Math 530 Spring 2022, Lecture 18: Trees

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

1. Trees and arborescences (cont'd)

1.1. Spanning arborescences vs. spanning trees

Last time, we proved the BEST theorem, which connected the # of Eulerian circuits in a digraph with the # of spanning arborescences of the same digraph. Now we are looking for a way to compute the latter.

For example, let us try to do this for digraphs of the form G^{bidir} where G is a multigraph. I claim that the spanning arborescences of G^{bidir} rooted to a given vertex r are just the spanning trees of G in disguise:

Proposition 1.1.1. Let $G = (V, E, \varphi)$ be a multigraph. Fix a vertex $r \in V$. Recall that the arcs of G^{bidir} are the pairs $(e, i) \in E \times \{1, 2\}$. Identify each spanning tree of G with its edge set, and each spanning arborescence of G^{bidir} with its arc set.

If B is a spanning arborescence of G^{bidir} rooted to r, then we set

$$\overline{B} := \{e \mid (e, i) \in B\}.$$

(Recall that we are identifying spanning arborescences with their arc sets, so that "$(e, i) \in B$" means "(e, i) is an arc of B".)

Then:

(a) If B is a spanning arborescence of G^{bidir} rooted to r, then $\overline{B}$ is a spanning tree of G.

(b) The map

$$\left\{\text{spanning arborescences of } G^{\text{bidir}} \text{ rooted to } r\right\} \to \{\text{spanning trees of } G\}, \qquad B \mapsto \overline{B}$$

is a bijection.

Example 1.1.2. Here is a multigraph G (on the left) with the corresponding multidigraph G^{bidir} (on the right):

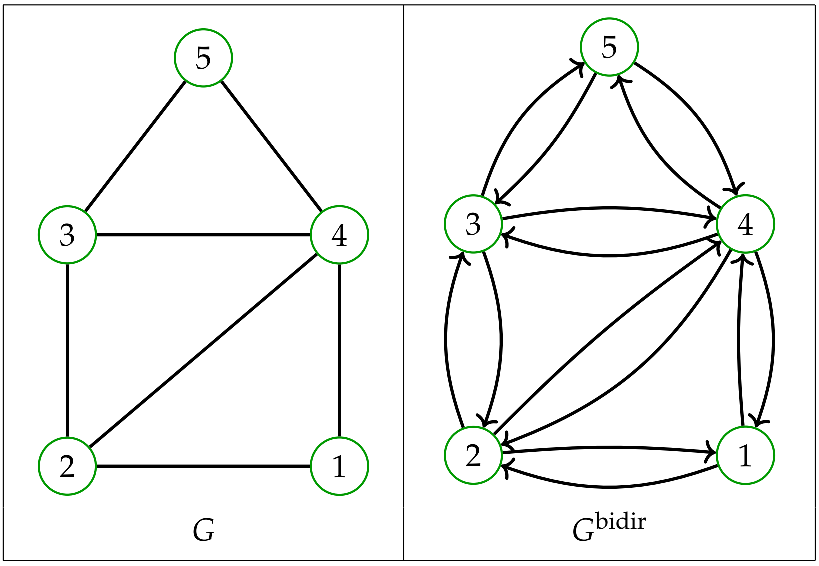

Here is a spanning arborescence B of G^{bidir} rooted to 1, and the corresponding spanning tree $\overline{B}$ of G:

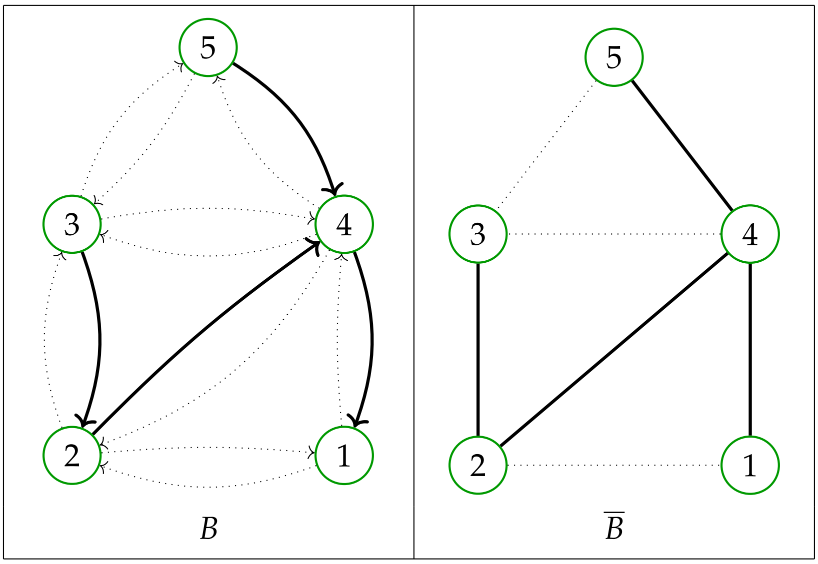

(here, the arcs of G^{bidir} that **don't** belong to B, as well as the edges of G that **don't** belong to $\overline{B}$, have been drawn as dotted arrows). It is fairly easy to see how B can be reconstructed from $\overline{B}$: You just need to replace each edge of $\overline{B}$ by the appropriately directed arc (namely, the one that is "directed towards 1").

