ARBORESCENCES OF DERIVED GRAPHS

CHRISTOPHER DOWD, VALERIE ZHANG, AND SYLVESTER ZHANG

1. Introduction

1.1. **Arborescences.** Let $\Gamma = (V, E, \mathrm{wt})$ be an edge-weighted quiver—that is, a directed multigraph with a function on the edges $\mathrm{wt} : E \to R$, where R is some ring. We usually abbreviate "edge-weighted" to "weighted," and instead of saying "quiver" or "directed multigraph" we usually say "graph"—any instance in which we wish to consider only directed or only simple graphs will be explicitly noted. We consider the weights of the edges of G to be indeterminates, treating the weight $\mathrm{wt}(e)$ of an edge e as a variable; let the set of such variables be denoted $\mathrm{wt}(E)$. We denote the set of outgoing edges of a vertex v by $E_i(v)$, and the set of ingoing edges of v by $E_t(v)$ (here i and t stand for initial and terminal, respectively). We denote the initial vertex of an edge e by e_i and terminal vertex of e by e_t. Via abuse of notation, we will sometimes make the statement $e = (v, w)$, which means $e_i = v$ and $e_t = w$. However, when Γ is not necessarily simple, there may be more than one edge satisfying these properties, so this statement is not actually a statement of equality.

Definition 1.1. An *arborescence* T of Γ rooted at $v \in V$ is a spanning tree directed towards v. That is, for all vertices w, there exists a directed path from w to v through T. [1] We denote the set of arborescences of Γ rooted at vertex v by $\mathcal{T}_v(\Gamma)$. The *weight of an arborescence* $\mathrm{wt}(T)$ is the product of the weights of its edges:

$$\mathrm{wt}(T) = \prod_{e \in T} \mathrm{wt}(e)$$

We denote by $A_v(\Gamma)$ the sum of the weights of all arborescences of Γ rooted at v:

$$A_v(\Gamma) = \sum_{T \in \mathcal{T}_v(\Gamma)} \mathrm{wt}(T)$$

$A_v(\Gamma)$ is either zero or a homogeneous polynomial of degree $|V| - 1$ in the edge weights of G.

1.2. **The Laplacian Matrix and the Matrix Tree Theorem.** The Matrix Tree Theorem, also known as Kirchoff's Theorem, yields a way of computing $A_v(\Gamma)$ through the *Laplacian matrix* of Γ.

Definition 1.2. The *Laplacian matrix* $L(\Gamma)$ of a graph Γ is the difference of the weighted degree matrix D and the weighted adjacency matrix A of Γ:

$$L(\Gamma) = D(\Gamma) - A(\Gamma)$$

Here, the weighted degree matrix is the diagonal matrix whose i-th entry is

$$d_{ii} = \sum_{e \in E_i(v)} \mathrm{wt}(e)$$

, and the weighted adjacency matrix has entries defined by

$$a_{ij} = \sum_{e = (v_i, v_j)} \mathrm{wt}(e)$$

Received by the editors July 26, 2019.

[1] Most other authors define an arborescence rooted at v to be a spanning tree directed *away* from v, so that v is the unique source rather than the unique sink. Our convention is consistent with the study of R-systems.

Again, since we will always be working with weighted graphs in this paper, we will usually drop the word "weighted" when talking about the Laplacian matrix. Note the ordering convention of the definition of the Laplacian: when we label the vertices of Γ as $v_1, v_2, \ldots$, we will always assume that v_1 corresponds to the first row and column of $L(\Gamma)$, that v_2 corresponds to the second row and column of $L(\Gamma)$, and so on.

Theorem 1.3. (Matrix Tree Theorem) *The sum of the weights of arborescences rooted at v_i is equal to the the minor of $L(\Gamma)$ obtained by removing the i-th row and column:*

$$A_{v_i}(\Gamma) = \det L_i^i(\Gamma)$$

1.3. Covering graphs.

Definition 1.4. A *k-fold cover* of $\Gamma = (V, E)$ is a graph $\tilde{\gamma} = (\tilde{V}, \tilde{E})$ that is a k-fold covering space of G in the topological sense that preserves edge weight. More concretely, there exists a projection map $\pi : \tilde{\Gamma} \to \Gamma$ such that

(1) π maps vertices to vertices and edges to edges;

(2) $|\pi^{-1}(v)| = |\pi^{-1}(e)| = k$ for all $v \in V, e \in E$;

(3) For all $\tilde{e} \in \tilde{E}$, we have $\mathrm{wt}(\tilde{e}) = \mathrm{wt}(\pi(\tilde{e}))$;

(4) Local homeomorphism: restriction of π to a neighborhood of $\tilde{v} \in \tilde{V}$ is a graph isomorphism onto the neighborhood of $\pi(\tilde{v}) \in V$. Here we mean "neighborhood" in the graph-theoretical sense, i.e. the neighborhood of v is the set $E_i(v)$ union the set of vertices adjacent to v.

