Math 4707 Spring 2018 (Darij Grinberg): midterm 3 with solutions [preliminary version]

Contents

Recall the following:

- If $n \in \mathbb{N}$, then $[n]$ denotes the n-element set $\{1, 2, \ldots, n\}$.

0.1. Ordering acyclic digraphs

See Spring 2017 Math 5707 Homework set #2 (or our class notes from April 23) for the definition of a "multidigraph". We will refer to multidigraphs simply as *digraphs*. Thus, any digraph D is a triple (V, A, φ), where V and A are finite sets, and where φ is a map $A \to V \times V$.

Definition 0.1. A digraph is said to be *acyclic* if it has no cycles.

For example, the digraph

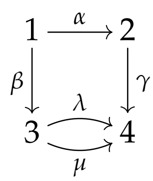

(1)

is acyclic, whereas the digraph

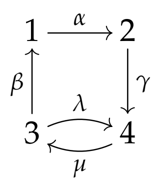

is not acyclic (both $(1, \alpha, 2, \gamma, 4, \mu, 3, \beta, 1)$ and $(3, \lambda, 4, \mu, 3)$ are cycles of this latter digraph).

[Acyclic digraphs are often called "dags", apparently because the proper abbreviation "adgs" would be harder to pronounce.]

Exercise 1. Let $D = (V, A, \varphi)$ be an acyclic digraph. Prove that there is a list $(v_1, v_2, \ldots, v_n)$ of elements of V such that

- each element of V appears exactly once in this list $(v_1, v_2, \ldots, v_n)$;
- whenever i and j are two elements of $[n]$ such that some arc of D has source v_i and target v_j, we must have $i < j$.

(In other words, prove that there is a list consisting of all vertices of V, which contains each of them exactly once, and with the property that the source of any arc must appear before the target of this arc in the list. For example, if D is the digraph (1), then there are two such lists: $(1, 2, 3, 4)$ and $(1, 3, 2, 4)$.)

Exercise 1 is a fundamental result of far-reaching use. In the situation of Exercise 1, a list $(v_1, v_2, \ldots, v_n)$ of elements of V satisfying the two required conditions is called a *topological ordering or topological sorting* of D. Any Linux distribution comes with a built-in utility for computing a topological sorting of a digraph, `tsort`. For example, if I want to compute a topological sorting of the digraph

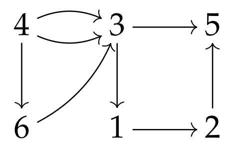

(2)

(whose arcs I have not labeled because the names of the arcs don't matter), I just create a text file (say, `arcs.txt`) that lists its arcs (each arc on a separate line, in the format "`source target`"):

```
4 3
4 3
3 5
4 6
6 3
3 1
1 2
2 5
```

and then run "`tsort arcs.txt`", which yields the following output:

```
$ tsort arcs.txt
4
6
3
1
2
5
```

This is telling me that $(4, 6, 3, 1, 2, 5)$ is a topological ordering of the digraph (2). And indeed, it is.

Why would Linux provide such a utility? The specific reasons seem to be historical, but the algorithm itself is highly useful. For example, you want to install a piece of software through a package manager (e.g., `apt-get install`), but that software depends on some other software that is not installed, and that latter software in turn depends on some further software... These dependencies are known (each program declares a list of other programs it depends on, or at least its direct dependencies), so the package manager learns a hopefully acyclic digraph whose vertices are software packages and whose arcs stand for dependencies (i.e., an arc with source s and target t means that program s needs to be installed before t). Now the package manager needs to know in what order the packages may be installed. That order has to be a topological sorting of the digraph.

I shall outline one of many possible solutions for Exercise 1:

Solution to Exercise 1 (sketched). For each vertex v of D, we let $\operatorname{Anc}(v)$ be the set of all $w \in V$ such that there exists a path from w to v in D. The elements of $\operatorname{Anc}(v)$ are called the *ancestors* of v (whence the notation Anc).

Now, let $(v_1, v_2, \ldots, v_n)$ be a list of all vertices of D in the order of increasing $|\operatorname{Anc}(v)|$, where ties are resolved arbitrarily. In other words, let $(v_1, v_2, \ldots, v_n)$ be a list of all vertices of D (with each vertex appearing exactly once in this list) satisfying

$$|\operatorname{Anc}(v_1)| \leq |\operatorname{Anc}(v_2)| \leq \cdots \leq |\operatorname{Anc}(v_n)| . \tag{3}$$

(Such a list clearly exists.)

The digraph D is acyclic, and thus has no cycles.

Now, we claim the following:

> *Claim 1:* Each element of V appears exactly once in this list $(v_1, v_2, \ldots, v_n)$.
>
> *Claim 2:* Whenever i and j are two elements of $[n]$ such that some arc of D has source v_i and target v_j, we must have $i < j$.

Once these two claims are proven, the exercise will clearly follow, since the very assertion of the exercise was the existence of a list $(v_1, v_2, \ldots, v_n)$ of elements of V satisfying Claim 1 and Claim 2. So it remains to prove Claim 1 and Claim 2.

[*Proof of Claim 1:* Claim 1 is obvious from the construction of the list $(v_1, v_2, \ldots, v_n)$.]

[*Proof of Claim 2:* Let i and j be two elements of $[n]$ such that some arc of D has source v_i and target v_j. We must prove that $i < j$.

Assume the contrary. Thus, $i \geq j$, so that $j \leq i$. Hence, from (3), we obtain $|\operatorname{Anc}(v_j)| \leq |\operatorname{Anc}(v_i)|$.

We know that some arc of D has source v_i and target v_j. Consider such an arc, and denote it by a.

But $\operatorname{Anc}(v_i) \subseteq \operatorname{Anc}(v_j)$ [1]. However, $v_j \in \operatorname{Anc}(v_j)$ [2] and therefore $\operatorname{Anc}(v_i) \neq \operatorname{Anc}(v_j)$ [3]. Combining this with $\operatorname{Anc}(v_i) \subseteq \operatorname{Anc}(v_j)$, we conclude that $\operatorname{Anc}(v_i)$ is a **proper** subset of $\operatorname{Anc}(v_j)$. Hence, $|\operatorname{Anc}(v_i)| < |\operatorname{Anc}(v_j)|$. This contradicts $|\operatorname{Anc}(v_j)| \leq |\operatorname{Anc}(v_i)|$. This contradiction shows that our assumption was false. Thus, we have $i < j$. This proves Claim 2.]