Proof of Proposition 1.1.1. This is an exercise in yak-shaving (and we have, in fact, shaved a very similar yak in Lecture 16; the only difference is that we are no longer dealing with trees in isolation, but rather with spanning trees of G).

(a) Let B be a spanning arborescence of G^{bidir} rooted to r. Then, B^{und} is a tree (by the implication A'1$\Longrightarrow$A'3 in Theorem 1.1.3 from Lecture 17). However, it is easy to see that $B^{\text{und}} \cong \overline{B}$ as multigraphs (indeed, each vertex v of B^{und} corresponds to the same vertex v of $\overline{B}$, whereas any edge (e,i) of B^{und} corresponds to the edge e of $\overline{B}$) [1]. Thus, $\overline{B}$ is a tree (since B^{und} is a tree)[2], therefore a spanning tree of G (since $\overline{B}$ is clearly a spanning subgraph of G). This proves Proposition 1.1.1 **(a)**.

(b) We must prove that this map is surjective and injective.

Surjectivity: Let T be a spanning tree of G. Then, the multidigraph $T^{r\rightarrow}$ (defined in Definition 1.1.6 in Lecture 16) is an arborescence rooted from r (by Lemma 1.1.9 in Lecture 16). Reversing each arc in this arborescence $T^{r\rightarrow}$, we obtain a new multidigraph $T^{r\leftarrow}$, which is thus an arborescence rooted to r. Unfortunately, $T^{r\leftarrow}$ is not a subdigraph of G^{bidir}, for a rather stupid reason: The arcs of $T^{r\leftarrow}$ are elements of E, whereas the arcs of G^{bidir} are pairs of the form (e,i) with $e \in E$ and $i \in \{1,2\}$.

Fortunately, this is easily fixed: For each arc e of $T^{r\leftarrow}$, we let e' be the arc (e,i) of G^{bidir} that has the same source as e (and thus the same target as e). This is uniquely determined, since the arcs $(e,1)$ and $(e,2)$ of G^{bidir} have different sources[3]. If we replace each arc e of $T^{r\leftarrow}$ by the corresponding arc e' of G^{bidir}, then we obtain a spanning subdigraph S of G^{bidir} that is an arborescence rooted to r (since $T^{r\leftarrow}$ is an arborescence rooted to r, and we have only replaced its arcs by equivalent ones with the same sources and the same targets). In other words, we obtain a spanning arborescence S of G^{bidir} rooted to r. It is easy to see that $\overline{S} = T$. Hence, the map

$$\left\{\text{spanning arborescences of } G^{\text{bidir}} \text{ rooted to } r\right\} \to \{\text{spanning trees of } G\}, \quad B \mapsto \overline{B}$$

[1] Here we need to use the fact that for each edge e of $\overline{B}$, **exactly one** of the two pairs $(e,1)$ and $(e,2)$ is an edge of B^{und}. But this is easy to check: At least one of the two pairs $(e,1)$ and $(e,2)$ must be an arc of B (since e is an edge of $\overline{B}$). In other words, at least one of the two pairs $(e,1)$ and $(e,2)$ must be an edge of B^{und}. But both of these pairs cannot be edges of B^{und} at the same time (since this would create a cycle, but B^{und} is a tree and thus has no cycles). Hence, exactly one of these pairs is an edge of B^{und}, qed.

[2] Alternatively, you can prove this as follows: The vertex r is a to-root of B (since B is an arborescence rooted to r). Thus, for each $v \in V$, there is a path from v to r in B. By "projecting" this path onto $\overline{B}$ (that is, replacing each arc (e,i) of this path by the corresponding edge e of $\overline{B}$), we obtain a path from v to r in $\overline{B}$. This shows that the multigraph $\overline{B}$ is connected. Furthermore, the definition of $\overline{B}$ shows that $\left|\overline{B}\right| \leq |B| = |V| - 1$ (by Statement A'2 in Theorem 1.1.3 in Lecture 17, since B is an arborescence rooted to r). Hence, $\left|\overline{B}\right| < |V|$. Thus, we can apply the implication T5$\Longrightarrow$T1 of the Tree Equivalence Theorem (Theorem 1.2.4 in Lecture 14) to conclude that $\overline{B}$ is a tree.

