Math 372 lecture for Wednesday, Week 11

Matrix-tree theorem

Let $G = (V, E)$ be a directed graph. For this means that V is a finite set of vertices and E is a finite multiset of directed edges. A directed edge is an ordered pair (u, v) with $u, v \in V$. It is not necessary to assume G is connected in any sense. A *subgraph* of G is a directed graph whose vertices and directed edges form sub(multi)sets of V and E, respectively.

The Laplacian for a directed graph. For $v \in V$, let $\mathrm{outdeg}_G(v)$ be the number of edges in G of the form (v, w) for some $w \in V$. Fix an ordering of the vertices $v_1, \cdots, v_n$, and define

$$D = \mathrm{diag}(\mathrm{outdeg}_G(v_i))_{i=1,\dots,n}.$$

Define the directed $n \times n$ adjacency matrix A by letting A_{ij} be the number of directed edges starting at v_i, i.e., of the form (v_i, w) for some $w \in V$. Then the *Laplacian matrix* for G is

$$L := D - A^t,$$

where A^t is the transpose of A. (In the literature, what we are calling L would often be called the transpose of the Laplacian.) The columns of L define firing rules in a chip-firing game on G.

Definition. A *(directed) spanning tree of G rooted at $s \in V$* is a subgraph T such that for all $v \in V$, there exists a unique directed path in T from v to s. The vertex s is the *root* or *sink* of the tree.

If T is a directed spanning tree of G rooted at s, then one may show that (i) T contains *all* of the vertices of G (hence, the word "spanning"); (ii) T contains no directed cycles; and (iii) for all vertices v of G, the outdegree of v in T is 0 if $v = s$, and is 1, otherwise. In particular, T contains no multiple edges.

Example. The graph pictured below has three directed edges and one undirected edge:

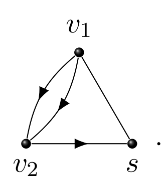

The determinant of its reduced Laplacian with respect to s is

$$\det\begin{pmatrix} 3 & 0 \\ -2 & 1 \end{pmatrix} = 3,$$

which is the number of spanning trees rooted at s, as shown below:

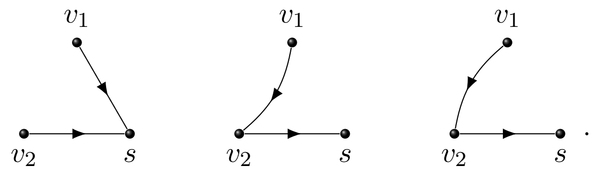

Note that second two trees are different since the multiple edges of the form (v_1, v_2) are counted as distinct.

Let L be the Laplacian matrix of G relative to an ordering of the vertices, $v_1, \ldots, v_n$. For each $k \in \{1, \ldots, n\}$, let $L^{(k)}$ denote the $(n-1) \times (n-1)$ matrix formed by removing the k-th row and column of L. This is the *reduced Laplacian* for G with respect to v_k.

Matrix-tree Theorem. The determinant of $L^{(k)}$ is the number of spanning trees of G rooted at v_k.

Proof. Since a permutation of the vertices induces a corresponding permutation of the rows and columns of L, it suffices to consider the case $k = n$. For ease of notation, we write $\widetilde{L} := L^{(n)}$. Letting a_{ij} denote the number of times (v_i, v_j) appears as an edge of G, we have

$$\widetilde{L} = \begin{pmatrix} \sum_{i \neq 1} a_{1i} & -a_{21} & -a_{31} & \ldots & -a_{n-1,1} \\ -a_{12} & \sum_{i \neq 2} a_{2i} & -a_{32} & \ldots & -a_{n-1,2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -a_{1,n-1} & -a_{2,n-1} & -a_{3,n-1} & \ldots & \sum_{i \neq n-1} a_{n-1,i} \end{pmatrix}$$

where $\sum_{i \neq k} a_{ki}$ denotes the sum over $i \in \{1, \ldots, n\} \setminus \{k\}$. Each column encodes the rule for reverse-firing the corresponding vertex. $\square$