Having proven both Claim 1 and Claim 2, we are done solving Exercise 1 (as explained above). □

0.2. Watersheds in digraphs

A *simple digraph* means a pair (V, A), where V is a finite set, and where A is a subset of $V \times V$. We identify every simple digraph (V, A) with the multidigraph (V, A, ι), where ι is the map sending each $(u, v) \in A$ to $(u, v) \in V \times V$. Thus, simple digraphs are the same as multidigraphs whose arcs are already pairs of vertices, the first entry being the source and the second entry being the target. (So the relation between simple digraphs and multidigraphs is the same as the relation between simple graphs and multigraphs.)

[1] *Proof.* Let $t \in \operatorname{Anc}(v_i)$. Thus, t is an element of V such that there exists a path from t to v_i in D (by the definition of $\operatorname{Anc}(v_i)$). Consider this path; denote it by $\mathbf{p}$. Extend this path $\mathbf{p}$ further by the arc a and the vertex v_j; it then becomes a walk from t to v_j. Hence, there exists a walk from t to v_j. Thus, there exists a path from t to v_j as well (since we can turn any walk into a path by removing redundant parts). Hence, t is an element of V such that there exists a path from t to v_j in D. In other words, $t \in \operatorname{Anc}(v_j)$ (by the definition of $\operatorname{Anc}(v_j)$).

Now, forget that we fixed t. We thus have proven that $t \in \operatorname{Anc}(v_j)$ for each $t \in \operatorname{Anc}(v_i)$. In other words, $\operatorname{Anc}(v_i) \subseteq \operatorname{Anc}(v_j)$.

[2] *Proof.* There exists a path from v_j to v_j in D (namely, the trivial path (v_j)). Hence, v_j is an element w of V such that there exists a path from w to v_j in D. In other words, $v_j \in \operatorname{Anc}(v_j)$ (by the definition of $\operatorname{Anc}(v_j)$).

[3] *Proof.* Assume the contrary. Thus, $\operatorname{Anc}(v_i) = \operatorname{Anc}(v_j)$. Hence, $v_j \in \operatorname{Anc}(v_j) = \operatorname{Anc}(v_i)$. In other words, v_j is an element w of V such that there exists a path from w to v_i in D (by the definition of $\operatorname{Anc}(v_i)$). In other words, v_j is an element of V, and there exists a path from v_j to v_i in D. Consider this path. Extend this path by the arc a and the vertex v_j. The result will be a cycle from v_j to v_j. Thus, the digraph D has a cycle. This contradicts the fact that the digraph D has no cycles. This contradiction shows that our assumption was false, qed.

Example 0.2. Consider the following simple digraph:

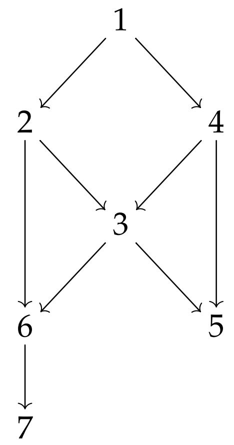

. (4)

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 2 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 simple digraph:

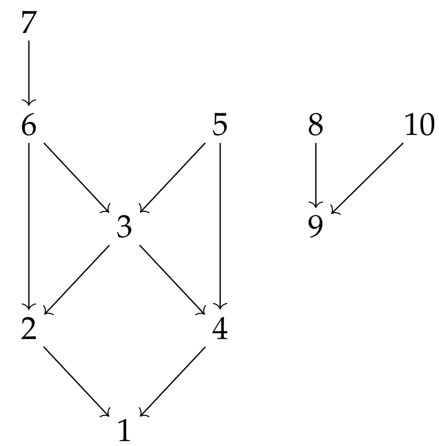

. (5)

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 (4) 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 is acyclic). 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 2. Let D be an acyclic multidigraph. 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 $u \longrightarrow v$ 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.

The so-called *no-watershed condition* says that 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 $w \overset{*}{\longrightarrow} t$.

Assume that the no-watershed condition holds. Prove that for each vertex p of D, there exists **exactly one** sink q of D such that $p \overset{*}{\longrightarrow} q$.

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

Exercise 2 is a well-known result rewritten in the language of digraphs. The result is known as *Newman's diamond lemma*, or, more precisely, its variant where the well-foundedness condition is replaced by finiteness and acyclicity. Note that the "no-watershed condition" is commonly called the *diamond condition* or *local confluence*. My talk [Grinbe17a] gives an introduction into the subject, including an outline of a solution to Exercise 2. Other references are collected in the MathOverflow post #289300 [MO289300]. Most authors working in this subject don't use the concept of multidigraphs, but instead speak of "abstract rewriting systems"; the vertices of D are then called the "objects", the relation $\longrightarrow$ is called the "reduction relation", and the sinks of D are called the "normal forms". The claim of Exercise 2 (and its more general version, in which D may have infinitely many vertices and arcs, but we requires the extra condition that there exist no infinite paths $v_0 \longrightarrow v_1 \longrightarrow v_2 \longrightarrow \cdots$) is used, e.g., in the theory of Gröbner bases and in the abstract study of computation (e.g., the Church-Rosser property for λ-calculus).

0.3. Arborescences of a wheel

Definition 0.3. Let D be a digraph, and let u be a vertex of D.

(a) Then, D is called a *u-arborescence* if and only if for each vertex v of D, there is a **unique** walk from u to v in D.

(b) Assume that D is a multidigraph (V, A, ϕ). A *u-arborescence of D* means a subset B of A such that the multidigraph $(V, B, \phi\mid_B)$ is a u-arborescence.

I believe that what I just called a "u-arborescence" is the same as what Vic called "arborescence with root u", except that maybe the arcs are pointing in the opposite direction. My notion of a "u-arborescence" is equivalent to what is called a "directed tree with root u" in [Sahi13], but the equivalence is not entirely obvious to prove.