[3] *Proof.* The edge e of T is not a loop (because T is a tree, but a tree cannot have any loops). Hence, its two endpoints are distinct. Thus, the arcs $(e,1)$ and $(e,2)$ of G^{bidir} have different sources (since their sources are the two endpoints of e).

sends S to T. This shows that T is a value of this map. Since we have proved this for every spanning tree T of G, we have thus shown that this map is surjective.

Injectivity: The main idea is that, in order to recover a spanning arborescence B back from the corresponding spanning tree $\overline{B}$, we just need to "orient the edges of the tree towards r". Here are the (annoyingly long) details:

Let B and C be two sparbs[4] such that $\overline{B} = \overline{C}$. We must show that $B = C$.

Assume the contrary. Thus, $B \neq C$. Let T be the tree $\overline{B} = \overline{C}$. Thus, each edge e of T corresponds to either an arc $(e, 1)$ or an arc $(e, 2)$ in B (since $T = \overline{B}$), and likewise for C. Conversely, each arc (e, i) of B or of C corresponds to an edge e of T. Hence, from $B \neq C$, we see that there must exist an edge e of T such that

- **either** we have $(e, 1) \in B$ and $(e, 2) \in C$,
- **or** we have $(e, 1) \in C$ and $(e, 2) \in B$.

Consider this edge e. We WLOG assume that $(e, 1) \in B$ and $(e, 2) \in C$ (else, we can just swap B with C). Let the arc $(e, 1)$ of G^{bidir} have source s and target t, so that $(e, 2)$ has source t and target s. The edge e thus has endpoints s and t.

Since B is an arborescence rooted to r, the vertex r is a to-root of B. Hence, there exists a path $\mathbf{p}$ from s to r in B. This path $\mathbf{p}$ must begin with the arc $(e, 1)$ [5]. Projecting this path $\mathbf{p}$ down onto T, we obtain a path $\overline{\mathbf{p}}$ from s to r in T. (By the word "projecting", we mean replacing each arc (e, i) by the corresponding edge e. Clearly, doing this to a path in B yields a path in T, because $T = \overline{B}$.) Since the path $\mathbf{p}$ begins with the arc $(e, 1)$, the "projected" path $\overline{\mathbf{p}}$ begins with the edge e. Thus, in the tree T, the path from s to r begins with the edge e (because this path must be the path $\overline{\mathbf{p}}$). As a consequence, t must be the second vertex of this path (since the edge e has endpoints s and t), so that removing the first edge from this path yields the path from t to r. Thus, $d(t, r) = d(s, r) - 1$, where d denotes distance on the tree T. Hence, $d(t, r) < d(s, r)$.

A similar argument (but with the roles of B and C swapped, as well as the roles of s and t swapped, and the roles of $(e, 1)$ and $(e, 2)$ swapped) shows that $d(s, r) < d(t, r)$. But this contradicts $d(t, r) < d(s, r)$.

This contradiction shows that our assumption was false. Thus, we have proved that $B = C$.

[4]Henceforth, "sparb" is short for "spanning arborescence of G^{bidir} rooted to r".

[5]*Proof.* Since r is a to-root of B, we know that there exists a path from t to r in B. Let $\mathbf{t}$ be this path. Extending this path $\mathbf{t}$ by the vertex s and the arc $(e, 1)$ (which we both insert at the start of $\mathbf{t}$), we obtain a walk $\mathbf{t}'$ from s to r in B. (So, if $\mathbf{t} = (t, \ldots, r)$, then $\mathbf{t}' = (s, (e, 1), t, \ldots, r)$.)

However, B is an arborescence rooted to r. Thus, Statement A'4 in the Dual Arborescence Equivalence Theorem (Theorem 1.1.3 in Lecture 17) shows that for each vertex $v \in V$, the digraph B has a unique walk from v to r. Hence, in particular, B has a unique walk from s to r. Thus, $\mathbf{p} = \mathbf{t}'$ (since both $\mathbf{p}$ and $\mathbf{t}'$ are walks from s to r in B). Since $\mathbf{t}'$ begins with the arc $(e, 1)$, we thus conclude that $\mathbf{p}$ begins with the arc $(e, 1)$.

