 3. Let $G = (V, E, \varphi)$ be a connected multigraph. Let $v \in V$ be any vertex.

(a) Pick any $w \in V$ such that $d(v, w)$ is maximum among all $w \in V$. Prove that w is a non-cut vertex of G.

(b) Let $n = |V|$. Prove that

$$\sum_{u\in V} d(v, u) \le \binom{n}{2}.$$

Solution to Exercise 3. **(a)** We must prove that $G \setminus w$ is connected or has no vertices. If $\operatorname{V}(G) = \{w\}$, then $G \setminus w$ has no vertices, and we are done. Thus, for the rest of this proof, we assume that $\operatorname{V}(G) \neq \{w\}$. Thus, G has more than one vertex; therefore, $w \neq v$ (since w has maximum distance from v).

Let $x \in \operatorname{V}(G) \setminus \{w\}$. Since G is connected, there exists a shortest path from v to x. Fix such a path P.

We claim that P does not pass through w. Indeed, assume the contrary. Then, since $x \neq w$, the vertex w is a proper internal vertex of the path P. Therefore, the initial segment of P from v to w is strictly shorter than P, and thus

$$d(v,w) < d(v,x).$$

This contradicts the maximality of $d(v,w)$. Hence our claim is proved.

Thus, the path P avoids w; therefore, P is also a path from v to x in $G \setminus w$. We have proved that every vertex $x \in \operatorname{V}(G) \setminus \{w\}$ is path-connected with v inside $G \setminus w$. Hence any two vertices of $G \setminus w$ are path-connected to each other inside $G \setminus w$ (namely, through v). Therefore, $G \setminus w$ is connected. So w is a non-cut vertex.

(b) We shall prove part **(b)** by induction on n.

If $n = 1$, then $\operatorname{V}(G) = \{v\}$, so that

$$\sum_{u \in \operatorname{V}(G)} d(v,u) = d(v,v) = 0 = \binom{1}{2}.$$

Thus part **(b)** is proved in this case.

Now assume that $n > 1$, and let us prove part **(b)**. Choose a vertex $w \in \operatorname{V}(G)$ such that $d(v,w)$ is maximum among all $w \in \operatorname{V}(G)$. Part **(a)** shows that w is non-cut. Hence the multigraph $H = G \setminus w$ is connected. Clearly, $\operatorname{V}(H)$ has $n-1$ elements.

Now let $u \in \operatorname{V}(G) \setminus \{w\}$. We claim that every shortest path from v to u in G avoids w. Indeed, if a shortest path from v to u in G passed through w, then $u \neq w$ would force w to be a proper internal vertex of this path, and we would obtain $d(v,w) < d(v,u)$, contradicting the maximality of $d(v,w)$.

Thus, a shortest path from v to u in G is already a path in H. Consequently,

$$d_H(v,u) = d_G(v,u) \qquad \text{for every } u \in \operatorname{V}(G) \setminus \{w\}. \tag{5}$$

Here, of course, the subscript indicates in which multigraph the distance is taken.

Applying the induction hypothesis to the connected multigraph H (which has $n-1$ vertices and vertex set $\operatorname{V}(G) \setminus \{w\}$), we obtain

$$\sum_{u \in \operatorname{V}(G) \setminus \{w\}} d_H(v,u) \leq \binom{n-1}{2}.$$

Using (5), we can rewrite this as

$$\sum_{u \in \operatorname{V}(G) \setminus \{w\}} d_G(v,u) \leq \binom{n-1}{2}. \tag{6}$$

Also, since a path cannot visit more than n vertices, we have $d_G(v,w) \leq n-1$. Therefore,

$$\sum_{u \in \mathrm{V}(G)} d_G(v,u) = \sum_{u \in \mathrm{V}(G) \setminus \{w\}} d_G(v,u) + d_G(v,w)$$
$$\leq \binom{n-1}{2} + (n-1) \qquad \text{(by } d_G(v,w) \leq n-1 \text{ and (6))}$$
$$= \binom{n}{2}.$$

This proves part **(b)** and thus completes the solution. □

Exercise 4. Let G be a connected multigraph. Let H be a connected subgraph of G. Prove that the number of non-cut vertices of H is $\leq$ the number of non-cut vertices of G.