Let $\mathfrak{S}_{n-1}$ be the permutation group on $\{1, \ldots, n-1\}$. Recall that the *sign* of $\sigma \in \mathfrak{S}_{n-1}$ is $\mathrm{sgn}(\sigma) := (-1)^t$ where t is the number of factors in any expression for σ as a product of transpositions—it records whether an even or odd number of swaps is required to create the permutation σ. Let $\mathrm{Fix}(\sigma)$ be the set of *fixed points* of σ:

$$\mathrm{Fix}(\sigma) := \{i \in \{1, \ldots, n-1\} : \sigma(i) = i\}.$$

Then

$$\det \widetilde{L} = \sum_{\sigma \in \mathfrak{S}_{n-1}} \mathrm{sgn}(\sigma) \widetilde{L}_{\sigma(1),1} \cdots \widetilde{L}_{\sigma(n-1),n-1}, \tag{1}$$

where

$$\widetilde{L}_{\sigma(k),k} = \begin{cases} \sum_{i \neq k} a_{k,i} & \text{if } k \in \mathrm{Fix}(\sigma) \\ -a_{k,\sigma(k)} & \text{otherwise.} \end{cases}$$

The idea now is to expand (1) into signed monomials in the a_{ij} and to think of each monomial as a directed graph by identifying a_{ij} with the directed edge (v_i, v_j) labeled with the number of times this edge appears in G, i.e., with a_{ij}, itself:

$$v_i \xrightarrow{a_{ij}} v_j \quad .$$

We then show that after cancellation due to the signs of the permutations, the remaining monomials correspond exactly to the distinct spanning trees rooted at v_n. Each monomial itself—a product of various a_{ij}—is an integer which counts the number of times its corresponding spanning tree occurs as a spanning tree of G. (Recall that since G may have repeated edges, a spanning tree may occur more than once.)

We pause now for an extended example. For those readers wishing to skip ahead, the end of the example is marked with a line.

Example. Take $n = 10$ and $\sigma = (2,7)(3,5,9) \in \mathfrak{S}_9$. The set of fixed points is $\mathrm{Fix}(\sigma) = \{1,4,6,8\}$ and $\mathrm{sgn}(\sigma) = \mathrm{sgn}((2,7))\ \mathrm{sgn}((3,5,9)) = (-1)\cdot 1 = -1$. The term in the expansion of $\det \widetilde{L}$ corresponding to σ is

$$\begin{aligned}
&\mathrm{sgn}(\sigma)\widetilde{L}_{\sigma(1),1}\widetilde{L}_{\sigma(2),2}\cdots\widetilde{L}_{\sigma(9),9}\\
&\quad = (-1)(a_{1,2}+\cdots+a_{1,10})(-a_{2,7})(-a_{3,5})(a_{4,1}+\cdots+a_{4,10})'\\
&\qquad \cdot(-a_{5,10})(a_{6,1}+\cdots+a_{6,10})'(-a_{7,2})(a_{8,1}+\cdots+a_{8,10})'(-a_{9,3}),
\end{aligned}$$

where the prime symbol on a factor indicates the term of the form $a_{i,i}$ should be omitted from the enclosed sum. Continuing,

$$\begin{aligned}
= (-1)\Big[&\overbrace{(a_{1,2}+\dots)}^{\sigma(1)=1}\overbrace{(a_{4,1}+\dots)'}^{\sigma(4)=4}\overbrace{(a_{6,1}+\dots)'}^{\sigma(6)=6}\overbrace{(a_{8,1}+\dots)'}^{\sigma(8)=8}\Big]\\
&\cdot\Big[\underbrace{(-a_{2,7})(-a_{7,2})}_{(2,7)}\underbrace{(-a_{3,5})(-a_{5,9})(-a_{9,3})}_{(3,5,9)}\Big].
\end{aligned}$$

Question: Which monomials, identified with directed graphs, appear in the expansion of the above?