Forget that we fixed B and C. We thus have shown that if B and C are two sparbs such that $\overline{B} = \overline{C}$, then $B = C$. In other words, our map

$$\begin{aligned}\left\{\text{spanning arborescences of } G^{\text{bidir}} \text{ rooted to } r\right\} &\to \{\text{spanning trees of } G\}, \\ B &\mapsto \overline{B}\end{aligned}$$

is injective.

We have now shown that this map is both surjective and injective. Hence, it is a bijection. This proves Proposition 1.1.1 **(b)**. □

1.2. The matrix-tree theorem

So counting spanning trees in a multigraph is a particular case of counting spanning arborescences (rooted to a given vertex) in a multidigraph. But how do we do either? Let us begin with some simple examples:

Example 1.2.1. There is only one spanning tree of the complete graph K_1:

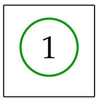

.

There is only one spanning tree of the complete graph K_2:

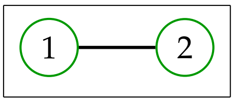

.

There are 3 spanning trees of the complete graph K_3:

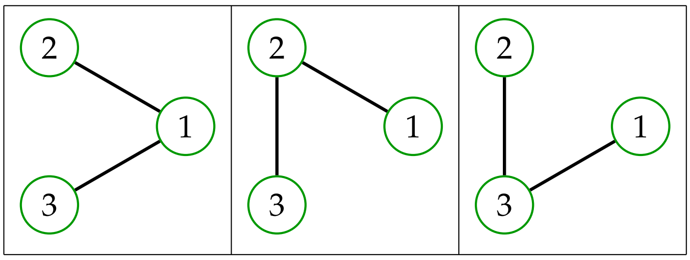

.

(They are all isomorphic, but still distinct.)

There are 16 spanning trees of the complete graph K_4:

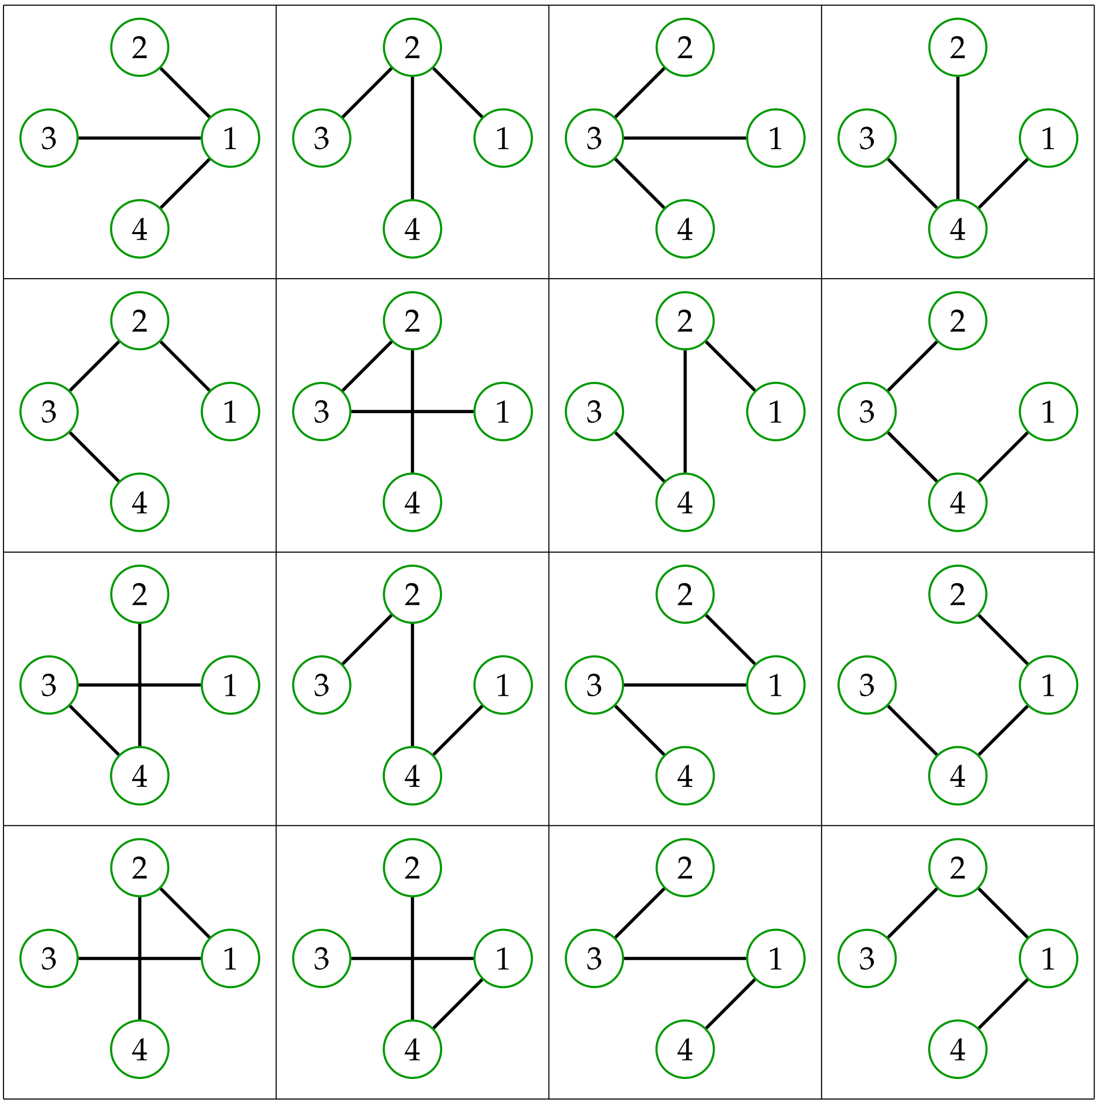

(There are only two non-isomorphic ones among them.)

This example suggests that the # of spanning trees of a complete graph K_n is n^{n-2}.

This is indeed true, and we will prove this later. For now, however, let us address the more general problem of counting spanning arborescences of an arbitrary digraph D.

First, we introduce a notation:

Definition 1.2.2. We will use the **Iverson bracket notation**: If $\mathcal{A}$ is any logical statement, then we set

$$[\mathcal{A}] := \begin{cases} 1, & \text{if } \mathcal{A} \text{ is true;} \\ 0, & \text{if } \mathcal{A} \text{ is false.} \end{cases}$$

For example, $[K_2 \text{ is a tree}] = 1$ whereas $[K_3 \text{ is a tree}] = 0$.

Definition 1.2.3. Let M be a matrix. Let i and j be two integers. Then,

$M_{i,j}$ will mean the entry of M in row i and column j;
$M_{\sim i, \sim j}$ will mean the matrix M with row i removed and column j removed.

For example,

$$\begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix}_{2,3} = f \qquad \text{and} \qquad \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix}_{\sim 2, \sim 3} = \begin{pmatrix} a & b \\ g & h \end{pmatrix}.$$

We shall now assign a matrix to (more or less) any multidigraph:[6]

Definition 1.2.4. Let $D = (V, A, \psi)$ be a multidigraph. Assume that $V = \{1, 2, \ldots, n\}$ for some $n \in \mathbb{N}$.

For any $i, j \in V$, we let $a_{i,j}$ be the # of arcs of D that have source i and target j.

The **Laplacian** of D is defined to be the $n \times n$-matrix $L \in \mathbb{Z}^{n \times n}$ whose entries are given by

$$L_{i,j} = \left(\deg^+ i\right) \cdot \underbrace{[i = j]}_{\substack{\text{This is also} \\ \text{known as } \delta_{i,j}}} - a_{i,j} \qquad \text{for all } i, j \in V.$$

In other words, it is the matrix

$$L = \begin{pmatrix} \deg^+ 1 - a_{1,1} & -a_{1,2} & \cdots & -a_{1,n} \\ -a_{2,1} & \deg^+ 2 - a_{2,2} & \cdots & -a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ -a_{n,1} & -a_{n,2} & \cdots & \deg^+ n - a_{n,n} \end{pmatrix}.$$

[6]Recall that the symbol "#" means "number".

Example 1.2.5. Let D be the digraph

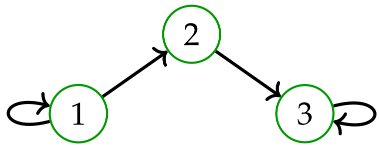

.

Then, its Laplacian is

$$\begin{pmatrix} 2-1 & -1 & -0 \\ -0 & 1-0 & -1 \\ -0 & -0 & 1-1 \end{pmatrix} = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{pmatrix}.$$

One thing we notice from this example is that loops do not matter at all to the Laplacian L. Indeed, a loop with source i and target i counts once in $\deg^+ i$ and once in $a_{i,i}$, but these contributions cancel out.

Here is a simple property of Laplacians:

Proposition 1.2.6. Let $D = (V, A, \psi)$ be a multidigraph. Assume that $V = \{1, 2, \ldots, n\}$ for some positive integer n.

Then, the Laplacian L of D is singular; i.e., we have $\det L = 0$.