Solution to Exercise 4. For any connected multigraph K, let $c(K)$ denote the number of non-cut vertices of K.

We shall use three simple lemmas:

Lemma 0.8. Let K be a connected subgraph of the connected multigraph G. Assume that $\mathrm{V}(K) \neq \mathrm{V}(G)$. Then there exists a vertex $z \in \mathrm{V}(G) \setminus \mathrm{V}(K)$ that is adjacent in G to a vertex of K.

Proof. Choose any vertex $x \in \mathrm{V}(G) \setminus \mathrm{V}(K)$ (this exists since $\mathrm{V}(K) \neq \mathrm{V}(G)$). Since G is connected and $\mathrm{V}(K) \neq \varnothing$ (because K is connected), there exists a path from x to some vertex of K. Among all such paths, choose one of shortest length:

$$\mathbf{x} = (x_0, x_1, \ldots, x_\ell).$$

Then $x_\ell \in \mathrm{V}(K)$ and $x_0 = x \notin \mathrm{V}(K)$. Hence not all vertices of this path belong to $\mathrm{V}(K)$. Let j be the largest index such that $x_j \notin \mathrm{V}(K)$. Then $j < \ell$ (since $x_\ell \in \mathrm{V}(K)$), so that $x_{j+1} \in \mathrm{V}(K)$ (since j is the largest). Thus the vertex $z := x_j$ belongs to $\mathrm{V}(G) \setminus \mathrm{V}(K)$ and is adjacent in G to the vertex $x_{j+1} \in \mathrm{V}(K)$ (since x_{j+1} follows x_j on the path $\mathbf{x}$). This proves the lemma. □

Lemma 0.9. Let K be a connected multigraph. Let K' be a multigraph obtained from K by adding some extra edges but no new vertices. Then $c(K') \geq c(K)$.

Proof. Let x be a non-cut vertex of K. Then $K \setminus x$ is connected or has no vertices. But $K' \setminus x$ is obtained from $K \setminus x$ merely by adding some edges. Hence $K' \setminus x$ is again connected or has no vertices. Therefore x is a non-cut vertex of K'. Thus every non-cut vertex of K is also a non-cut vertex of K', and so $c(K') \geq c(K)$. □

Lemma 0.10. Let K be a connected multigraph. Let K' be a multigraph obtained from K by adjoining one new vertex z together with a new edge that connects this vertex z with a vertex w of K. Then $c(K') \geq c(K)$.

Proof. If K has only one vertex, then the claim is clear (since $c(K) = 1$ and $c(K') = 2$). Thus, from now on, we assume that K has at least two vertices.

The new vertex z is itself non-cut in K', because removing z leaves the connected multigraph K.

Now let x be a non-cut vertex of K such that $x \neq w$. Then, the graph $K \setminus x$ is connected or has no vertices. Since it cannot have no vertices (because K has at least two vertices), we thus see that it is connected.

The graph $K' \setminus x$ is obtained from this connected graph $K \setminus x$ by adding the new vertex z and a new edge that connects this vertex z with the existing vertex w (here we use $x \neq w$, which ensures that w is really a vertex of $K \setminus x$). This shows that $K' \setminus x$ is again connected (since a connected graph remains connected if we add a new vertex and join it to an existing vertex). Hence the vertex x remains non-cut in K'.

Forget that we fixed x. We thus have shown that any non-cut vertex x of K satisfying $x \neq w$ remains non-cut in K'. Thus, we cannot lose more than one non-cut vertex when we go from K to K'. However, we always gain the new non-cut vertex z. Thus, $c(K') \geq c(K)$. □

Now we solve the exercise. We must prove that $c(H) \leq c(G)$, that is, $c(G) \geq c(H)$.

If $\operatorname{V}(G) = \operatorname{V}(H)$, then G is obtained from H merely by adding edges. In this case, Lemma 0.9 immediately yields $c(G) \geq c(H)$, and we are done.