For an example, the simple digraph

$$\begin{array}{ccc} 1 & \longrightarrow & 2 \\ \downarrow & & \\ 3 & \longrightarrow & 5 \\ \downarrow & & \downarrow \\ 4 & & 6 \end{array}$$

is a 1-arborescence. Thus, the set

$$\{(1,2), (1,3), (3,4), (3,5), (5,6)\}$$

is a 1-arborescence of the simple digraph

$$\begin{array}{ccc} 1 & \longrightarrow & 2 \\ \downarrow & & \downarrow \\ 3 & \longrightarrow & 5 \\ \downarrow & & \downarrow \\ 4 & \longrightarrow & 6 \end{array}.$$

Lemma 0.4. Let D be a digraph. Let u be a vertex of D. Let B be a u-arborescence of D.

(a) No arc in B has target u.

(b) Let v be a vertex of D distinct from u. Then, exactly one arc in B has target v.

Proof of Lemma 0.4 (sketched). Write the digraph D in the form $D = (V, A, \phi)$. The multidigraph $(V, B, \phi\mid_B)$ is a u-arborescence (since B is a u-arborescence of D). In other words, for each vertex $v \in V$, there is a **unique** walk from u to v in $(V, B, \phi\mid_B)$. Let us denote this walk by $\mathbf{w}_v$. Of course, $\mathbf{w}_v$ is a walk in the digraph D as well (since $B \subseteq A$).

(a) Assume the contrary. Thus, some arc $b \in B$ has target u. Consider this b. Let v be the source of this arc b. We can extend the walk $\mathbf{w}_v$ (which is a walk from

u to v in $(V, B, \phi\mid_B))$ by the arc b (which is an arc from v to u in $(V, B, \phi\mid_B))$ and the vertex u. The result is a walk from u to u in $(V, B, \phi\mid_B)$ having nonzero length. But there also exists a walk from u to u in $(V, B, \phi\mid_B)$ having zero length (namely, the walk (u)). Thus, we have found **two different** walks from u to u in $(V, B, \phi\mid_B)$. This contradicts the fact that there is a **unique** walk from u to u in $(V, B, \phi\mid_B)$, namely $\mathbf{w}_u$ (because this is how $\mathbf{w}_u$ was defined). This contradiction shows that our assumption was false. Thus, Lemma 0.4 **(a)** is proven.

(b) Recall that $\mathbf{w}_v$ is a walk from u to v in $(V, B, \phi\mid_B)$. This walk $\mathbf{w}_v$ has at least one arc (since v is distinct from u), and thus has a last arc. Let b be this last arc. Thus, b is an arc in B that has target v. Hence, **at least one** arc in B has target v. But we must prove that **exactly one** arc in B has target v. Thus, it remains to prove that b is the **only** arc in B that has target v.

Indeed, let c be any arc in B that has target v. We must prove that $c = b$.

Let w be the source of the arc c. Then, $\mathbf{w}_w$ is a walk from u to w in $(V, B, \phi\mid_B)$ (by the definition of $\mathbf{w}_w$). We can extend this walk $\mathbf{w}_w$ (which is a walk from u to w in $(V, B, \phi\mid_B))$ by the arc c (which is an arc from w to v in $(V, B, \phi\mid_B))$ and the vertex v. The result is a walk from u to v in $(V, B, \phi\mid_B)$ whose last arc is c. This walk must therefore be $\mathbf{w}_v$ (because $\mathbf{w}_v$ was defined to be the **unique** walk from u to v in $(V, B, \phi\mid_B))$. Hence, c (being the last arc of this walk) must be the last arc of the walk $\mathbf{w}_v$. But we already know that b is the last arc of the walk $\mathbf{w}_v$ (since this is how b was defined). Hence, $c = b$.

Now, forget that we fixed c. We thus have shown that if c is any arc in B that has target v, then $c = b$. Thus, b is the **only** arc in B that has target v. As explained, this concludes the proof of Lemma 0.4 **(b)**. □

Lemma 0.5. Let $D = (V, A, \phi)$ be a multidigraph. Let u be a vertex of D. Let B be a u-arborescence of D. Then, the multidigraph $(V, B, \phi\mid_B)$ has no cycles.

Proof of Lemma 0.5 (sketched). Let $\mathbf{c}$ be a cycle of the multidigraph $(V, B, \phi\mid_B)$. We must derive a contradiction.

Let v be the starting point of $\mathbf{c}$. The multidigraph $(V, B, \phi\mid_B)$ is a u-arborescence (since B is a u-arborescence of D). Hence, there is a **unique** walk from u to v in $(V, B, \phi\mid_B)$ (by the definition of a u-arborescence). Let us denote this walk by $\mathbf{w}$. Now, concatenating this walk $\mathbf{w}$ with the cycle $\mathbf{c}$, we obtain a new walk from u to v in $(V, B, \phi\mid_B)$. This new walk must be $\mathbf{w}$ (because $\mathbf{w}$ is the **unique** walk from u to v in $(V, B, \phi\mid_B))$. But this contradicts the fact that this new walk is longer than $\mathbf{w}$ (since it was obtained by concatenating $\mathbf{w}$ with the cycle $\mathbf{c}$).

Forget that we fixed $\mathbf{c}$. We thus have obtained a contradiction for each cycle $\mathbf{c}$ of the multidigraph $(V, B, \phi\mid_B)$. Thus, the multidigraph $(V, B, \phi\mid_B)$ has no cycles. This proves Lemma 0.5. □

Notice that Lemma 0.4 has something like a converse, which we can use to characterize u-arborescences:

Lemma 0.6. Let $D = (V, A, \phi)$ be a multidigraph. Let B be a subset of A. Let u be a vertex of D. Assume the following:

- No arc in B has target u.
- For each vertex v of D distinct from u, there is exactly one arc in B that has target v.
- For each vertex v of D distinct from u, there is at least one walk from u to v in $(V, B, \phi\mid_B)$.

Then, B is a u-arborescence of D.

We leave the proof of this lemma to the reader.

