

Math 4707 Fall 2017 (Darij Grinberg): midterm 3

due date: Monday 11 December 2017 at the beginning of class, or before that by email or moodle

Please solve **at most 4** of the 10 exercises!

0.1. One last binomial sum

Exercise 1. Let $n \in \mathbb{N}$. Prove that

$$\sum_{k=0}^n \binom{-2}{k} = (-1)^n \left\lfloor \frac{n+2}{2} \right\rfloor.$$

0.2. The Cartesian product of two permutations

We have defined the sign of a permutation of $[n]$ for an $n \in \mathbb{N}$. But we can, more generally, define the sign of a permutation of **any** finite set. This would be difficult to define directly; instead, we define it by reducing it to a permutation of $[n]$ as follows:

Definition 0.1. Let X be a finite set. We want to define the sign of any permutation of X .

Fix a bijection $\phi : X \rightarrow [n]$ for some $n \in \mathbb{N}$. (Such a bijection always exists.) For every permutation σ of X , set

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

Here, the right hand side is well-defined because $\phi \circ \sigma \circ \phi^{-1}$ is a permutation of $[n]$.

It is not hard to check (see [Grinbe16, Exercise 5.12 (a)]) that $(-1)_{\phi}^{\sigma}$ depends only on the permutation σ of X , but not on the bijection ϕ . (In other words, for a given σ , any two different choices of ϕ will lead to the same $(-1)_{\phi}^{\sigma}$.)

This allows us to define the *sign* of the permutation σ to be $(-1)_{\phi}^{\sigma}$ for any choice of bijection $\phi : X \rightarrow [n]$. We denote this sign simply by $(-1)^{\sigma}$. (When $X = [n]$, then this sign is clearly the same as the sign $(-1)^{\sigma}$ we defined before, because we can pick the bijection $\phi = \text{id}$.)

(In contrast, we could **not** define the length $\ell(\sigma)$ of a permutation σ of X , because different bijections ϕ can lead to different values of $\ell(\phi \circ \sigma \circ \phi^{-1})$.)

The sign of a permutation σ of a finite set X has the following properties (see [Grinbe16, Exercise 5.12]):

- The permutation $\text{id} : X \rightarrow X$ satisfies $(-1)^{\text{id}} = 1$.
- We have $(-1)^{\sigma \circ \tau} = (-1)^{\sigma} \cdot (-1)^{\tau}$ for any two permutations σ and τ of X .

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 permutation $\sigma \times \tau$ of the set $U \times V$ by setting

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

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

(b) Prove that $\sigma \times \tau = (\sigma \times \text{id}) \circ (\text{id} \times \tau)$.

(c) Prove that $(-1)^{\sigma \times \tau} = ((-1)^\sigma)^{|V|} ((-1)^\tau)^{|U|}$. (All the signs here are well-defined due to Definition 0.1.)

(Can you find a slick proof for part (c) that involves no endless stream of trivial lemmas?)

0.3. Non-cut vertices I

See solutions to Spring 2017 Math 5707 homework set #2 (specifically, Section 0.1) for definitions of simple graphs, multigraphs, digraphs and multidigraphs. Note, in particular, that all of these are assumed to be finite (i.e., they have finitely many vertices and finitely many edges).

Recall that a *multigraph* is defined to be a triple (V, E, φ) , where V and E are two finite sets and φ is a map $E \rightarrow \mathcal{P}_2(V)$ (sending each “edge” $e \in E$ to an unordered pair of two distinct “vertices”). The elements of V are called the *vertices* of the multigraph; the elements of E are called its *edges*.

Definition 0.2. Let $G = (V, E, \varphi)$ and $G' = (V', E', \varphi')$ be two multigraphs. We say that G' is a *subgraph* of G if and only if $V' \subseteq V$, $E' \subseteq E$ and $(\varphi'(e) = \varphi(e) \text{ for each } e \in E')$.

Thus, a subgraph of a multigraph G is simply a multigraph obtained from G by removing some vertices and some edges¹, provided that for each vertex we remove, all edges containing that vertex are also removed. For example, the 2-vertex graph $1 \text{ --- } 2$ has 5 subgraphs: itself; the subgraph obtained by removing the edge (but leaving both vertices intact); the two subgraphs obtained by removing one vertex (along with the edge); and finally the subgraph obtained by removing everything.

Definition 0.3. Let $G = (V, E, \varphi)$ be a multigraph. Let $v \in V$ be a vertex. Then, $G \setminus v$ shall denote the subgraph $(V \setminus \{v\}, E', \varphi|_{E'})$ of G , where E' is the set of all edges $e \in E$ that don't contain the vertex v . In other words, $G \setminus v$ is the subgraph of G obtained by removing the vertex v and all edges containing v .