Answer: For each fixed point i of σ, we get a choice of any edge of the form $v_i \xrightarrow{a_{ij}} v_j$ where $j \in \{1,\dots,10\}$ and $j \neq i$. For each non-trivial cycle of σ, there is only one choice:

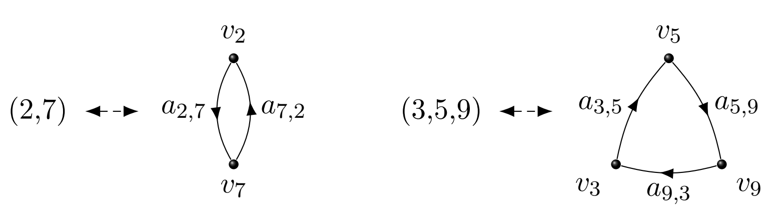

Figure 1 considers three monomials coming from the expansion of the term in $\det \widetilde{L}$ corresponding to σ. Each monomial m corresponds to a directed graph G_m. Column F shows the part of G_m that comes from choices for the fixed points of σ, and column C shows the part that comes from the nontrivial cycles. Note that, as in example (c), these two parts may share vertices. There may be an edge connecting a fixed point vertex to a cycle vertex in G_m. Example (b) shows that it is not necessary for v_{10} to occur in G_m. In general, v_{10} does not appear if and only if each vertex in G_m has a path to a directed cycle (since the outdegree for each non-root vertex is 1).

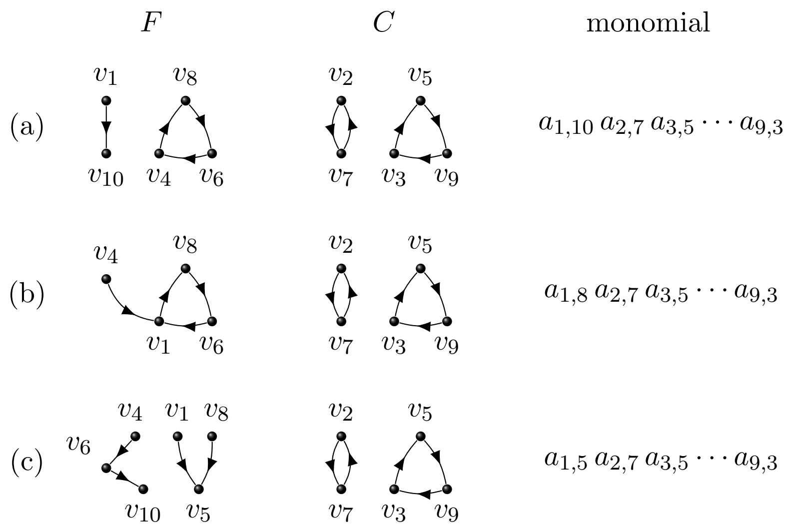

Figure 1: Monomials and corresponding graphs for Example .

Finally, it is important to determine the sign of each monomial corresponding to σ in the expansion of $\det \widetilde{L}$. The sign is determined by $\operatorname{sgn}\sigma$ and by the number of factors of the form $-a_{ij}$ that go into the calculation of the monomial. With these two considerations in mind, it is straightforward to see that the resulting sign is $(-1)^{\#\text{ non-trivial cycles of }\sigma}$. For instance, consider the monomial in example (a) in Figure 1. It appears in the expansion of $\det \widetilde{L}$ in the term

$$\operatorname{sgn}((2,7)(3,5,9))\ \underbrace{a_{1,10}a_{4,8}a_{6,4}a_{8,6}}_{\operatorname{Fix}(\sigma)=\{1,4,6,8\}}\ \underbrace{(-a_{2,7})(-a_{7,2})}_{(2,7)}\underbrace{(-a_{3,5})(-a_{5,9})(-a_{9,3})}_{(3,5,9)}.$$

Each cycle of σ ultimately contributes a factor of -1:

$$\operatorname{sgn}((2,7))(-a_{2,7})(-a_{7,2}) = -1\cdot a_{2,7}a_{7,2}$$
$$\operatorname{sgn}((3,5,9))(-a_{3,5})(-a_{5,9})(-a_{9,3}) = -1\cdot a_{3,5}a_{5,9}a_{9,3}.$$