Thus we may assume that $\operatorname{V}(G) \neq \operatorname{V}(H)$. We shall now recursively construct a sequence $(K_0, K_1, \ldots, K_m)$ of connected subgraphs of G such that $K_0 = H$ and $\operatorname{V}(K_m) = \operatorname{V}(G)$, and such that each K_i (for $i > 0$) is obtained from K_{i-1} by adjoining a new vertex and a new edge that joins this new vertex with an existing vertex of K_{i-1}.

Set $K_0 = H$. If K_i has already been constructed and satisfies $\operatorname{V}(K_i) \neq \operatorname{V}(G)$, then Lemma 0.8 (applied to $K = K_i$) shows that there exists a vertex $z \in \operatorname{V}(G) \setminus \operatorname{V}(K_i)$ that is adjacent in G to a vertex of K_i. Let K_{i+1} be the multigraph obtained from K_i by adjoining the new vertex z together with an edge of G joining z to a vertex of K_i. Then, K_{i+1} is connected. Thus Lemma 0.10 yields

$$c(K_{i+1}) \geq c(K_i).$$

Repeating this process, we eventually obtain a connected spanning subgraph K_m of G having the same vertex set as G, and satisfying

$$c(K_m) \geq c(K_{m-1}) \geq c(K_{m-2}) \geq \cdots \geq c(K_0) = c(H). \tag{7}$$

Finally, the multigraph G is obtained from K_m by adding further edges and no new vertices (since K_m is a spanning subgraph of G). Hence Lemma 0.9 yields

$$c(G) \geq c(K_m). \tag{8}$$

Combining (7) with (8), we find

$$c(G) \geq c(H).$$

In other words, the number of non-cut vertices of H is at most the number of non-cut vertices of G. This proves Exercise 4. □

When do transpositions generate all permutations?

Exercise 5. Let $G = (V, E, \varphi)$ be a connected multigraph.

For each $e = \{u, v\} \in \mathcal{P}_2(V)$, let t_e be the permutation of V that switches u with v and leaves all other elements unchanged.

An *E-transposition* means a permutation of the form t_e for some $e \in \varphi(E)$.

Prove that every permutation of V can be written as a composition of some E-transpositions.

Solution to Exercise 5. Given two distinct elements $u, v \in V$, the permutation $t_{\{u,v\}}$ (that is, the transposition of V that swaps u with v) shall be denoted $t_{u,v}$. We first prove two lemmas.

Lemma 0.11. Let u, v, w be three distinct elements of V. Then,

$$t_{u,v} \circ t_{v,w} \circ t_{u,v} = t_{u,w}.$$

Proof. This can be shown by directly verifying that both sides send u, v, w to w, v, u, respectively, while leaving all other elements of V fixed (since each of the transpositions $t_{u,v}$, $t_{v,w}$, $t_{u,w}$ leaves these other elements fixed). □

Lemma 0.12. Let k be a positive integer, and let $v_0, v_1, \ldots, v_k$ be $k+1$ distinct elements of the set V. For each $i \in \{1, 2, \ldots, k\}$, set $e_i = \{v_{i-1}, v_i\}$. Then

$$t_{v_0,v_k} = t_{e_1} \circ t_{e_2} \circ \cdots \circ t_{e_{k-1}} \circ t_{e_k} \circ t_{e_{k-1}} \circ \cdots \circ t_{e_2} \circ t_{e_1} \tag{9}$$

(where the subscripts on the right hand side are $e_1, e_2, \ldots, e_k$ followed by $e_{k-1}, e_{k-2}, \ldots, e_1$).

Proof. We shall prove Lemma 0.12 by induction on k.

The base case is $k = 1$. In this case, (9) reduces to $t_{v_0,v_1} = t_{e_1}$, which is true because $e_1 = \{v_0, v_1\}$.