1.4. **An invariant ratio.** Motivating our research is a corollary of Galashin and Pylavskyy, which they discovered in their study of R-systems.

Proposition 1.5. (Corollary of Galashin–Pylyavskyy, Theorem 2.3 in [GP17]) *Let Γ be a simple and strongly connected graph, and let $\tilde{\Gamma}$ be any covering graph of Γ. Let v be a vertex of Γ. Then the ratio*

$$\frac{A_{\tilde{v}}(\tilde{\Gamma})}{A_v(\Gamma)}$$

is a well-defined as a rational function in the edge weights of Γ, and furthermore this ratio is independent of the choice of vertex v and its lift $\tilde{v}$.

Proof. Since Γ is strongly connected, there is at least one arborescence rooted at v, so $A_v(\Gamma)$ is not the zero polynomial, showing well-definition.

To show that this ratio is independent of our choice of vertex, we define two sets of parameters $X = (X_v)_{v \in V}$ and $X' = (X'_v)_{v \in V}$ that satisfy

$$\sum_{(u,v) \in E} \mathrm{wt}(u, v) \frac{X_v}{X_{u'}} = \sum_{(v,w) \in E} \mathrm{wt}(v, w) \frac{X_w}{X'_v}. \tag{1}$$

We wish to find solutions to this relation when X is the all-ones vector. We rewrite (1) as

$$\sum_{(u,v) \in E} \mathrm{wt}(u, v) \frac{X_v}{X_u} \frac{X_u}{X_{u'}} = \sum_{(v,w) \in E} \mathrm{wt}(v, w) \frac{X_w}{X_v} \frac{X_v}{X_{v'}}.$$

To solve, we treat the above equation as a linear system in the variables $\left(\frac{X_v}{X'_v}\right)_{v\in V}$. The matrix M of this system is

$$m_{vu} = \begin{cases} \sum_{(v,w)\in E} \mathrm{wt}(v,w)\frac{X_w}{X_v} & u = v \\ -\mathrm{wt}(u,v)\frac{X_v}{X_u} & (u,v) \in E \\ 0 & \text{else} \end{cases}$$

$$= \begin{cases} \sum_{(v,w)\in E} \mathrm{wt}(v,w) & u = v \\ -\mathrm{wt}(u,v) & (u,v) \in E \\ 0 & \text{else} \end{cases}$$

This is the transpose of the weighted Laplacian matrix $L(\Gamma)$ of Γ, albeit with modified edge weights: if $e = (u,v)$, then the new weight of e is $\mathrm{wt}(u,v)\frac{X_v}{X_u}$. The Matrix Tree Theorem implies that L has rank $|V|-1$, since the minor obtained by removing any row and column corresponding to a vertex v is nonzero, since $A_v(\Gamma) \neq 0$ for Γ strongly connected. Thus, M has a one-dimensional kernel as well, so there is a unique solution up to scaling for $M\left(\frac{X'}{X}\right) = 0$. We may check that $\frac{X_v}{X'_v} = A_v(\Gamma)$ is such a solution—the inner product of row v of M with the vector $(A_v(\Gamma))_{v\in\Gamma}$ is

$$A_v(\Gamma)\sum_{w\neq v}\mathrm{wt}(v,w) - \sum_{u\neq v} A_u(\Gamma)\,\mathrm{wt}(u,v)$$

However, given a pair (T,e) consisting of an arborescence T rooted at u and an edge $e = (u,v)$, we may obtain a pair $(\varphi(T),\varphi(e))$ consisting of an arborescence $\varphi(T)$ rooted at v and an edge $\varphi(e) = (u,v)$. $\varphi(T)$ is obtained by appending e to T, and then removing the unique outgoing edge $\varphi(e)$ of v in T. φ is invertible, and we have $\mathrm{wt}(T)\,\mathrm{wt}(e) = \mathrm{wt}(T')\,\mathrm{wt}(e')$, so in expression (2) above the positive and negative terms cancel, showing that the expression is indeed zero, as desired. Thus, $\frac{X_v}{X'_v} = A_v(\Gamma)$ is the unique solution up to scaling to the system (1).