Proof. The sum of all columns of L is the zero vector, because for each $i \in V$ we have

$$\sum_{j=1}^{n} \left(\left(\deg^+ i\right) \cdot [i = j] - a_{i,j} \right) = \underbrace{\sum_{j=1}^{n} \left(\deg^+ i\right) \cdot [i = j]}_{\substack{=\deg^+ i \\ \text{(since only the addend} \\ \text{for } j=i \text{ can be nonzero)}}} - \underbrace{\sum_{j=1}^{n} a_{i,j}}_{\substack{=\deg^+ i \\ \text{(since this is counting} \\ \text{all arcs with source } i)}}$$
$$= \deg^+ i - \deg^+ i = 0.$$

In other words, we have $Le = 0$ for the vector $e := (1, 1, \ldots, 1)^T$. Thus, this vector e lies in the kernel (aka nullspace) of L, and so L is singular.

(Note that we used the positivity of n here! If $n = 0$, then e is the zero vector, because a vector with 0 entries is automatically the zero vector.) $\square$

Proposition 1.2.6 shows that the determinant of the Laplacian of a digraph is not very interesting. But there is a general rule that when a matrix has determinant 0, its largest nonzero minors (= determinants of submatrices) often carry some interesting information; they are "the closest the matrix has" to a nonzero determinant. In the case of the Laplacian, they turn out to count spanning arborescences:

Theorem 1.2.7 (Matrix-Tree Theorem)**.** Let $D = (V, A, \psi)$ be a multidigraph. Assume that $V = \{1, 2, \ldots, n\}$ for some positive integer n.

Let L be the Laplacian of D. Let r be a vertex of D. Then,

$$(\text{\# of spanning arborescences of } D \text{ rooted to } r) = \det\left(L_{\sim r, \sim r}\right).$$

Before we prove this, some remarks:

- The determinant $\det\left(L_{\sim r, \sim r}\right)$ is the (r, r)-th entry of the adjugate matrix of L.
- The $V = \{1, 2, \ldots, n\}$ assumption is a typical "WLOG assumption": If you have an arbitrary digraph D, you can always rename its vertices as $1, 2, \ldots, n$, and then this assumption will be satisfied. Thus, Theorem 1.2.7 helps you count the spanning arborescences of any digraph. That said, you can also drop the $V = \{1, 2, \ldots, n\}$ assumption from Theorem 1.2.7 if you are okay with matrices whose rows and columns are indexed not by numbers but by elements of an arbitrary finite set[7].

Now, let us use the Matrix-Tree Theorem to count the spanning trees of K_n. This should help you get an intuition for the theorem before we come to its proof.

We fix a positive integer n. Let L be the Laplacian of the multidigraph K_n^{bidir} (where K_n, as we recall, is the complete graph on the set $\{1, 2, \ldots, n\}$). Then, each vertex of K_n^{bidir} has outdegree $n - 1$, and thus we have

$$L = \begin{pmatrix} n-1 & -1 & \cdots & -1 \\ -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & n-1 \end{pmatrix}$$

(this is the $n \times n$-matrix whose diagonal entries are $n - 1$ and whose off-diagonal entries are -1). By Proposition 1.1.1 **(b)** (applied to $G = K_n$ and $r = 1$), there is a bijection between $\left\{\text{spanning arborescences of } K_n^{\text{bidir}} \text{ rooted to } 1\right\}$ and

[7] Such matrices are perfectly fine, just somewhat unusual and hard to write down (which row do you put on top?). See `https://mathoverflow.net/questions/317105` for details.

$\{$spanning trees of $K_n\}$. Hence, by the bijection principle, we have

$$\begin{aligned}
&(\text{\# of spanning trees of } K_n) \\
&= \left(\text{\# of spanning arborescences of } K_n^{\text{bidir}} \text{ rooted to } 1\right) \\
&= \det\left(L_{\sim 1,\sim 1}\right) \qquad \left(\text{by Theorem 1.2.7, applied to } D = K_n^{\text{bidir}} \text{ and } r = 1\right) \\
&= \det \underbrace{\begin{pmatrix} n-1 & -1 & \cdots & -1 \\ -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & n-1 \end{pmatrix}}_{\text{an } (n-1)\times(n-1)\text{-matrix}}.
\end{aligned}$$

How do we compute this determinant? Here are three ways:

- The most elementary approach is using row transformations:

$$\begin{aligned}
&\det \begin{pmatrix} n-1 & -1 & \cdots & -1 \\ -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & n-1 \end{pmatrix} \\
&= \det \begin{pmatrix} n-1 & -1 & -1 & -1 & \cdots & -1 \\ -n & n & 0 & 0 & \cdots & 0 \\ -n & 0 & n & 0 & \cdots & 0 \\ -n & 0 & 0 & n & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ -n & 0 & 0 & 0 & \cdots & n \end{pmatrix} \qquad \left(\begin{array}{c}\text{here, we have} \\ \text{subtracted the 1st row} \\ \text{from each other row}\end{array}\right) \\
&= n^{n-2} \det \begin{pmatrix} n-1 & -1 & -1 & -1 & \cdots & -1 \\ -1 & 1 & 0 & 0 & \cdots & 0 \\ -1 & 0 & 1 & 0 & \cdots & 0 \\ -1 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ -1 & 0 & 0 & 0 & \cdots & 1 \end{pmatrix} \qquad \left(\begin{array}{c}\text{here, we have} \\ \text{factored out} \\ \text{an } n \text{ from each} \\ \text{row except for} \\ \text{the first row}\end{array}\right) \\
&= n^{n-2} \underbrace{\det \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ -1 & 1 & 0 & 0 & \cdots & 0 \\ -1 & 0 & 1 & 0 & \cdots & 0 \\ -1 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ -1 & 0 & 0 & 0 & \cdots & 1 \end{pmatrix}}_{\substack{=1 \\ \text{(since the matrix is triangular} \\ \text{with diagonal entries } 1,1,\ldots,1)}} \qquad \left(\begin{array}{c}\text{here, we have added the 2nd,} \\ \text{3rd, etc. rows to the 1st row}\end{array}\right) \\
&= n^{n-2}.
\end{aligned}$$

- The so-called **matrix determinant lemma** says that for any $m \times m$-matrix $A \in \mathbb{R}^{m \times m}$, any column vector $u \in \mathbb{R}^{m \times 1}$ and any row vector $v \in \mathbb{R}^{1 \times m}$, we have

  $$\det(A + uv) = \det A + v(\operatorname{adj} A)\, u.$$

  This helps us compute our determinant, since

  $$\begin{aligned} &\begin{pmatrix} n-1 & -1 & \cdots & -1 \\ -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & n-1 \end{pmatrix} \\ &= \underbrace{\begin{pmatrix} n & 0 & \cdots & 0 \\ 0 & n & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & n \end{pmatrix}}_{=A} + \underbrace{\begin{pmatrix} -1 \\ -1 \\ \vdots \\ -1 \end{pmatrix}}_{=u} \underbrace{\begin{pmatrix} 1 & 1 & \cdots & 1 \end{pmatrix}}_{=v}. \end{aligned}$$

- Here is an approach that is heavier on linear algebra (specifically, eigenvectors and eigenvalues[8]):

  Let $(e_1, e_2, \ldots, e_{n-1})$ be the standard basis of the $\mathbb{R}$-vector space $\mathbb{R}^{n-1}$ (so that e_i is the column vector with its i-th coordinate equal to 1 and all its other coordinates equal to 0). Then, we can find the following $n-1$ eigenvectors of our $(n-1) \times (n-1)$-matrix $\begin{pmatrix} n-1 & -1 & \cdots & -1 \\ -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & n-1 \end{pmatrix}$:

  - the $n-2$ eigenvectors $e_1 - e_i$ for all $i \in \{2, 3, \ldots, n-1\}$, each of them with eigenvalue n (check this!);
  - the eigenvector $e_1 + e_2 + \cdots + e_{n-1}$ with eigenvalue 1 (check this!).

  Since these $n-1$ eigenvectors are linearly independent (check this!), they form a basis of $\mathbb{R}^{n-1}$. Hence, our matrix is similar to the diagonal matrix with diagonal entries $\underbrace{n, n, \ldots, n}_{n-2 \text{ times}}, 1$ (by [Treil17, Chapter 4, Theorem 2.1]), and therefore has determinant $\underbrace{nn \cdots n}_{n-2 \text{ times}} 1 = n^{n-2}$.

There are other ways as well. Either way, the result we obtain is n^{n-2}. Thus, we have proved (relying on the Matrix-Tree Theorem, which we haven't yet proved):

[8] See [Treil17, Chapter 4] for a refresher.

Theorem 1.2.8 (Cayley's formula)**.** Let n be a positive integer. Then, the # of spanning trees of the complete graph K_n is n^{n-2}.

In other words:

Corollary 1.2.9. Let n be a positive integer. Then, the # of simple graphs with vertex set $\{1, 2, \ldots, n\}$ that are trees is n^{n-2}.