Now let $k > 1$, and assume that Lemma 0.12 has already been proved for $k-1$ in place of k. Applying this induction hypothesis to the k distinct elements $v_1, v_2, \ldots, v_k$, we obtain

$$t_{v_1,v_k} = t_{e_2} \circ t_{e_3} \circ \cdots \circ t_{e_{k-1}} \circ t_{e_k} \circ t_{e_{k-1}} \circ \cdots \circ t_{e_3} \circ t_{e_2}. \tag{10}$$

Now, $e_1 = \{v_0, v_1\}$, so that $t_{e_1} = t_{v_0,v_1}$. Furthermore,

$$\begin{aligned} & t_{e_1} \circ t_{e_2} \circ \cdots \circ t_{e_{k-1}} \circ t_{e_k} \circ t_{e_{k-1}} \circ \cdots \circ t_{e_2} \circ t_{e_1} \\ &= \underbrace{t_{e_1}}_{=t_{v_0,v_1}} \circ \underbrace{\left(t_{e_2} \circ t_{e_3} \circ \cdots \circ t_{e_{k-1}} \circ t_{e_k} \circ t_{e_{k-1}} \circ \cdots \circ t_{e_3} \circ t_{e_2}\right)}_{\substack{=t_{v_1,v_k} \\ \text{(by (10))}}} \circ \underbrace{t_{e_1}}_{=t_{v_0,v_1}} \\ &= t_{v_0,v_1} \circ t_{v_1,v_k} \circ t_{v_0,v_1} = t_{v_0,v_k} \end{aligned}$$

(by Lemma 0.11, applied to $u = v_0$, $v = v_1$ and $w = v_k$). Thus (9) is proved for the present k. This completes the induction. □

We next show that every transposition $t_{a,b}$ with $a, b \in V$ can be written as a composition of E-transpositions.

Indeed, let $a, b \in V$ be two distinct vertices of G. Since G is connected, there exists a path $(a = v_0, v_1, \ldots, v_k = b)$ in G from a to b. Since this is a path, the vertices $v_0, v_1, \ldots, v_k$ are distinct. For every $i \in \{1, 2, \ldots, k\}$, set $e_i = \{v_{i-1}, v_i\}$. Then each e_i belongs to $\varphi(E)$ (since v_{i-1} and v_i are consecutive vertices of our path and thus adjacent in G), so that t_{e_i} is an E-transposition. Thus, Lemma 0.12 shows that t_{v_0,v_k} is a composition of E-transpositions. In other words, $t_{a,b}$ is a composition of E-transpositions (since $a = v_0$ and $b = v_k$).

Thus we have shown that every transposition of the set V is a composition of E-transpositions.

But every permutation of a finite set can be written as a composition of transpositions[4]. Hence, every permutation of V can be written as a composition of transpositions of V, while each of the latter transpositions, in turn, is a composition of E-transpositions (by the preceding paragraph). Thus, every permutation of V can be written as a composition of E-transpositions. This is exactly what had to be proved. □

Watersheds in digraphs

The following exercise is a result known as the finite case of *Newman's lemma* or *diamond lemma*:

Exercise 6. Let D be a finite multidigraph having no cycles. A vertex v of D is called a *sink* if there is no arc of D with source v.

If u and v are vertices of D, then:

[4]See, e.g., [Grinbe15, Exercise 5.15 **(b)**] for the proof of this fact

> - we write $u \longrightarrow v$ if and only if D has an arc with source u and target v;
> - we write $u \overset{*}{\longrightarrow} v$ if and only if D has a path from u to v.
>
> Assume that the following no-watershed condition holds:
>
> For any three vertices u, v and w of D satisfying $u \longrightarrow v$ and $u \longrightarrow w$, there exists a vertex t of D such that $v \overset{*}{\longrightarrow} t$ and $w \overset{*}{\longrightarrow} t$.
>
> Prove that for each vertex p of D, there exists a **unique** sink q of D such that $p \overset{*}{\longrightarrow} q$.

Solution to Exercise 6. For each vertex x of D, let $h(x)$ be the length of a longest path starting at x. This number is well-defined, since any path of D has length at most $|\operatorname{V}(D)| - 1$ (and of course, at least one path starting at x exists: the trivial path (x)). We note that D has no cycles, so each walk of D is a path (since otherwise, it would contain a cycle). Thus, the paths of D are the same as the walks of D.

The relation $\overset{*}{\longrightarrow}$ on the vertex set of D is called *reachability*: We say that a vertex q is *reachable* from a vertex p if $p \overset{*}{\longrightarrow} q$.