However, equation (1) is defined locally—the solutions to X_v and X'_v is depend only on neighborhood of v. Because any cover of a graph is locally homeomorphic to the original graph, any solution to (1) will hold both for vertices in the original graph and for its lifts in the cover. Therefore, choose one representative $\tilde{v} \in \tilde{V}$ for each vertex $v \in V$. Then we have two solutions to (1)—one from repeating the previous solution, and the other from using the local homeomorphism property:

$$\frac{X_{\tilde{v}}}{X'_{\tilde{v}}} = A_{\tilde{v}}(\tilde{\Gamma});$$
$$\frac{X_{\tilde{v}}}{X'_{\tilde{v}}} = A_v(\Gamma)$$

But we already know that any solution to (1) is unique up to scaling. Therefore, $(A_{\tilde{v}}(\tilde{\Gamma}))_{v\in V} = \lambda(A_v(\Gamma))_{v\in V}$ for some scalar λ, so that

$$\frac{A_{\tilde{v}}(\tilde{\Gamma})}{A_v(\Gamma)} = \lambda \tag{2}$$

for any $v \in V$, as desired.

Similarly, the local homeomorphism property also makes it clear that it does not matter which lift $\tilde{v}$ of v we use. ∎

Galashin and Pylyavskyy are concerned with the unique solution to the system (1), which they call the *R-system* of the graph Γ. They worked exclusively over strongly connected simple graphs, but in Section 4 we will extend Proposition 1.5 to arbitrary multigraphs.

Given the existence of this invariant ratio, we ask if there exists a nice explicit formula that computes it. This is our primary research question for this paper. While Galashin-Pylyavskyy holds for arbitrary covering graphs including non-regular covers, we will restrict ourselves to exploring special types of covers known as *derived graphs*.

1.5. **Voltage graphs and derived graphs.** In order to define a derived graph, we must first define the notion of *voltage*:

Definition 1.6. A *weighted G-voltage graph* $\Gamma = (V, E, \mathrm{wt}, \nu)$ is a weighted quiver with each edge e also labeled by an element $\nu(e)$ of a finite group G. This labeling is called a *voltage* of G with respect to G. Note that the voltage of an edge e is distinct from the weight of e.

Definition 1.7. Given a G-voltage graph Γ, we may construct an $|H|$-fold covering graph of G known as the *derived graph* $\tilde{\Gamma} = (\tilde{V}, \tilde{E}, \mathrm{wt})$. The derived graph of a voltage graph is a graph with vertex set $\tilde{V} = V \times H$. For each edge $e = (v, w)$ in Γ with voltage $g \in G$, we add to $\tilde{G}$ the $|H|$ edges

$$\{[v \times x, w \times (hx)] : x \in H\}$$

to create the edge set $\tilde{E}$.

Example 1.8. Let $H = \mathbb{Z}/3\mathbb{Z} = \{1, h, h^2\}$, and let G be the following H-volted graph, where edges labeled (x, y) have edge weight x and voltage y:

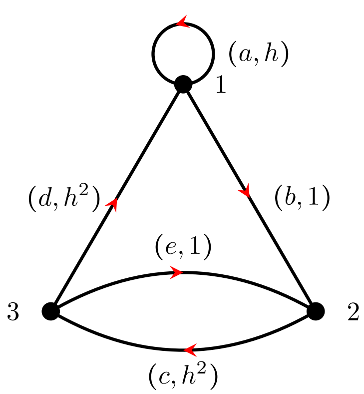

Then the derived graph $\tilde{G}$, with vertices $(v, x) = v_x$ and with edges labeled by weight, is