We now return to the proof. The monomials in the expansion of (1) correspond exactly with signed, weighted, ordered pairs (F, C) of graphs F and C formed as follows:

1. Choose a subset $X \subseteq \{1, \ldots, n-1\}$ (representing the fixed points of some σ).
2. Make any loopless, directed (not necessarily connected) graph F with vertices $\{1, \ldots, n\}$ such that
$$\operatorname{outdeg}_F(i) = \begin{cases} 1 & \text{if } i \in X \\ 0 & \text{if } i \notin X. \end{cases}$$
3. Let C be any vertex-disjoint union of directed cycles of length at least 2 (i.e., no loops) such that C contains all of the vertices $\{1, \ldots, n-1\} \setminus X$.

Each of these ordered pairs of graphs (F, C) is associated with an element of $\mathfrak{S}_{n-1}$, with the vertices of outdegree one in F determining the fixed points and with C determining the cycles. In general, this is a many-to-one relationship, given the choices in step (2). The *weight* of (F, C), denoted $\operatorname{wt}(F, C)$ is the product of its labels—those a_{ij} such that (v_i, v_j) occurs in either F or C—multiplied by $(-1)^\gamma$ where γ is the number of cycles in C. For instance, for each of the three examples in Figure 1, the number of cycles in C is 2, so the weight is just the listed monomial, without a sign change. With this notion of weight, it then follows in general that
$$\det \widetilde{L} = \sum_{(F,C)} \operatorname{wt}(F, C).$$

Let Ω denote the set of ordered pairs (F, C), constructed as above, but such that either F or C contains a directed cycle. We show that the monomials corresponding to elements of Ω cancel in pairs in the expansion of (1) by constructing a *sign reversing transposition* on Ω. Given $(F, C) \in \Omega$, pick the cycle γ of the disjoint union $F \sqcup C$ with the vertex of smallest index. Then if γ is in F, move it to C, and vice versa. Formally, if the cycle is in F, define $F' = F \setminus \{\gamma\}$ and $C' = C \cup \{\gamma\}$, and if it is in C, define $F' = F \cup \{\gamma\}$ and $C' = C \setminus \{\gamma\}$. This defines a transposition $(F, C) \mapsto (F', C')$ on Ω such that $\operatorname{wt}(F, C) = -\operatorname{wt}(F', C')$ since the number of cycles in C differs from the number of cycles in C' by one. See Figure 2 for an example. It follows that in the sum $\sum \operatorname{wt}(F, C)$, terms paired by the transposition cancel,

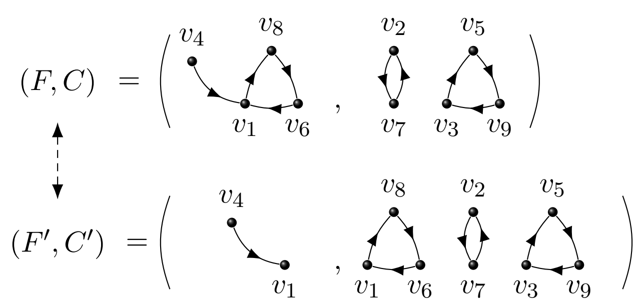

Figure 2: Sign reversing transposition.

leaving only those terms $\operatorname{wt}(F, C)$ for which the transposition is undefined, i.e., those (F, C)

such that the graph $F \sqcup C$ contains no cycles. In this case, the corresponding permutation is the identity permutation, $C = \emptyset$, and F is a spanning tree rooted at v_n. The weight, $\mathrm{wt}(F, C)$, counts the number of times this spanning tree occurs as a spanning tree of G due to G having multiple edges.

Consequences of the matrix-tree theorem

We now obtain several corollaries of the matrix-tree theorem.

Corollary. Let G be an undirected multigraph. Then the order of $\mathrm{Jac}(G)$ is the number of directed spanning trees of G rooted at the sink.