We note that if $u \longrightarrow v$ is an arc of D, then

$$h(u) > h(v). \tag{11}$$

(Indeed, each path starting at v can be extended to a path starting at u by adding this arc $u \longrightarrow v$ at its front (since any walk of D is a path). Of course, the latter path is longer than the former. Thus, whatever path exists starting at v, there will exist a longer path starting at u. Thus, $h(u) > h(v)$. This proves (11).)

Thus, we can see that if u, v are two vertices of D such that $u \overset{*}{\longrightarrow} v$, then

$$h(u) \geq h(v). \tag{12}$$

(Indeed, $u \overset{*}{\longrightarrow} v$ means that there is a path from u to v in D, and then (11) shows that the vertices $p_0, p_1, \ldots, p_k$ of this path satisfy $h(u) = h(p_0) > h(p_1) > h(p_2) > \cdots > h(p_k) = h(v)$. This yields (12).)

We must prove the following statement:

> *Statement A:* Let p be a vertex of D. Then, there exists a unique sink q such that $p \overset{*}{\longrightarrow} q$.

Proof of Statement A. We shall prove Statement A by strong induction on $h(p)$:

Base case: Assume that $h(p) = 0$. Then there is no path of positive length starting at p, hence no walk of positive length starting at p either (since paths are the same as walks in D). Thus, there is no arc with source p. Hence p itself is a sink. Clearly, $p \overset{*}{\longrightarrow} p$, and p is the unique sink reachable from p (for lack of arcs). So Statement A is proved in this case.

Induction step: Let $N > 0$, and assume that Statement A is already proved for every vertex p with $h(p) < N$. Now let p be a vertex with $h(p) = N$. We must prove Statement A for p.

Since $N > 0$, the vertex p is not a sink. Hence there exists at least one arc $p \longrightarrow r$. Consider this arc and its target r. Thus, (11) shows that $h(p) > h(r)$. Therefore, $h(r) < h(p) = N$, so that we can apply the induction hypothesis to r instead of p, and conclude that there exists a unique sink reachable from r. Let us denote this sink by q_r.

We now claim that the sink q_r does not depend on the chosen outgoing arc $p \longrightarrow r$. Indeed, let $p \longrightarrow r$ and $p \longrightarrow s$ be two outgoing arcs from p. By the no-watershed condition, there exists a vertex t such that $r \overset{*}{\longrightarrow} t$ and $s \overset{*}{\longrightarrow} t$. Then, (12) shows $h(r) \geq h(t)$, so that $h(t) \leq h(r) < N$, and therefore the induction hypothesis yields that there is a unique sink reachable from t; call it q_t.

Now $r \overset{*}{\longrightarrow} t \overset{*}{\longrightarrow} q_t$, so q_t is a sink reachable from r. But recall that there is only one sink reachable from r, namely q_r. Hence $q_r = q_t$. Likewise, $q_s = q_t$, where q_s denotes the unique sink reachable from s (this exists for the same reason as q_r). Therefore $q_r = q_s$. This proves our claim that q_r does not depend on r.

Thus all outgoing neighbors of p lead to the same sink. Let q denote this common sink. Then $p \overset{*}{\longrightarrow} q$, because p has an outgoing arc to some vertex r, and then $r \overset{*}{\longrightarrow} q$.

It remains to prove uniqueness (i.e., that q is the only sink reachable from p). Let q' be any sink satisfying $p \overset{*}{\longrightarrow} q'$. Thus, $q' \neq p$, since p is not a sink. Hence the path from p to q' has positive length, so it begins with some arc $p \longrightarrow r$. Then q' is a sink reachable from r. But by the induction hypothesis, the unique sink reachable from r is $q_r = q$. Thus $q' = q$.

So q is the unique sink reachable from p. This completes the induction step, and therefore the proof of Statement A. □

Thus, the exercise is solved. □

Acyclic orientations and source pushing

We next recall the relevant definitions.

Definition 0.13. Let $G = (V, E, \psi)$ be a multigraph.