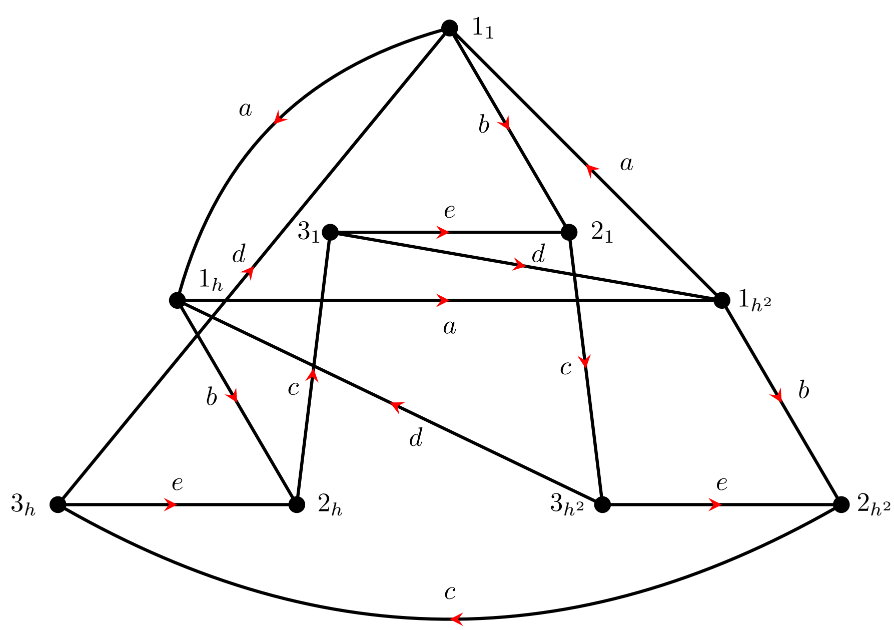

One might think that derived graphs are an overly special subset of general covering graphs, but they turn out to be surprisingly general.

Definition 1.9. Given a graph Γ and a covering graph $\tilde{\Gamma}$, the *deck group* $\mathrm{Aut}(\pi)$ of $\tilde{\Gamma}$ is the group of automorphisms on $\tilde{\Gamma}$ that preserve the fibers of the projection map π.

Definition 1.10. A *regular cover* $\tilde{\Gamma}$, sometimes known as a *Galois cover*, of a graph Γ is a covering graph whose deck group is transitive on each fiber $\pi^{-1}(v)$ for each $v \in V$.

Theorem 1.11. (Theorems 3 and 4 in [GT75]) *Every regular cover $\tilde{\Gamma}$ of a graph Γ may be realized as a derived cover of Γ with voltage assignments in $\mathrm{Aut}(\pi)$. Conversely, every derived graph is a regular cover.*

The main focus of this paper is to explore the relationship between the arborescences of a voltage graph Γ and the arborescences of its derived graph $\tilde{\Gamma}$, and thus explore this relationship for arbitrary regular covers. Our goal is to give an explicit formula for the ratio described by Proposition 1.5. Specifically, we examine how to express this ratio in terms of the determinant of a variant of the Laplacian matrix of Γ.

2. The voltage Laplacian

We wish to define a matrix similar to the Laplacian matrix that tracks all the relevant information in an G-voltage graph. In order to do so in general, we need to extend our field of coefficients in order to codify the data given by the voltage function ν. Following the language of Reiner and Tseng in [RT13]:

Definition 2.1. The *reduced group algebra* of a finite group G over a field F is the quotient

$$\overline{F[G]} = \frac{F[G]}{\left\langle \sum_{g \in G} h \right\rangle},$$

where $F[G]$ is the group algebra of G over F. That is, we quotient the group algebra by the all-ones vector with respect to the basis given by G.

For simplicity, in the remainder of this paper we take $F = \mathbb{Q}$. In practice, we will always be dealing with integers. Note that if $G \cong \mathbb{Z}/2\mathbb{Z}$, then $\overline{\mathbb{Q}[G]} \cong \mathbb{Q}$, with the non-identity element of G identified with -1. For this reason, we will sometimes refer to $\mathbb{Z}/2\mathbb{Z}$-voltage graphs as *signed* graphs, and the voltage of an edge of such a graph as the sign of that edge. Similarly, if $G \cong \mathbb{Z}/p\mathbb{Z}$, with p prime, then $\overline{\mathbb{Q}[G]} \cong \mathbb{Q}(\zeta_p)$, where ζ_p is a primitive p-th root of unity and the generator h of G is identified with ζ_p. The fact that the reduced group alebgra of prime cyclic G is actually a field extension over $\mathbb{Q}$ will be important later in giving us nice formulas for the ratio of arborescences described in the introduction.

We now define a generalization of the Laplacian matrix that takes into account voltages:

Definition 2.2. The *voltage adjacency matrix* $\mathscr{A}(G)$ has entries given by

$$a_{ij} = \sum_{e=(v_i,v_j)\in E} \nu(e)\mathrm{wt}(e),$$

where we consider $\nu(e)$ as an element of the reduced group algebra $\overline{\mathbb{Q}[G]}$. That is, the i,j-th entry consists of sum of the *volted* weights of all edges going from the i-th vertex to the j-th vertex. The *voltage Laplacian matrix* $\mathscr{L}(\Gamma)$ is defined as

$$\mathscr{L}(\Gamma) = D(\Gamma) - \mathscr{A}(\Gamma)$$

where $D(\Gamma)$ is the (unvolted) weighted degree matrix as described in Definition 1.2.