For example, if G is the 3-vertex graph $1 \text{ --- } 2 \text{ --- } 3$, then $G \setminus 1$ is the 2-vertex graph $2 \text{ --- } 3$, whereas $G \setminus 2$ is the 2-vertex graph $1 \text{ --- } 3$.

¹“Some” may mean “none”, and may also mean “all” (as well as anything inbetween).

Definition 0.4. Let $G = (V, E, \varphi)$ be a multigraph. A vertex $v \in V$ is said to be *non-cut* if the multigraph $G \setminus v$ is connected or has no vertices.

For example, if G is the 3-vertex graph $1 \text{ --- } 2 \text{ --- } 3$, then the non-cut vertices of G are 1 and 3.

Definition 0.5. Let G be a multigraph. Let $(v_0, e_1, v_1, e_2, v_2, \dots, v_{k-1}, e_k, v_k)$ be a walk in G . Then, the *length* of this walk is defined to be k (that is, the number of edges).

Definition 0.6. Let $G = (V, E, \varphi)$ be a connected multigraph. Let $v \in V$ and $w \in V$ be two vertices. Then, $d(v, w)$ (the *distance* between v and w) is defined as the smallest length of a path from v to w . (This is also the smallest length of a walk from v to w , because every walk from v to w can be trimmed down to a path of the same or smaller length.)

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

Exercise 3 (a) is particularly important, as it guarantees that any connected multigraph with at least one vertex has a non-cut vertex. This allows proving properties of connected multigraphs by induction on the number of vertices.

Note that the inequality in Exercise 3 (b) is sharp (i.e., equality can hold): If V is the n -vertex graph $1 \text{ --- } 2 \text{ --- } 3 \text{ --- } \dots \text{ --- } n$ and if $v = 1$, then $\sum_{u \in V} \underbrace{d(v, u)}_{=u-1} =$

$$0 + 1 + \dots + (n-1) = \binom{n}{2}.$$

0.4. Non-cut vertices II: subgraphs

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 $\leq$ to the number of non-cut vertices of G .

0.5. When do transpositions generate all permutations?

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

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

An *E-transposition* shall mean 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*.

[You are allowed to use the result of Exercise 3 **(a)** here even if you have not solved that exercise.]


Note that Exercise 5 generalizes Exercise 3 on Math 4990 homework set #7, because the simple graph $([n], \{\{1, 2\}, \{1, 3\}, \dots, \{1, n\}\})$ (for $n > 0$) is connected.

Exercise 5 also generalizes the fact that every permutation of $[n]$ can be written as a composition of simple transpositions $s_i = t_{i, i+1}$, because the simple graph $([n], \{\{1, 2\}, \{2, 3\}, \dots, \{n-1, n\}\})$ (for $n > 0$) is connected.

Exercise 5 also has a converse: If $G = (V, E, \varphi)$ is a multigraph such that every permutation of V can be written as a composition of some E -transpositions, then G is connected or V is empty. This is not hard to check.

0.6. Watersheds in digraphs

Example 0.7. Consider the following digraph:




Imagine a game chip placed initially at the vertex 1. The chip is allowed to move along the arcs of the digraph (from source to target). For example, the chip can first move along the arc $(1,2)$ to 2, then along the arc $(2,3)$ to 3, then along the arc $(3,5)$ to 5. Once it arrives at 5, it can no longer move, because there are no arcs with source 5. We say that 5 is a *sink* for this reason (see Exercise 6 below for the precise definition).

Alternatively, the chip could have moved along the arc $(1,2)$ to 2, then along the arc $(2,6)$ to 6, then along the arc $(6,7)$ to 7. At this point it would again be stuck, since 7 is a sink.

Thus, the chip can get stuck in **two different sinks**, depending on the path it takes. (It will always get stuck in **some** sink, because our digraph has no cycles.)

Now, consider the following digraph:



This time, any chip starting at any given vertex will necessarily get stuck at **the same sink** no matter what path it takes (either the sink 1, if it started at one of the vertices 1,2,3,4,5,6,7; or the sink 9, if it started at one of the vertices 8,9,10). How can we show this without checking all possible paths?

One criterion, which is clearly necessary, is that there are no “watershed vertices”: i.e., there is no vertex u from which the chip can take two different arcs

(u, v) and (u, w) such that v and w “never meet again” (i.e., there exists no vertex reachable both from v and from w). For example, the digraph (1) has a “watershed vertex” (namely, 3, because the arcs $(3, 5)$ and $(3, 6)$ lead to the vertices 5 and 6 which “never meet again”).

The next exercise claims that this condition is also sufficient (as long as our digraph has no cycles). That is, if there are no “watershed vertices” and no cycles, then the sink at which a chip gets stuck is uniquely determined by the vertex it started at (rather than by the path it took).