Two much less trivial properties of arborescences are the following:

Theorem 0.7. Let $D = (V, A, \phi)$ be a multidigraph. For any two vertices i and j of D, we let $a_{i,j}$ be the number of arcs of D having source i and target j. For any vertex i of D, we let w_i be the number of i-arborescences of D. Then, each vertex i of D satisfies

$$\sum_{j \in V} a_{i,j} w_j = \left(\sum_{k \in V} a_{k,i} \right) w_i.$$

Theorem 0.7 is essentially [Sahi13, Theorem 1] (restated without using matrices, and with weighted arcs instead of parallel arcs: e.g., instead of having 4 arcs with source i and target j, [Sahi13, Theorem 1] would speak of an arc of weight 4 from i to j). Here is the main idea of its proof:

Hint to the proof of Theorem 0.7. For any two vertices i and j of D, we let $A_{i,j}$ be the set of all arcs of D having source i and target j. Note that this set satisfies $|A_{i,j}| = a_{i,j}$.

For any vertex i of D, we let W_i be the set of all i-arborescences of D. This set W_i satisfies $|W_i| = w_i$.

Now, let i be a vertex of D. A subset B of A is said to be *univalent* if for each $v \in V$, there is exactly one arc in B whose target is v. A subset B of A is said to be an *i-unicycle* if it has the following properties:

- The subset B is univalent.
- Each cycle of the digraph $(V, B, \phi\mid_B)$ contains the vertex i.
- The digraph $(V, B, \phi\mid_B)$ has exactly one cycle from i to i.

(Note that these properties imply that the digraph $(V, B, \phi\mid_B)$ has only one cycle, up to cyclic rotation. Therefore the name "i-unicycle".)

If B is an i-unicycle, then

- the unique cycle from i to i of the digraph $(V, B, \phi\mid_B)$ will be denoted by $z(B)$;
- the first arc of this cycle will be denoted by $f(B)$;
- the last arc of this cycle will be denoted by $\ell(B)$;
- the target of $f(B)$ will be denoted by $\mathfrak{j}(B)$;
- the source of $\ell(B)$ will be denoted by $\mathfrak{k}(B)$.

The following facts are important and easy to check:

> *Claim 1:* Let v be a vertex of D. Let β be an arc of D with target v such that i is either the source or the target of v. Let C be a v-arborescence of D. Then, $C \cup \{\beta\}$ is a i-unicycle.
>
> *Claim 2:* Let B be an i-unicycle. Let β be an arc of the cycle $z(B)$, and let v be the target of β. Then, $B \setminus \{\beta\}$ is a v-arborescence of D.

Now, we can prove that the map

$$\{i\text{-unicycles}\} \to \bigcup_{j\in V} \left(A_{i,j} \times W_j\right),$$
$$B \mapsto (f(B), B \setminus \{f(B)\})$$

is well-defined (i.e., most importantly, if B is an i-unicycle, then $B \setminus \{f(B)\}$ is a $\mathfrak{j}(B)$-arborescence of D (by Claim 2)) and bijective (indeed, its inverse map sends a pair $(\beta, C) \in A_{i,j} \times W_j$ for any $j \in V$ to the i-unicycle $C \cup \{\beta\}$ (which is well-defined by Claim 1)). Thus, it is a bijection from $\{i\text{-unicycles}\}$ to $\bigcup\limits_{j\in V} \left(A_{i,j} \times W_j\right)$. Hence,

$$\left|\{i\text{-unicycles}\}\right| = \left|\bigcup_{j\in V} \left(A_{i,j} \times W_j\right)\right| = \sum_{j\in V} \underbrace{\left|A_{i,j}\right|}_{=a_{i,j}} \cdot \underbrace{\left|W_j\right|}_{=w_j}$$
$$\left(\begin{array}{c}\text{since the sets } A_{i,j} \times W_j \text{ for different } j \in V \\ \text{are disjoint (because the sets } A_{i,j} \text{ for different } j \in V \\ \text{are disjoint)}\end{array}\right)$$
$$= \sum_{j\in V} a_{i,j} w_j. \qquad (6)$$

On the other hand, we can prove that the map

$$\{i\text{-unicycles}\} \to \left(\bigcup_{k\in V} A_{k,i}\right) \times W_i,$$
$$B \mapsto (\ell(B), B \setminus \{\ell(B)\})$$

is well-defined (i.e., most importantly, if B is an i-unicycle, then $B \setminus \{\ell(B)\}$ is an i-arborescence of D (by Claim 2)) and bijective (indeed, its inverse map sends a pair $(\beta, C) \in A_{k,i} \times W_i$ for any $k \in V$ to the i-unicycle $C \cup \{\beta\}$ (which is well-defined by Claim 1)). Thus, it is a bijection from $\{i\text{-unicycles}\}$ to $\left(\bigcup_{k \in V} A_{k,i}\right) \times W_i$. Hence,

$$\begin{aligned}
|\{i\text{-unicycles}\}| &= \left|\left(\bigcup_{k\in V} A_{k,i}\right) \times W_i\right| = \left(\sum_{k\in V} \underbrace{|A_{k,i}|}_{=a_{k,i}}\right) \cdot \underbrace{|W_i|}_{=w_i} \\
&\qquad \left(\begin{array}{c}\text{since the sets } A_{k,i} \text{ for different } k \in V \\ \text{are disjoint}\end{array}\right) \\
&= \left(\sum_{k\in V} a_{k,i}\right) w_i .
\end{aligned}$$

Comparing this with (6), we obtain $\sum_{j\in V} a_{i,j} w_j = \left(\sum_{k\in V} a_{k,i}\right) w_i$. This proves Theorem 0.7. □

Definition 0.8. Let $D = (V, A, \phi)$ be a multidigraph. Let v be a vertex of D.

(a) The *indegree* $\operatorname{indeg} v$ of v denotes the number of arcs $a \in A$ having target v.

(b) The *outdegree* $\operatorname{outdeg} v$ of v denotes the number of arcs $a \in A$ having source v.

Theorem 0.9. Let D be a multidigraph. For any vertex i of D, we let w_i be the number of i-arborescences of D. Assume that each vertex v of D satisfies $\operatorname{indeg} v = \operatorname{outdeg} v$. Then, $w_i = w_j$ for any two vertices i and j of D.