Note that in the special case $G \cong \{1\}$ we have $\mathscr{L}(\Gamma) = L(\Gamma)$.

Since we consider the edge weights of Γ as indeterminates, we treat the entries of $\mathscr{L}(G)$ as elements of $\overline{\mathbb{Z}[G]}[\mathrm{wt}(E)] \subset \overline{\mathbb{Q}[G]}[\mathrm{wt}(E)]$—that is, the polynomial ring of edge weights with coefficients in the group algebra.

Example 2.3. Let Γ the $\mathbb{Z}/3\mathbb{Z}$-voltage graph in Example 1.8. Under the identification $\overline{\mathbb{Q}[\mathbb{Z}/3\mathbb{Z}]} \cong \mathbb{Q}(\zeta_3)$, the voltage Laplacian of Γ is

$$\begin{aligned}\mathscr{L}(\Gamma) &= \begin{bmatrix} a+b & 0 & 0 \\ 0 & c & 0 \\ 0 & 0 & d+e \end{bmatrix} - \begin{bmatrix} \zeta_3 a & b & 0 \\ 0 & 0 & \zeta_3^2 c \\ \zeta_3^2 d & e & 0 \end{bmatrix} \\ &= \begin{bmatrix} (1-\zeta_3)a+b & -b & 0 \\ 0 & c & -\zeta_3^2 c \\ -\zeta_3^2 d & -e & d+e \end{bmatrix}\end{aligned}$$

3. The determinant of the voltage Laplacian counts vector fields

Definition 3.1. A *vector field* γ of a directed graph Γ is a subgraph of Γ such that every vertex of Γ has outdegree 1 in γ. Similarly to arborescences, we define the *weight* of a vector field $\mathrm{wt}(\gamma) := \prod_{e\in\gamma} \mathrm{wt}(e)$ of a vector field be the product of its edge weights, so that $\mathrm{wt}(\gamma)$ is a degree $|V|$ monomial with respect to the edge weights of Γ.

The determinant of $\mathscr{L}(\Gamma)$ counts vector fields of Γ in the following way:

Theorem 3.2. *Let G be an abelian group, and let Γ be an edge-weighted Γ-voltage quiver. Write $C(\gamma)$ for the set of cycles in a vector field γ (such cycles are necessarily disjoint), and let the* voltage *of a cycle c be given by $\nu(c) = \prod_{e\in c} \nu(e) \in \overline{\mathbb{Q}[G]}[\mathrm{wt}(E)]$—this product is well-defined since G is abelian. Then*

$$\sum_{\gamma\subseteq\Gamma} \left[\mathrm{wt}(\gamma) \prod_{c\in C(\gamma)} (1-\nu(c)) \right] = \det \mathscr{L}(\Gamma)$$

where the sum ranges over all vector fields γ of Γ.

The following special case of the above, with $G = \mathbb{Z}/2\mathbb{Z}$, will be useful later in our first proof of the ratio of arborescences formula for signed graphs:

Corollary 3.3. *Suppose that Γ is a $\mathbb{Z}/2\mathbb{Z}$-voltage graph, i.e. a signed graph. Define a* negative vector field *γ of Γ to be a vector field such that every cycle c of γ has an odd number of negative edges, so that $\nu(c) = -1$. Denote the set of negative vector fields of G by $\mathcal{N}(\Gamma)$. Then*

$$\sum_{\gamma\in\mathcal{N}(\Gamma)} 2^{\#C(\gamma)} \mathrm{wt}(\gamma) = \det \mathscr{L}(\Gamma)$$

We present two proofs of Theorem 3.2: one original, as far as we are aware, and the other dating back to Chaiken.

The first proof proceeds by deletion-contraction, and requires the following lemma in order to run smoothly:

Lemma 3.4. *Let Γ be as in Theorem 3.2 with voltage function $\nu : E \to \overline{\mathbb{Q}[G]}$, let v be any vertex of Γ, and let $g \in G$. We define a new voltage function $\nu_{v,g}$ given by*

$$\nu_{v,g}(e) = \begin{cases} \nu(e): & \text{if } e = (v,v) \\ g\nu(e): & \text{if } e \in E_i(v), e \notin E_t(v) \\ g^{-1}\nu(e): & \text{if; } e \in E_t(v), e \notin E_i(v) \\ \nu(e): & \text{else} \end{cases}$$