Proof. We have seen that $\mathrm{Jac}(G) \simeq \mathbb{Z}^{n-1}/\operatorname{im}\widetilde{L}$ where $\widetilde{L}$ is the reduced Laplacian of G with respect to any vertex. Let D be the Smith normal form for $\widetilde{L}$, and write $D = U\widetilde{L}V$ for some integer invertible matrices U and V. Then we have seen that U defines an isomorphism $\mathbb{Z}^{n-1}/\operatorname{im}\widetilde{L} \simeq \mathbb{Z}^{n-1}/\operatorname{im}(D)$, and $\mathbb{Z}^{n-1}/\operatorname{im}(D) \simeq \prod_{i=1}^{n-1} \mathbb{Z}/d_i\mathbb{Z}$ where the d_i are the diagonal entries of D. Therefore, the number of elements in the group $\mathbb{Z}^{n-1}/\operatorname{im}(D)$ is $\det(D)$. Since U and V are invertible over the integers, we have $\det(D) = \det(U\widetilde{L}V) = \pm\widetilde{L}$. By the matrix-tree theorem $\widetilde{L}$ is nonnegative since it counts the number of spanning trees of G, and we know $\det(D)$ is nonnegative. Putting this all together:

$$|\mathrm{Jac}(G)| = |\mathbb{Z}^{n-1}/\operatorname{im}(D)| = \det(D) = \det(\widetilde{L}) = \# \text{ spanning trees of } G.$$

$\square$

A *tree on n labeled vertices* is a connected undirected graph with n labeled vertices and no cycles.

Corollary. (Cayley's formula) The number of trees on n labeled vertices is n^{n-2}.

Proof. Let $\widetilde{L}$ be the reduced Laplacian matrix of the complete graph K_n. Then by the matrix-tree theorem, Cayley's number is $\det(\widetilde{L})$. In homework, we showed this determinant is n^{n-2}. $\square$

Let $L^{(ij)}$ denote the matrix obtained from the Laplacian L of G by removing the i-th row and j-th column.

Corollary. The ij-th cofactor, $(-1)^{i+j}\det L^{(ij)}$, of L is the number of directed spanning trees rooted at the j-th vertex.

Proof. We have $(-1)^{i+j}\det L^{(ij)} = \det L^{(jj)}$ since the sum of the rows of L is the zero vector (exercise for the reader). The result then follows from the matrix-tree theorem.

Corollary. Let G be a directed graph with n vertices. Suppose the Laplacian matrix of G has eigenvalues $\mu_1, \ldots, \mu_n$ with $\mu_n = 0$. For each vertex v, let κ_v be the number of directed spanning trees of G rooted at v. Then,

$$\mu_1 \cdots \mu_{n-1} = \sum_v \kappa_v.$$

In other words, the product of these eigenvalues is the total number of rooted trees.

Proof. We may assume the vertex set is $1, \ldots, n$, with the i-th column of L corresponding to vertex i. First note that since L is singular, a zero eigenvalue μ_n always exists. Factoring the characteristic polynomial of L, we have

$$\det(L - I_n x) = (\mu_1 - x) \cdots (\mu_n - x).$$

We calculate the coefficient of x in this expression in two ways. Since $\mu_n = 0$, by expanding the right-hand side, we see the coefficient is $-\mu_1 \ldots \mu_{n-1}$. Now consider the left-hand side. For each i, let r_i denote the i-th row of L, and let e_i denote the i-th standard basis vector. Then

$$\det(L - I_n x) = \det(r_1 - e_1 x, \ldots, r_n - e_n x).$$

By multilinearity of the determinant, letting $\widetilde{L}_i$ denote the reduced Laplacian with respect to vertex i, the coefficient of x is

$$\sum_{i=1}^{n} \det(r_1, \ldots, r_{i-1}, -e_i, r_{i+1}, \ldots, r_n) = -\sum_{i=1}^{n} \det(\widetilde{L}_i) = -\sum_{i=1}^{n} \kappa_i.$$

$\square$

Remark. If G is undirected, or more generally, if G is Eulerian (which means that for each vertex v, the number of edges directed into v is equal to the number of edges directed out of v) then the number of spanning trees rooted at a vertex is independent of the particular vertex. Calling this number κ, the previous Corollary says in this case that

$$\kappa = \frac{\mu_1 \ldots \mu_{n-1}}{n}.$$