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

If u and v are any two vertices of D , then:

- we write $u \rightarrow v$ if and only if D has an **arc** with source u and target v ;
- we write $u \xrightarrow{*} v$ if and only if D has a **path** from u to v .

The so-called *no-watershed condition* says that for any three vertices u, v and w of D satisfying $u \rightarrow v$ and $u \rightarrow w$, there exists a vertex t of D such that $v \xrightarrow{*} t$ and $w \xrightarrow{*} t$.

Assume that the no-watershed condition holds. Prove that for each vertex p of D , there exists a **unique** sink q of D such that $p \xrightarrow{*} q$.

[Hint: Induction on the “height” of p (that is, the length of a longest path starting at p).]

0.7. Acyclic orientations and source pushing


Roughly speaking, an *orientation* of a multigraph G is a way to assign to each edge of G a direction (thus making it an arc). If the resulting **multidigraph** has no cycles, then this orientation will be called *acyclic*. A rigorous way to state this definition is the following:

Definition 0.8. Let $G = (V, E, \psi)$ be a multigraph.

(a) An *orientation* of G is a map $\phi : E \rightarrow V \times V$ such that each $e \in E$ has the following property: If we write $\phi(e)$ in the form $\phi(e) = (u, v)$, then $\psi(e) = \{u, v\}$.

(b) An orientation ϕ of G is said to be *acyclic* if and only if the multidigraph (V, E, ϕ) has no cycles.


Example 0.9. Let $G = (V, E, \psi)$ be the following multigraph:



Then, the following four maps ϕ are orientations of G :

- the map sending a to $(1,2)$, sending b to $(1,2)$, sending c to $(3,2)$, and sending d to $(1,3)$;
- the map sending a to $(2,1)$, sending b to $(1,2)$, sending c to $(3,2)$, and sending d to $(3,1)$;
- the map sending a to $(1,2)$, sending b to $(1,2)$, sending c to $(2,3)$, and sending d to $(1,3)$;
- the map sending a to $(1,2)$, sending b to $(1,2)$, sending c to $(2,3)$, and sending d to $(3,1)$.

Here are the multidigraphs (V, E, ϕ) corresponding to these four maps (in the order mentioned):



Only the first and the third of these orientations ϕ are acyclic (since only the first and the third of these multidigraphs have no cycles).

Definition 0.10. Let $G = (V, E, \psi)$ be a multigraph.

Let ϕ be an orientation of G .


A vertex $v \in V$ is said to be a *source* of ϕ if no arc of the multidigraph (V, E, ϕ) has target v . Exercise 6 (a) on Math 5707 (Spring 2017) homework set #5 shows that if ϕ is acyclic and if $V \neq \emptyset$, then there exists a source of ϕ .

If v is a source of ϕ , then we can define a new orientation ϕ' of G as follows:


- For each $e \in E$ satisfying $v \in \psi(e)$, we set $\phi'(e) = (u, v)$, where u is chosen such that $\phi(e) = (v, u)$.
- For all other $e \in E$, we set $\phi'(e) = \phi(e)$.

(Roughly speaking, this simply means that ϕ' is obtained by ϕ by reversing the directions of all edges that contain v .) We say that this new orientation ϕ' is obtained from ϕ by *pushing the source* v .

Example 0.11. Let $G = (V, E, \psi)$ be the following multigraph:




Consider the orientation ϕ of G for which the multidigraph (V, E, ϕ) looks as follows:




(Formally speaking, this is the orientation ϕ that sends the edges a, b, c, d, e, f to the pairs $(1, 2), (3, 2), (1, 4), (3, 5), (4, 5), (5, 4)$, respectively.)

This orientation ϕ has two sources 1 and 3. We can transform this orientation by pushing the source 1; this results in the following orientation ϕ' (shown here by drawing the multidigraph (V, E, ϕ')):




This new orientation ϕ' has a single source, 3. If we push this source, we obtain a new orientation ϕ'' , which looks as follows (again, represented by the multidigraph (V, E, ϕ'')):



This orientation ϕ'' , in turn, has a single source, 2. If we push this source, we

obtain a new orientation ϕ''' , which looks as follows (again, represented by the multidigraph (V, E, ϕ''')):



This orientation ϕ''' has no sources, and thus cannot be transformed any further by pushing sources.

The preceding example suggests some questions: For example, given an orientation of a multigraph, can we keep pushing sources indefinitely, or will we eventually end up at an orientation that has no more sources? The following is easy to see:

Proposition 0.12. Let ϕ be an acyclic orientation of a multigraph $G = (V, E, \psi)$. Let v be a source of ϕ . Then, the orientation obtained from ϕ by pushing the source v is again acyclic.