Then

(a) for any cycle c of Γ, we have $\nu(c) = \nu_{v,g}(c)$; and

(b) the determinant of the voltage Laplacian of Γ with respect to the voltage ν is equal to the determinant of the voltage Laplacian of Γ with respect to $\nu_{v,g}$. That is,

$$\det \mathscr{L}(V, E, \mathrm{wt}, \nu) = \det \mathscr{L}(V, E, \mathrm{wt}, \nu_{v,g})$$

Proof. (a) If c does not contain the vertex v, or if c is a loop at v, then the voltages of all edges in c remain unchanged. Otherwise, c contains exactly one ingoing edge e of v and one outgoing edge f of v, so that

$$\begin{aligned}\nu_{v,h}(c) &= \frac{\nu(c)}{\nu(e)\nu(f)}[g\nu(e)][g^{-1}\nu(f)] \\ &= \nu(c)\end{aligned}$$

as desired.

(b) The matrix $\mathscr{L}(V, E, \mathrm{wt}, \nu)$ may be transformed into the matrix $\mathscr{L}(V, E, \mathrm{wt}, \nu_{v,g})$ by multiplying the row corresponding to v by h and multiplying the column corresponding to v by g^{-1}, so the determinant remains unchanged.

■

This lemma will allow us some freedom to change the voltage of Γ as needed.

Proof. (First proof of Theorem 3.2). Denote the left-hand side of the theorem as

$$\Omega(\Gamma) := \sum_{\gamma \subseteq \Gamma} \left[\mathrm{wt}(\gamma) \prod_{c \in C(\gamma)} (1 - \nu(c)) \right]$$

We proceed by deletion-contraction. Our base case will be when the only edges of Γ are loops. When this happens, $\mathscr{L}(\Gamma)$ is diagonal, with

$$\ell_{ii} = \sum_{e=(v_i,v_i) \in E} (1 - \nu(e))\,\mathrm{wt}(e).$$

Thus we have

$$\det \mathscr{L}(\Gamma) = \prod_{i=1}^{|V|} \left(\sum_{e=(v_i,v_i) \in E} [1 - \nu(e)]\,\mathrm{wt}(e) \right)$$

If we expand the product above, each term will correspond to a unique combination of one loop per vertex of Γ. But such combinations are precisely the vector fields of Γ, so we obtain

$$\det \mathscr{L}(\Gamma) = \Omega(\Gamma)$$

For the inductive step, assume that there exists at least one edge e between distinct vertices, and assume that the proposition holds for graphs with fewer non-loop edges than Γ. Using the lemma, we may change the voltage of Γ so that e has voltage 1 without changing either $\Omega(\Gamma)$ or $\det \mathscr{L}(\Gamma)$. Without loss of generality, let $v_1 = e_i$ and $v_2 = e_t$.

If γ is a vector field of Γ, then γ either contains e or it does not. In the latter case, γ is also a vector field of $\Gamma \backslash e$. Clearly all such γ arise uniquely from a vector field of $\Gamma \backslash e$. Therefore, there is a weight-preserving bijection between the vector fields of Γ not containing e and the vector fields of $\Gamma \backslash e$.

Otherwise, if $e \in \gamma$, then no other edge of the form (v_1, v_j) is in γ. We define a special type of contraction: let $\Gamma/_1 e := (\Gamma/e) \backslash E_i(v_1)$. That is, we contract along e, and delete all other edges originally in $E_i(v_1)$. Note that the contraction process merges vertices v_1 and v_2 into a "supervertex," which we denote v_{12}.

Then vector field γ descends uniquely to a vector field $\overline{\gamma}$ on $\Gamma/_1 e$. Every vector field $\overline{\gamma}$ in $\Gamma/_1 e$ corresponds uniquely to a vector field of Γ containing e, obtained by letting the unique edge coming out the supervertex v_{12} in $\overline{\gamma}$ be the unique edge coming out of the vertex v_2 in γ, and letting e be the unique edge with source at v_1 in γ. This inverse map shows that the vector fields of Γ containing e are in bijection with the vector fields of $\Gamma/_1 e$. This bijection is weight-preserving up to a factor of $\mathrm{wt}(e)$. Finally, note that γ and its contraction $\overline{\gamma}$ have the same number of cycles, with the same voltages. If a cycle contains e in γ, then that cycle is made one edge shorter in $\overline{\gamma}$, but still has positive length since e is assumed to not be a loop. If c is a cycle containing e in γ, then since e has voltage 1, the cycle voltage $\nu(c/e)$ of the contracted version of c is equal to the cycle voltage before contraction. Thus, we may write

$$\Omega(\Gamma) = \Omega(\Gamma\backslash e) + \mathrm{wt}(e)\Omega(\Gamma/_1 e)$$