Theorem 0.9 appears in various sources (e.g., [Berge91, Chapter 11, §3, Corollary 1]), and can be derived from the BEST theorem (see, e.g., [Klings11, Theorem 2][4] or [Bollob98, §I.3, Theorem 13][5] or [JacGou79, Corollary 4.1]). The main idea is that if D is a multidigraph with the property that each vertex of D satisfies $\operatorname{indeg} v =$

[4]Klingsberg's [Klings11, §2] gives a self-contained proof of the BEST theorem. If you want to read it, keep in mind that the condition that D be "connected" (I am not sure whether he means "strongly connected" or "weakly connected" by this word) is never used in the proof and thus can (and should) be removed from the statement of [Klings11, Theorem 2]. Also, keep in mind that he works with "trees flowing into" a vertex u; this is almost the same as u-arborescences, but instead of requiring a unique walk from u to any vertex v, he requires a unique walk from any vertex v to u. Thus, of course, the concept of a tree flowing into u becomes the concept of a u-arborescence once you reverse each arc of your digraph (i.e., swap the source with the target). Also, Klingsberg defines a circuit to be an equivalence class of circuits up to cyclic rotation.

[5]When reading Bollobás's [Bollob98, §I.3], keep in mind that he works with "spanning trees oriented towards" a vertex u; this is almost the same as u-arborescences, but instead of requiring a unique walk from u to any vertex v, he requires a unique walk from any vertex v to u. This is exactly what Klingsberg calls a "tree flowing into" u. Also, Bollobás counts cyclic rotations of a circuit as identical.

outdeg v, then the BEST theorem shows that the total number of Eulerian circuits of D (counted up to cyclic rotation) equals

$$w_i \cdot \prod_{v \in V} (\operatorname{outdeg} v - 1)!$$

for any given $i \in V$. Since this number clearly doesn't depend on i, we thus conclude that any two vertices i and j of such a multidigraph D satisfy

$$w_i \cdot \prod_{v \in V} (\operatorname{outdeg} v - 1)! = w_j \cdot \prod_{v \in V} (\operatorname{outdeg} v - 1)!,$$

and therefore $w_i = w_j$, which proves Theorem 0.9.

Exercise 3. Let m be a positive integer. Let W_m be the simple digraph with $m+1$ vertices $0, 1, \ldots, m$ and the following $3m$ arcs:

$$\begin{aligned} &(1,2), (2,3), \ldots, (m-1, m), (m, 1); \\ &(0,i) \qquad \text{for each } i \in [m]; \\ &(i,0) \qquad \text{for each } i \in [m]. \end{aligned}$$

(Visually speaking, W_m consists of a cycle that traverses m vertices $1, 2, \ldots, m$, as well as a "center vertex" 0 which is joined to each of these m vertices by one edge in each direction. For example, here is how W_6 looks like:

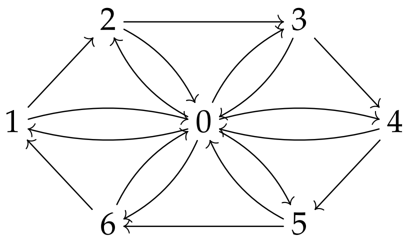

.

And here is a 0-arborescence of W_6:

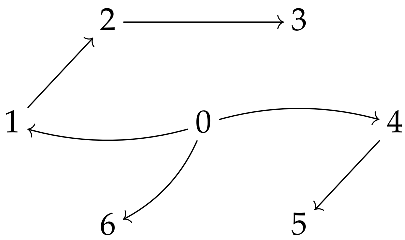

, (7)

where of course we draw a 0-arborescence B by drawing the digraph (V, B) with V being the vertex set of W_m.)

(a) Compute the number of 0-arborescences of W_m.

(b) Let $i \in [m]$. Compute the number of i-arborescences of W_m.

[For example, if $m = 3$, then both answers are 7, and the 0-arborescenses of W_3 are

$$\begin{aligned} &\{(0,1), (0,2), (0,3)\}, && \{(0,1), (0,2), (2,3)\}, \\ &\{(0,1), (0,3), (1,2)\}, && \{(0,1), (1,2), (2,3)\}, \\ &\{(0,2), (0,3), (3,1)\}, && \{(0,2), (2,3), (3,1)\}, \\ &\{(0,3), (1,2), (3,1)\}. \end{aligned}$$

]

Solution to Exercise 3 (sketched). Let V be the set of all vertices of W_m. We classify the $3m$ arcs of W_m into the following three classes:

- The arcs $(1,2), (2,3), \ldots, (m-1, m), (m, 1)$ of W_m will be called the *circle-*

arcs. Note that there are m of these arcs. Note also that for any two vertices $i, j \in [m]$, there is exactly one path from i to j that uses only circle-arcs. This path will be denoted by $P_{i,j}$.

- The arcs $(0,1), (0,2), \ldots, (0,m)$ of W_m will be called the *centrifugal arcs*.
- The arcs $(1,0), (2,0), \ldots, (m,0)$ of W_m will be called the *centripetal arcs*.

Also, if $i \in [m]$, then i^- shall denote the source of the unique circle-arc whose target is i. Thus, $i^- = i - 1$ when $i > 1$, whereas $1^- = m$.

Also, if $i \in [m]$, then i^+ shall denote the source of the unique circle-arc whose source is i. Thus, $i^+ = i + 1$ when $i < m$, whereas $m^+ = 1$.

In the following, if u is a vertex of W_m, then "u-arborescence" will always mean "u-arborescence of W_m".

(a) This number is $2^m - 1$.

Proof. If S is any nonempty subset of $[m]$, then the *S-spider* shall mean the subset of A consisting of the following arcs:

- the k centrifugal arcs $(0,i)$ for all $i \in S$;
- the $m - k$ circle-arcs (i^-, i) for all $i \notin S$.

For example, here is how the $\{1,2,4\}$-spider in W_6 looks like:

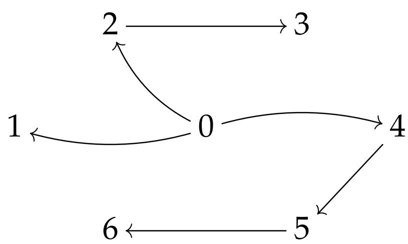

.