(a) An *orientation* of G is a map $\phi : E \to V \times V$ such that, whenever $\phi(e) = (u, v)$, we have $\psi(e) = \{u, v\}$. Given such an orientation ϕ and an edge $e \in E$, we say that e is *oriented from u toward v* in ϕ, where $(u, v) = \phi(e)$. Each orientation ϕ gives rise to a multidigraph (V, E, ϕ), which will often be called ϕ again (so, e.g., a "cycle in ϕ" means a cycle in this multidigraph).

(b) An orientation ϕ is called *acyclic* if the multidigraph (V, E, ϕ) has no directed cycles.

Definition 0.14. Let $G = (V, E, \psi)$ be a multigraph, and let ϕ be an orientation of G.

A vertex $v \in V$ is called a *source* of ϕ if no arc of the multidigraph (V, E, ϕ) has target v.

If v is a source of ϕ, then the orientation obtained from ϕ by *pushing the source* v is the orientation ϕ' defined by the following rules:

- if $e \in E$ satisfies $v \in \psi(e)$ and $\phi(e) = (v, u)$, then $\phi'(e) = (u, v)$;
- all other edges keep their orientations.

In other words, we reverse all edges containing v.

Exercise 7. Let $G = (V, E, \psi)$ be a connected multigraph. Set $n = |V|$.

Let $(\phi_0, \phi_1, \ldots, \phi_k)$ be a sequence of orientations of G, and let $(v_1, v_2, \ldots, v_k)$ be a sequence of vertices of G such that, for each $i \in \{1, 2, \ldots, k\}$, the orientation ϕ_i is obtained from ϕ_{i-1} by pushing the source v_i.

(a) Prove that if u and w are two mutually adjacent vertices of G, then between any two consecutive appearances of u in the sequence $(v_1, v_2, \ldots, v_k)$, the vertex w must appear at least once.

Now assume that $k > \dbinom{n}{2}$.

(b) Prove that each vertex of G appears at least once in the sequence $(v_1, v_2, \ldots, v_k)$.

(c) Prove that the orientations $\phi_0, \phi_1, \ldots, \phi_k$ are acyclic.

Solution to Exercise 7. **(a)** Let u and w be two mutually adjacent vertices of G. Assume that u appears in the sequence $(v_1, v_2, \ldots, v_k)$ at two positions $a < b$ (that is, $v_a = v_b = u$), and that these two appearances are consecutive, meaning that u does not appear anywhere between these two positions (though u might appear at other positions in this sequence). We must prove that w appears somewhere among $v_{a+1}, v_{a+2}, \ldots, v_{b-1}$.

Assume the contrary. Then w does not appear among $v_{a+1}, v_{a+2}, \ldots, v_{b-1}$.

When we pass from ϕ_{a-1} to ϕ_a, we push the source $u = v_a$. Therefore every edge joining u to w is oriented toward u in ϕ_a. Now, as long as neither u nor w is pushed, the orientations of these edges cannot change. Since neither u nor w appears among $v_{a+1}, v_{a+2}, \ldots, v_{b-1}$, it follows that in ϕ_{b-1} every edge joining u to w is still oriented toward u.

But $u = v_b$ is a source of ϕ_{b-1}, so no edge of ϕ_{b-1} can be oriented toward u. This contradiction proves that w must indeed appear between the two consecutive appearances of u. This proves part **(a)**.

(b) Assume that some vertex $v \in V$ does not appear in the sequence $(v_1, v_2, \ldots, v_k)$. For each $u \in V$, let N_u denote the number of times that u appears in this sequence. Thus, $N_v = 0$ by our assumption. Also, $k = \sum_{u \in V} N_u$.

We shall prove, by induction on $d(v,u)$, that

$$N_u \leq d(v,u) \qquad \text{for every } u \in V. \tag{13}$$

Base case: If $d(v,u) = 0$, then $u = v$, so $N_u = N_v = 0 = d(v,u)$. Thus (13) holds in this case.