By inductive hypothesis, since $\Gamma\backslash e$ and $\Gamma/_1 e$ have strictly fewer non-loop edges than Γ, we have

$$\Omega(\Gamma\backslash e) + \mathrm{wt}(e)\Omega(\Gamma/_1 e) = \det \mathscr{L}(\Gamma\backslash e) + \mathrm{wt}(e)\det \mathscr{L}(\Gamma/_1 e)$$

Note that $\mathscr{L}(\Gamma\backslash e)$ is equal to $\mathscr{L}(\Gamma)$ with $\mathrm{wt}(e)$ deleted from both the $1,1$- and $1,2$-entries. Therefore, via expansion by minors, we obtain

$$\det L(\Gamma\backslash e) + \mathrm{wt}(e)\det L_1^1(\Gamma) + \mathrm{wt}(e)\det L_1^2(\Gamma) = \det \mathscr{L}(\Gamma) \tag{3}$$

where $\mathscr{L}_i^j(\Gamma)$ is the submatrix of $\mathscr{L}(\Gamma)$ obtained by removing the i-th row and the j-th column.

To construct $\mathscr{L}(\Gamma/_1 e)$ from $\mathscr{L}(\Gamma)$, we disregard the first row of $\mathscr{L}(\Gamma)$, since the special contraction $\Gamma/_1 e$ simply removes the outgoing edges $E_i(v_1)$. Then, we combine the first two columns of $\mathscr{L}(\Gamma)$ by making their sum the first column of $\mathscr{L}(\Gamma/_1 e)$, since when we perform a contraction that merges v_1 and v_2 into v_{12}, we also have $E_t(v_1)\cup E_t(v_2) = E_t(v_{12})$. Thus $\mathscr{L}(\Gamma/_1 e)$ is a $(|V|-1)\times(|V|-1)$ matrix that agrees with both $\mathscr{L}_1^1(\Gamma)$ and $\mathscr{L}_1^2(\Gamma)$ on its last $|V|-2$ columns, and whose first column is the sum of the first columns of $\mathscr{L}_1^1(\Gamma)$ and $\mathscr{L}_1^2(\Gamma)$. Therefore,

$$\det \mathscr{L}(\Gamma/_1\, e) = \det \mathscr{L}_1^1(\Gamma) + \det \mathscr{L}_1^2(\Gamma)$$

Substituting into (3), we obtain

$$\begin{aligned}\det \mathscr{L}(\Gamma) &= \det \mathscr{L}(\Gamma\backslash e) + \mathrm{wt}(e)\det \mathscr{L}(\Gamma/_1\, e)\\ &= \Omega(\Gamma\backslash e) + \mathrm{wt}(e)\Omega(\Gamma/_1\, e)\\ &= \Omega(\Gamma)\end{aligned}$$

as desired. ■

The second proof of the theorem follows a style similar to Chaiken's proof of the Matrix Tree Theorem in [Cha82]. Chaiken actually proves a more general identity, which he calls the "All-Minors Matrix Tree Theorem," that gives a combinatorial formula for any minor of the voltage Laplacian. We do not reproduce such generality here, but instead follow a simplified version of his proof, more along the lines of Stanton and White's version of Chaiken's proof of the Matrix Tree Theorem [SW86].

Proof. (Second proof of Theorem 3.2) (Chaiken). Let Γ have n vertices. For simplicity, assume that Γ has no multiple edges, since we can always decomposed $\det \mathscr{L}(\Gamma)$ into a sum of determinants of voltage Laplacians of simple subgraphs of Γ, which also partitions the sum given in the theorem. We also assume that Γ is a complete bidirected graph, since we can ignore edges not in Γ by just considering them to have edge weight 0. Write $\mathscr{L}(\Gamma) = (\ell_{ij})$, write $D(\Gamma) = (d_{ij})$, and write $\mathscr{A}(\Gamma) = (a_{ij})$, so that $\ell_{ij} = \delta_{ij}d_{ii} - a_{ij}$. Then the determinant of $\mathscr{L}(\Gamma)$ may be decomposed as

$$\det \mathscr{L}(\Gamma) = \det(\delta_{ij}d_{ii} - a_{ij}) = \sum_{S\subseteq[n]}\left[\sum_{\pi\in P(S)}(-1)^{\#C(\pi)}\,\mathrm{wt}_\nu(\pi)\prod_{i\in[n]-S} d_{ii}\right]$$

where $P(S)$ denotes the set of permutations of S, the set $C(\pi)$ is set of cycles of π, and $\mathrm{wt}_\nu(\pi) \coloneqq \prod_{i\in S} a_{i,\pi(i)}$. The product of the d_{ii} may be rewritten as a sum over functions $[n]-S\to[n]$, yielding