This proposition shows that if we start with an acyclic orientation of a multigraph (with at least one vertex), then we can keep pushing sources indefinitely (since the orientation always remains acyclic, and thus there always will be sources to push). The next exercise (specifically, Exercise 7 (c)) yields a converse (for connected multigraphs): If we can keep pushing sources indefinitely (or, even, if we can keep pushing sources for more than $\binom{n}{2}$ times in a row), then our orientation must have been acyclic.

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

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

(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 $(v_1, v_2, \dots, v_k)$, the vertex w must appear at least once.

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

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

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

[**Hint:** For part (b), assume that some vertex v does not appear in the sequence $(v_1, v_2, \dots, v_k)$. Then, argue that any vertex $u \in V$ appears at most $d(v, u)$ times in this sequence, using part (a). Then apply Exercise 3 (b). For part (c), first argue that any cycle existing in **one** of the orientations $\phi_0, \phi_1, \dots, \phi_k$ would automatically exist in **all** of these orientations.]

[You may use Exercise 3 (b) even if you have not solved this exercise.]

Exercise 8. Let $G = (V, E, \psi)$ be a connected multigraph.

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 . (More rigorously: We write $\phi \xrightarrow{v} \phi'$ if and only if there exist a sequence $(\phi_0, \phi_1, \dots, \phi_k)$ of orientations of G and a sequence $(v_1, v_2, \dots, v_k)$ of vertices of G distinct from v such that for each $i \in \{1, 2, \dots, k\}$, the orientation ϕ_i is obtained from ϕ_{i-1} by pushing the source v_i (in particular, this is saying that v_i is a source of ϕ_{i-1}), and such that $\phi_0 = \phi$ and $\phi_k = \phi'$.)

If ϕ is an orientation of G , then we say that ϕ is *v-fleeing* if ϕ has no source other than v . (Note that ϕ may or may not have v as a source.)

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

[**Hint:** Consider a new multidigraph O_v whose vertices are the orientations of G , and which has an arc from an orientation ϕ to an orientation ϕ' if and only if ϕ' can be obtained from ϕ by pushing a source different from v . Use Exercise 7 (b) to argue that this multidigraph O_v has no cycles, and then use Exercise 6.]

[You may use both exercises mentioned in the hint without solving them.]

0.8. On network flows

The following two exercises are probably best approached after reading [Grinbe17].

Exercise 9. State and prove the analogues of all definitions and results from [Grinbe17, §1.1–§1.8] for networks in which the digraph is replaced by a multidigraph (so the arcs are no longer just pairs of vertices, but rather arbitrary objects that map to pairs of vertices).

[**Hint:** You don't have to give the whole proofs, as long as you explain what needs to be modified and how. For example, what would a^{-1} be if a is an arc of a multidigraph? Alternatively, you can deduce the analogues from the results proven in [Grinbe17].]

Definition 0.13. Let G be a multigraph. A *spanning subgraph* of G means a subgraph H of G such that each vertex of G is a vertex of H .

Thus, a spanning subgraph of a multigraph G is obtained from G by removing some edges², but not removing any vertices.

Definition 0.14. Let $H = (V, E, \varphi)$ be a multigraph. The *degree function* of H means the function $\deg : V \rightarrow \mathbb{N}$ that sends each vertex $v \in V$ of H to its degree $\deg v$.

Exercise 10. Let $(G; X, Y)$ be a bipartite graph, where $G = (V, E, \varphi)$ is a multigraph. Let $\gamma : V \rightarrow \mathbb{N}$ be a function such that each $y \in Y$ satisfies $\gamma(y) \leq \deg y$. Prove that the following two statements are equivalent:

Statement 1: There exists a spanning subgraph of G with degree function γ . (In other words, there exists a spanning subgraph of G such that each vertex $v \in V$ is contained in exactly $\gamma(v)$ edges of this subgraph.)

Statement 2: Every subset S of X satisfies

$$\sum_{s \in S} \gamma(s) \leq \sum_{y \in Y} \min \{ \gamma(y), \deg y \},$$

and this inequality becomes an equality for $S = X$ (but not necessarily only for $S = X$).

References

- [Grinbe16] Darij Grinberg, *Notes on the combinatorial fundamentals of algebra*, 10 January 2019.
<http://www.cip.ifi.lmu.de/~grinberg/primes2015/sols.pdf>
 The numbering of theorems and formulas in this link might shift when the project gets updated; for a “frozen” version whose numbering is guaranteed to match that in the citations above, see <https://github.com/darijgr/detnotes/releases/tag/2019-01-10>.
- [Grinbe17] Darij Grinberg, *UMN, Spring 2017, Math 5707: Lecture 16 (flows and cuts in networks)*,
<http://www.cip.ifi.lmu.de/~grinberg/t/17s/5707lec16.pdf>.

²“Some” includes the options “none” and “all”.