Also, the 0-arborescence of W_6 shown in (7) is actually the $\{1,4,6\}$-spider.

It is easy to see that for every nonempty subset S of $[m]$, the S-spider is a 0-arborescence. Thus, we have found $2^m - 1$ many 0-arborescences (one for each nonempty subset S of $[m]$), and they are all distinct (indeed, we can recover S from the S-spider by looking at the centrifugal arcs appearing in the S-spider). Thus, the number of 0-arborescences of W_m is at least $2^m - 1$. In order to complete this proof, we only need to check that these are the only 0-arborescences. In other words, we only need to check that each 0-arborescence is an S-spider for some nonempty subset S of $[m]$.

Let B be a 0-arborescence. We want to show that B is an S-spider for some nonempty subset S of $[m]$. Indeed, let S be the set of all $i \in [m]$ such that B contains the centrifugal arc $(0,i)$. Then, S is a nonempty subset of $[m]$ (indeed, it must be nonempty, since some centrifugal arc needs to be used in order to escape 0). Clearly, B cannot contain any centripetal arcs (by Lemma 0.4 **(a)**, applied to $D = W_m$ and $u = 0$). Moreover, for each $i \in S$, the circle-arc (i^-, i) is not contained in B (since that would cause B to contain **two** different arcs with target i, which would contradict Lemma 0.4 **(b)**). Meanwhile, for each $i \in [m] \setminus S$, the circle-arc (i^-, i) is contained in B (by Lemma 0.4 **(b)** again, because the centrifugal arc $(0, i)$ is not in B). This shows that the arcs contained in B are precisely the arcs that are contained in the S-spider. Hence, B is the S-spider. This proves what we wanted to show. Thus, Exercise 3 **(a)** is solved.

(b) This number is $2^m - 1$ as well.

Proof. This can be shown similarly to how we proved part **(a)** above, but the argument will be much messier due to the asymmetry of the situation. Instead, here are two different proofs, using Theorem 0.7 and Theorem 0.9 instead.

First proof of the fact that the number of i-arborescences B of W_m is $2^m - 1$*:* For any vertex u of D, we let w_u be the number of u-arborescences of D. Then, Theorem 0.7 shows that

$$\sum_{j \in V} a_{i,j} w_j = \left(\sum_{k \in V} a_{k,i} \right) w_i. \tag{8}$$

But the digraph W_m has a rotational symmetry: Specifically, there is a digraph isomorphism from W_m to W_m that sends the vertices $0, 1, 2, \ldots, m$ to $0, i, i+1, \ldots, m, 1, 2, \ldots, i-1$, respectively. (Visually speaking, it just rotates the circumference of the circle so that vertex 1 falls onto vertex i.) Thus, $w_i = w_1$. Similarly, $w_j = w_1$ for each $j \in [m]$. In other words, $w_1 = w_2 = \cdots = w_m$. Now, (8) yields

$$\begin{aligned} \left(\sum_{k \in V} a_{k,i} \right) w_i &= \sum_{j \in V} a_{i,j} w_j \\ &= \underbrace{a_{i,0}}_{=1} \underbrace{w_0}_{\substack{=2^m - 1 \\ \text{(by our answer to} \\ \text{part (a))}}} + \underbrace{a_{i,i^+}}_{=1} \underbrace{w_{i^+}}_{\substack{=w_i \\ \text{(since } w_1 = w_2 = \cdots = w_m)}} + \sum_{\substack{j \in V; \\ j \notin \{0, i^+\}}} \underbrace{a_{i,j}}_{=0} w_j \\ &= (2^m - 1) + w_i, \end{aligned}$$

so that

$$(2^m - 1) + w_i = \underbrace{\left(\sum_{k \in V} a_{k,i} \right)}_{=2} w_i = 2 w_i.$$

Subtracting w_i from this equality, we find $2^m - 1 = w_i$, so that $w_i = 2^m - 1$, qed.

Second proof of the fact that the number of i-arborescences B of W_m is $2^m - 1$*:* Each vertex v of W_m satisfies $\operatorname{indeg} v = \operatorname{outdeg} v$ (namely, both $\operatorname{indeg} v$ and $\operatorname{outdeg} v$ equal $\begin{cases} m, & \text{if } v = 0; \\ 2, & \text{if } v \neq 0 \end{cases}$). For any vertex u of W_m, we let w_u be the number of u-arborescences of W_m. Then, Theorem 0.9 shows that $w_i = w_j$ for every vertex j of D. Applying this to $j = 0$, we obtain $w_i = w_0 = 2^m - 1$ (because in part **(a)** of this exercise, we have proven that the number of 0-arborescences of W_m is $2^m - 1$). Qed.

□

0.4. Back to undirected graphs

In the following exercises, we will use the following definitions:

Definition 0.10. For each $n \in \mathbb{N}$, we define the *n-th path graph* to be the simple graph

$$\begin{aligned} & (\{1,2,\ldots,n\}, \{\{i,i+1\} \mid i \in \{1,2,\ldots,n-1\}\}) \\ & = (\{1,2,\ldots,n\}, \{\{1,2\}, \{2,3\}, \ldots, \{n-1,n\}\}). \end{aligned}$$

This graph is denoted by P_n. It has n vertices and $n-1$ edges (unless $n = 0$, in which case it has 0 edges). Here is a drawing of P_4:

$$1 \text{ --- } 2 \text{ --- } 3 \text{ --- } 4 .$$

Definition 0.11. For each integer $n > 1$, we define the *n-th cycle graph* to be the simple graph

$$\begin{aligned} & (\{1,2,\ldots,n\}, \{\{i,i+1\} \mid i \in \{1,2,\ldots,n-1\}\} \cup \{n,1\}) \\ & = (\{1,2,\ldots,n\}, \{\{1,2\}, \{2,3\}, \ldots, \{n-1,n\}, \{n,1\}\}). \end{aligned}$$

This graph is denoted by C_n. It has n vertices and $\begin{cases} n, & \text{if } n \geq 3; \\ 1, & \text{if } n = 2 \end{cases}$ edges. Here is a drawing of C_6:

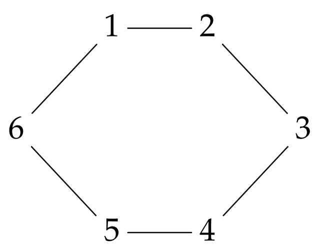

.