Proof. This is just Cayley's formula, since the simple graphs with vertex set $\{1, 2, \ldots, n\}$ that are trees are precisely the spanning trees of K_n. □

There are many ways to prove Cayley's formula. I can particularly recommend the two combinatorial proofs given in [Galvin21, §2.4 and §2.5], as well as Joyal's proof sketched in [Leinst19]. Most textbooks on enumerative combinatorics give one proof or another; e.g., [Stanle18, Appendix to Chapter 9] gives three. Cayley's formula also appears in Aigner's and Ziegler's best-of compilation of mathematical proofs [AigZie14, Chapter 33] with four different proofs. Note that some of the sources use a matrix-tree theorem for **undirected** graphs; this is a particular case of our matrix-tree theorem.

However, in order to complete our proof, we still need to prove the Matrix-Tree Theorem.

We shall do this next time (i.e., in Lecture 19). First, let us prepare with a simple lemma (yet another criterion for a digraph to be an arborescence):

Lemma 1.2.10. Let $D = (V, A, \psi)$ be a multidigraph. Let r be a vertex of D. Assume that D has no cycles. Assume moreover that D has no arcs with source r. Assume furthermore that each vertex $v \in V \setminus \{r\}$ has outdegree 1. Then, the digraph D is an arborescence rooted to r.

This lemma is precisely homework set #6 problem 4 **(b)**, at least after reversing all arcs. But let us give a self-contained proof here:

Proof of Lemma 1.2.10. Let u be any vertex of D. Let $\mathbf{p} = (v_0, a_1, v_1, a_2, v_2, \ldots, a_k, v_k)$ be a longest path of D that starts at u. [9] Thus, $v_0 = u$.

We shall show that $v_k = r$. Indeed, assume the contrary. Thus, $v_k \neq r$, so that $v_k \in V \setminus \{r\}$. Hence, the vertex v_k has outdegree 1 (since we assumed that each vertex $v \in V \setminus \{r\}$ has outdegree 1). Thus, there exists an arc b of D that has source v_k. Consider this arc b, and let w be its target. Thus, appending the arc b and the vertex w to the end of the path $\mathbf{p}$, we obtain a walk

$$\mathbf{w} = (v_0, a_1, v_1, a_2, v_2, \ldots, a_k, v_k, b, w)$$

of D that starts at u (since $v_0 = u$). Proposition 1.2.9 in Lecture 10 shows that this walk $\mathbf{w}$ either is a path or contains a cycle. Hence, $\mathbf{w}$ is a path (since D

[9] Such a path clearly exists, since the length-0 path (u) is a path of D that starts at u, and since a path of D cannot have length larger than $|V| - 1$.

has no cycles). Thus, $\mathbf{w}$ is a path of D that starts at u. Since $\mathbf{w}$ is longer than $\mathbf{p}$ (namely, longer by 1), this shows that $\mathbf{p}$ is not the longest path of D that starts at u. But this contradicts the very definition of $\mathbf{p}$.

This contradiction shows that our assumption was false. Hence, $v_k = r$. Thus, $\mathbf{p}$ is a path from u to r (since $v_0 = u$ and $v_k = r$). Therefore, the digraph D has a path from u to r (namely, $\mathbf{p}$).

Forget that we fixed u. We thus have shown that for each vertex u of D, the digraph D has a path from u to r. In other words, r is a to-root of D. Furthermore, we have $\deg^+ r = 0$ (since D has no arcs with source r), and each $v \in V \setminus \{r\}$ satisfies $\deg^+ v = 1$ (since we have assumed that each vertex $v \in V \setminus \{r\}$ has outdegree 1). In other words, the digraph D satisfies Statement A'6 from the dual arborescence equivalence theorem (Theorem 1.1.3 in Lecture 17). Therefore, it satisfies Statement A'1 from that theorem as well (since all six statements A'1, A'2, ..., A'6 are equivalent). In other words, D is an arborescence rooted to r. This proves Lemma 1.2.10. $\square$

References

[AigZie14] Martin Aigner, Günter M. Ziegler, *Proofs from the Book*, 6th edition, Springer 2018.

[Galvin21] David Galvin, *Basic Discrete Mathematics (Spring 2021)*.
`https://www3.nd.edu/~dgalvin1/60610/60610_S21/index.html`
Follow the overleaf link. Notes: `Course-notes.tex`; solved homework: `main.tex`.

[Leinst19] Tom Leinster, *The probability that an operator is nilpotent*, arXiv:1912.12562v2.

[Stanle18] Richard P. Stanley, *Algebraic Combinatorics: Walks, Trees, Tableaux, and More*, 2nd edition, Springer 2018.
See `https://math.mit.edu/~rstan/algcomb/errata2.pdf` for errata.

[Treil17] Sergei Treil, *Linear Algebra Done Wrong*, 4 September 2017.
`https://www.math.brown.edu/streil/papers/LADW/LADW.html`