Induction step: Now let $u \in V$ satisfy $d(v,u) > 0$. Choose a shortest path[5]

$$(v = x_0, x_1, \ldots, x_\ell = u)$$

from v to u. Then $\ell = d(v,u)$, and the penultimate vertex $x_{\ell-1}$ is adjacent to u. Thus, setting $w = x_{\ell-1}$, we obtain

$$d(v,w) = \ell - 1$$

(indeed, $(v = x_0, x_1, \ldots, x_{\ell-1} = w)$ must be a shortest path from v to w, because if there were a shorter path, then we could extend it by the edge $x_{\ell-1} \to x_\ell$ to obtain a path from v to u shorter than the above path $(v = x_0, x_1, \ldots, x_\ell = u)$).

Part **(a)** shows that between any two consecutive appearances of u in $(v_1, v_2, \ldots, v_k)$, there must be an appearance of w. Therefore,

$$N_u \leq N_w + 1$$

(since the N_u appearances of u create $N_u - 1$ gaps, and each of these gaps must contain an appearance of w by the preceding sentence). But $d(v,w) = \ell - 1 = d(v,u) - 1$ (since $\ell = d(v,u)$), so we can apply the induction hypothesis to w instead of u. We obtain $N_w \leq d(v,w)$. Hence,

$$N_u \leq N_w + 1 \leq d(v,w) + 1 = d(v,u)$$

(since $d(v,w) = d(v,u) - 1$). Thus (13) is proved.

Now summing (13) over all $u \in V$, we obtain

$$k = \sum_{u \in V} N_u \leq \sum_{u \in V} d(v,u).$$

Exercise 3 **(b)** yields $\sum_{u \in V} d(v,u) \leq \binom{n}{2}$. Therefore $k \leq \binom{n}{2}$, contradicting the assumption $k > \binom{n}{2}$.

This contradiction shows that no vertex can be missing from the sequence $(v_1, v_2, \ldots, v_k)$. Hence each vertex of G appears at least once. This proves part **(b)**.

(c) We first show that, for each $i \in \{1, 2, \ldots, k\}$, the orientations ϕ_{i-1} and ϕ_i either both have a directed cycle or both have none.

Indeed, fix such an i. Since ϕ_i is obtained from ϕ_{i-1} by pushing the source v_i, the only edges whose directions change are the edges that contain v_i.

[5]We write all paths as lists of vertices, since the edges are not relevant to our proof.

Now, because v_i is a source of ϕ_{i-1}, no directed cycle in ϕ_{i-1} can pass through v_i or even use an edge containing v_i. Hence every directed cycle in ϕ_{i-1} is still a directed cycle in ϕ_i.

Likewise, after the push, the vertex v_i becomes a sink of ϕ_i. Therefore no directed cycle in ϕ_i can pass through v_i or use an edge containing v_i. Hence every directed cycle in ϕ_i is already a directed cycle in ϕ_{i-1}.

Thus, the directed cycles of ϕ_{i-1} are precisely the directed cycles of ϕ_i. Repeating this for all i, we conclude that the orientations $\phi_0, \phi_1, \ldots, \phi_k$ all have the same set of directed cycles. Therefore, either all orientations $\phi_0, \phi_1, \ldots, \phi_k$ are acyclic, or all of them have a directed cycle in common.

Assume for the sake of contradiction that they all have a directed cycle in common. Choose one such cycle, and let x be a vertex on this cycle. Then x has an incoming arc in each ϕ_i (namely, the arc of our cycle that enters x), so x is not a source of ϕ_i. Consequently, x can never be one of the pushed vertices $v_1, v_2, \ldots, v_k$.

This contradicts part **(b)**, which says that every vertex of G appears at least once in the sequence $(v_1, v_2, \ldots, v_k)$. Hence our assumption was false. Therefore all orientations $\phi_0, \phi_1, \ldots, \phi_k$ are acyclic. □

Exercise 8. Let $G = (V, E, \psi)$ be a connected multigraph, and fix a vertex $v \in V$.

If ϕ and ϕ' are two orientations of G, then we write $\phi \xrightarrow{v} \phi'$ if and only if ϕ' is obtained from ϕ by repeatedly pushing sources without ever pushing the source v.

An orientation ϕ is called *v-fleeing* if ϕ has no source other than v.