0.5. Chromatic polynomials of complete bipartite graphs

For the definition and properties of the chromatic polynomial of a simple graph, see Exercise 4 on Spring 2017 Math 5707 midterm #2. In a nutshell:

- If $G = (V, E)$ is a simple graph, then the *chromatic polynomial* χ_G of G is a polynomial in a single indeterminate x (with integer coefficients) defined by

$$\chi_G = \sum_{F \subseteq E} (-1)^{|F|} x^{\operatorname{conn}(V,F)}.$$

  (Here, as usual, $\operatorname{conn} H$ denotes the number of connected components of any graph H.)

We could easily extend this definition to multigraphs, but I don't want to carry the extra notation around.

- The main property of chromatic polynomials in the following: If $G = (V, E)$ is a graph, and if $k \in \mathbb{N}$, then $\chi_G(k)$ is the number of all proper k-colorings[6] of G. Note that this property uniquely determines χ_G, because any polynomial in a single indeterminate x is uniquely determined by its values on all nonnegative integers.

- For any $n \in \mathbb{N}$, the complete graph K_n (that is, the simple graph with n vertices $1, 2, \ldots, n$ and all possible edges $\{i, j\}$ with $1 \leq i < j \leq n$) has chromatic polynomial $\chi_{K_n} = x(x-1)\cdots(x-n+1)$.

- If T is a tree with n vertices, then $\chi_T = x(x-1)^{n-1}$. Thus, in particular, for any positive integer n, the path graph P_n (see Definition 0.10) has characteristic polynomial $\chi_{P_n} = x(x-1)^{n-1}$ (since it is a tree with n vertices).

- If $n > 1$ is an integer, then the chromatic polynomial of the cycle graph C_n (see Definition 0.11) is $\chi_{C_n} = (x-1)^n + (-1)^n(x-1)$. (This is Exercise 2 **(a)** on Spring 2017 Math 5707 midterm #3.)

Definition 0.12. Let $n \in \mathbb{N}$ and $m \in \mathbb{N}$. The graph $K_{n,m}$ is defined to be the simple graph with $n + m$ vertices

$$1, 2, \ldots, n, -1, -2, \ldots, -m$$

and nm edges

$$\{i, -j\} \qquad \text{for all } i \in [n] \text{ and } j \in [m].$$

(Note that $(K_{n,m}; \{1, 2, \ldots, n\}, \{-1, -2, \ldots, -m\})$ is a bipartite graph, called the *complete bipartite graph*.)

For example, the graph $K_{2,3}$ is

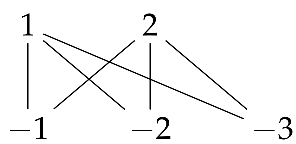

.

Exercise 4. Let $n \in \mathbb{N}$ and $m \in \mathbb{N}$.

(a) Prove that the chromatic polynomial of $K_{n,m}$ is

$$\chi_{K_{n,m}} = \sum_{i=0}^{n} \operatorname{sur}(n, i) \binom{x}{i} (x-i)^m.$$

[6]Recall that a *k-coloring* of G means a map $f : V \to \{1, 2, \ldots, k\}$. (The image $f(v)$ of a vertex $v \in V$ under this map is called the *color* of v under this k-coloring f.) A k-coloring f of G is said to be *proper* if each edge $\{u, v\}$ of G satisfies $f(u) \neq f(v)$. (In other words, a k-coloring f of G is proper if and only if no two adjacent vertices share the same color.)

(Recall that $\operatorname{sur}(n,i)$ denotes the number of all surjections from $[n]$ to $[i]$.)

(b) Prove that

$$\sum_{i=0}^{n} \operatorname{sur}(n,i) \binom{x}{i} (x-i)^m = \sum_{i=0}^{m} \operatorname{sur}(m,i) \binom{x}{i} (x-i)^n .$$

Remark 0.13. Applying Exercise 4 **(b)** to $n = 0$, we recover the identity

$$x^m = \sum_{i=0}^{m} \operatorname{sur}(m,i) \binom{x}{i},$$

which was Theorem 3.15 in the class of February 21.

0.6. Counting independent sets

Now let us return to independent sets of graphs.

Definition 0.14. Let G be a graph.

(a) An *independent set* of G means a set S of vertices of G such that no two distinct elements of S are adjacent.

(b) We let $\operatorname{ind} G$ be the number of all independent sets of G.

There are no good formulas for $\operatorname{ind} G$ in general, but we can always try to compute it when G is a particularly simple type of graph. For example, if G is the path graph P_n for some $n \in \mathbb{N}$, then the independent sets of G are precisely the lacunar subsets of $[n]$, and thus $\operatorname{ind} G$ is the Fibonacci number f_{n+2} (by Proposition 1.22 in the February 5 class). The independent sets of the cycle graph C_n are the lacunar subsets of $[n]$ which don't contain 1 and n simultaneously (i.e., they can contain at most one of 1 and n). The following definition will help counting them:

Definition 0.15. Let G be a graph. Let S be a set of vertices of G. Then, $G \setminus S$ will denote the graph obtained from G by removing all vertices in S (along with all edges that use these vertices).

(More rigorously: If G is a simple graph (V, E), then $G \setminus S$ is the simple graph $(V \setminus S, E')$, where E' is the set of all edges $e \in E$ such that no endpoint of e belongs to S. If G is a multigraph (V, E, φ), then $G \setminus S$ is the multigraph $(V \setminus S, E', \varphi\mid_{E'})$, where E' is the set of all edges $e \in E$ such that no endpoint of e belongs to S.)

For example, if G is the graph

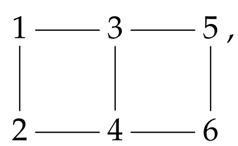

then $G \setminus \{2,3\}$ is the graph

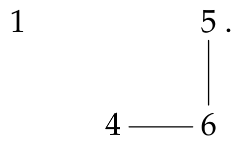

Exercise 5. **(a)** Let v be a vertex of a graph G. Let $N(v)$ be the set of all neighbors of v. Let $N^+(v) = \{v\} \cup N(v)$. Prove that

$$\operatorname{ind} G = \operatorname{ind}(G \setminus \{v\}) + \operatorname{ind}\left(G \setminus \left(N^+(v)\right)\right).$$

(b) Compute $\operatorname{ind}(C_n)$ for each $n \geq 2$ (in terms of the Fibonacci sequence).

Remark 0.16. **(a)** It is instructive to see what Exercise 5 **(a)** says when G is a path graph. Let $n > 1$ be an integer, and let G be the path graph P_n. Let $v \in [n]$ (so that v is a vertex of G). If $v = n$, then $G \setminus \{v\} = P_{n-1}$ and $G \setminus (N^+(v)) = P_{n-2}$ (since $N^+(v) = \{n, n-1\}$ in this case), so that Exercise 5 **(a)** yields

$$\operatorname{ind}(P_n) = \operatorname{ind}(P_{n-1}) + \operatorname{ind}(P_{n-2}).$$

This is precisely the recurrence equation of the Fibonacci numbers. Thus, we obtain a new (inductive) proof of the fact that

$$\operatorname{ind}(P_n) = f_{n+2} \qquad \text{for each } n \in \mathbb{N}. \tag{9}$$

However, we can also apply Exercise 5 **(a)** to another vertex v. Let $v \in \{2, 3, \ldots, n-1\}$. Then, $G \setminus \{v\}$ is a disconnected graph looking as follows:

$$1 — 2 — \cdots — v-1 \qquad v+1 — v+2 — \cdots — n\,.$$

We can treat this graph as a "disjoint union" of two path graphs, one of which is P_{v-1} while the other is P_{n-v} "in all but name" (its vertices are called $v+1, v+2, \ldots, n$ rather than $1, 2, \ldots, n-v$, but otherwise it is identical to P_{n-v}). To construct an independent set of $G \setminus \{v\}$, we thus just need to choose an independent set of the former path graph P_{v-1} and an independent set of the latter path grad P_{n-v} (with vertices renamed as $v+1, v+2, \ldots, n$), and take the union of these two independent sets. Hence,

$$\operatorname{ind}(G \setminus \{v\}) = \operatorname{ind}(P_{v-1}) \cdot \operatorname{ind}(P_{n-v}).$$

A similar argument shows that the graph $G \setminus (N^+(v))$ has the form

$$1 — 2 — \cdots — v-2 \qquad v+2 — v+3 — \cdots — n$$

(again, a "disjoint union" of two path graphs, which this time are P_{v-2} and P_{n-v-1}), and thus

$$\operatorname{ind}\left(G \setminus \left(N^+(v)\right)\right) = \operatorname{ind}(P_{v-2}) \cdot \operatorname{ind}(P_{n-v-1}).$$

Hence, Exercise 5 **(a)** becomes

$$\operatorname{ind} G = \underbrace{\operatorname{ind}\left(G \setminus \{v\}\right)}_{=\operatorname{ind}(P_{v-1})\cdot\operatorname{ind}(P_{n-v})} + \underbrace{\operatorname{ind}\left(G \setminus \left(N^{+}(v)\right)\right)}_{=\operatorname{ind}(P_{v-2})\cdot\operatorname{ind}(P_{n-v-1})}$$
$$= \operatorname{ind}(P_{v-1}) \cdot \operatorname{ind}(P_{n-v}) + \operatorname{ind}(P_{v-2}) \cdot \operatorname{ind}(P_{n-v-1}).$$

Since $G = P_n$, this rewrites as

$$\operatorname{ind}(P_n) = \operatorname{ind}(P_{v-1}) \cdot \operatorname{ind}(P_{n-v}) + \operatorname{ind}(P_{v-2}) \cdot \operatorname{ind}(P_{n-v-1}).$$

Using the equality (9), we can rewrite this as

$$f_{n+2} = f_{(v-1)+2} \cdot f_{(n-v)+2} + f_{(v-2)+2} \cdot f_{(n-v-1)+2} = f_{v+1} f_{n-v+2} + f_v f_{n-v+1}.$$

Applying this to $v = a$ and $n = a + b - 1$, we conclude that

$$f_{a+b+1} = f_{a+1} f_{b+1} + f_a f_b \qquad \text{for all } a \in \mathbb{N} \text{ and } b \in \mathbb{N}.$$

(To be more precise, we have only proven this in the case when $a > 1$ and $b > 1$; but all other cases are easy.) Thus, we have recovered the claim of Exercise 3 **(e)** on midterm #1.

(b) We can similarly count independent sets **of a given size**. For C_n, we find that the number of independent sets of C_n having size k (for a given $k \in \{0, 1, \ldots, n-1\}$) is $\dfrac{n}{n-k}\dbinom{n-k}{k}$. This can also be proven combinatorially; see [Stan11, Lemma 2.3.4].

References

[Berge91] Claude Berge, *Graphs*, 3rd revised edition, North-Holland 1991.

[Bollob98] Béla Bollobás, *Modern Graph Theory*, Graduate Texts in Mathematics #184, 1st edition, Springer 1998.

[Grinbe17a] Darij Grinberg, *The diamond lemma and its applications (talk slides)*, `http://www.cip.ifi.lmu.de/~grinberg/algebra/diamond-talk.pdf`

[JacGou79] D. M. Jackson, I. P. Goulden, *Sequence enumeration and the de Bruijn-van Aardenne-Ehrenfest-Smith-Tutte theorem*, Canad. J. Math. 31(1979), pp. 488–495, `https://doi.org/10.4153/CJM-1979-054-x`

[MO289300] Darij Grinberg, *MathOverflow post #289300*, `https://mathoverflow.net/a/289320/`

[Klings11] Paul Klingsberg, *The BEST Theorem*, `http://people.sju.edu/~pklingsb/best.pdf`

[Sahi13] Siddhartha Sahi, *Harmonic vectors and matrix tree theorems*, arXiv:1309.4047v1.

[Stan11] Richard Stanley, *Enumerative Combinatorics, volume 1*, Second edition, version of 15 July 2011. Available at `http://math.mit.edu/~rstan/ec/` .