Prove that for any orientation ϕ of G, there is a **unique** v-fleeing orientation ϕ' such that $\phi \xrightarrow{v} \phi'$.

Solution to Exercise 8. We define a multidigraph O_v as follows:

- The vertices of O_v are all orientations of G.
- There is an arc $\phi \longrightarrow \psi$ in O_v whenever ψ is obtained from ϕ by pushing a source of ϕ different from v.

Since G has only finitely many edges, it has only finitely many orientations. Therefore O_v is a finite multidigraph.

We shall prove two facts about O_v.

First fact: The multidigraph O_v has no directed cycles.

Proof. Assume the contrary. Then O_v has a directed cycle. Traversing this cycle repeatedly, we can obtain, for any positive integer N, a sequence of more than N pushes of sources, none of which is equal to v, starting from some orientation of G.

Choose $N = \binom{|V|}{2}$. Then Exercise 7 **(b)** shows that in such a long sequence, every vertex of G must appear at least once among the pushed sources. In particular, the vertex v must appear. But by construction of O_v, the vertex v is never pushed. This contradiction proves that O_v has no directed cycles. □

Second fact: The multidigraph O_v satisfies the no-watershed condition (see Exercise 6).

Proof. Indeed, let ϕ, ψ and ψ' be three vertices of O_v such that $\phi \longrightarrow \psi$ and $\phi \longrightarrow \psi'$. Then there exist two sources x and y of ϕ, both distinct from v, such that ψ is obtained from ϕ by pushing x, whereas ψ' is obtained from ϕ by pushing y.

If $x = y$, then $\psi = \psi'$. In this case, to verify the no-watershed condition for O_v, we can simply take $t = \psi$ itself, since then both $\psi \overset{*}{\longrightarrow} t$ and $\psi' \overset{*}{\longrightarrow} t$ hold trivially. Thus we may assume that $x \neq y$.

We claim that x and y are not adjacent in the orientation ϕ. In fact, if there were an edge between x and y, then this edge would have to be oriented either from x towards y or from y towards x. In the first case, y would not be a source; in the second case, x would not be a source. Both are impossible. Thus there is no edge between x and y.

Hence pushing x does not affect any edge containing y, and so y remains a source after x has been pushed. Likewise, x remains a source after y has been pushed.

Now let θ be the orientation obtained from ϕ by reversing all edges containing x and all edges containing y. This makes sense because no edge contains both x and y (as x and y are not adjacent). Pushing x first (in ϕ) and then pushing y produces exactly this orientation θ; and pushing y first (in ϕ) and then x also produces exactly this same orientation θ. In other words, these two pushes commute. Therefore, if we push y in ψ, and if we push x in ψ', we arrive at the same orientation θ. Thus

$$\psi \longrightarrow \theta \qquad \text{and} \qquad \psi' \longrightarrow \theta$$

in O_v. Therefore

$$\psi \overset{*}{\longrightarrow} \theta \qquad \text{and} \qquad \psi' \overset{*}{\longrightarrow} \theta$$

in O_v as well. Thus, the no-watershed condition holds for O_v. □

Now Exercise 6 applies to the finite acyclic multidigraph O_v. It shows that for any vertex ϕ of O_v, there is a unique sink ϕ' of O_v such that $\phi \overset{*}{\longrightarrow} \phi'$ in O_v. By the definition of the arcs of O_v, the relation $\phi \overset{*}{\longrightarrow} \phi'$ in O_v is precisely the relation $\phi \overset{v}{\longrightarrow} \phi'$ from the statement of the exercise.

Finally, an orientation ϕ' is a sink of O_v if and only if no source of ϕ' distinct from v can be pushed. This is equivalent to saying that ϕ' has no source other than v, i.e., that ϕ' is v-fleeing.

Thus, for every orientation ϕ of G, there is a unique v-fleeing orientation ϕ' such that $\phi \overset{v}{\longrightarrow} \phi'$. This solves the exercise. □

References

[Grinbe15] Darij Grinberg, *Notes on the combinatorial fundamentals of algebra*, 15 September 2022, arXiv:2008.09862v3.