$$\begin{aligned}\det \mathscr{L}(\Gamma) &= \sum_{S\subseteq[n]}\sum_{\pi\in P(S)}(-1)^{c(\pi)}\,\mathrm{wt}_\nu(\pi)\sum_{f:[n]-S\to[n]}\mathrm{wt}(f)\\ &= \sum_{S\subseteq[n]}\sum_{\pi\in P(S)}\sum_{f:[n]-S\to[n]}(-1)^{c(\pi)}\,\mathrm{wt}_\nu(\pi)\,\mathrm{wt}(f) \qquad (4)\end{aligned}$$

where $\mathrm{wt}(f)$ denotes the *unvolted* weight of the edge set corresponding to the function f, since this part of the product ultimately comes from the degree matrix. Thus, the determinant may be expressed as a sum of triples (S,π,f) of the above form—that is, we let S be an arbitrary subset of $[n]$, we let π be a permutation on S, and we let f be a function $[n]-S\mapsto[n]$.

The permutation π may always be decomposed into cycles, and f will sometimes have cycles as well—that is, sometimes we have $f^{(m)}(k) = k$ for some $k \in \mathbb{Z}$ and $k \in [n] - S$. We can "swap" cycles between π and f. Suppose c is a cycle of f that we want to swap into π. Let the subset of $[n]$ on which c is defined be denoted W. Then we may obtain from our old triple a new triple $(S \coprod W, \pi \coprod c, f|_{[n]-S-W})$, where $\pi \coprod c$ denotes the permutation on $S \coprod W$ given by $(\pi \coprod c)(v) = \pi(v)$ if $v \in S$ and $(\pi \coprod c)(v) = c(v)$ if $v \in W$. That is, we "move" C from f to π. Similarly, if c is a cycle of π, then we can obtain a new triple $(S - W, \pi|_{S-W}, f \coprod c)$. Note that these two operations are inverses.

This process is always weight-preserving—it does not matter whether c is considered as part of π or as part of f, since it will always contribute $\mathrm{wt}(c)$ to the product. However, one iteration of this map will swap the sign of $(-1)^{\#C(\pi)}$, and will also remove or add a factor from $\mathrm{wt}_\nu(\pi)$ corresponding to the voltage of c. If π and f have k cycles amongst both of them, then there are 2^k possibilities for swaps, yielding a free action of $(\mathbb{Z}/2\mathbb{Z})^k$. If we start from the case π is the empty partition, then the sign $(-1)^{\#C(\pi)}$ starts at 1. Every time we choose to swap a cycle c into π from f, we flip this sign and multiply by $\nu(c)$, effectively multiplying by $-\nu(c)$. Thus, the sum of terms in (4) coming from the orbit of the action of $(\mathbb{Z}/2\mathbb{Z})^k$ on (S, f, π) is

$$\mathrm{wt}(\pi)\,\mathrm{wt}(f) \prod_{c \in C(\pi) \cup C(f)} (1 - \nu(c))$$

where $\mathrm{wt}(\pi)$ is now unvolted. This orbit class corresponds to the contribution of one vector field γ of Γ to the overall sum, where γ is the unique vector field such that $\mathrm{wt}(\gamma) = \mathrm{wt}(\pi)\,\mathrm{wt}(f)$. Thus, summing over all orbit classes, we obtain the desired formula:

$$\det \mathscr{L}(\Gamma) = \sum_{\gamma \subseteq \Gamma} \left[\mathrm{wt}(\gamma) \prod_{c \in C(\gamma)} (1 - \nu(c)) \right]$$

■

Corollary 3.3 gives a fast proof of the Matrix Tree Theorem (Theorem 1.3). Indeed, Chaiken considers Theorem 3.2 to be a generalization of the Matrix Tree Theorem.

Proof. (Matrix Tree Theorem) The set $\mathcal{T}_{v_j}(\Gamma)$ of arborescences of the graph Γ rooted at v remains the same if we remove all edges in $E_i(v_j)$ and replace them with a single loop e, so let this be the case. We assign a $\mathbb{Z}/2\mathbb{Z}$-voltage to Γ: let all edges of Γ be positive except e, which is negative. Then the negative vector fields of Γ are precisely the arborescences of Γ plus the edge e—no other configurations are possible, since any cycle other than the loop e would be positive. Since every such negative vector field has exactly one cycle (the loop e), by the corollary the sum of the weights of the arborescences of Γ is given by

$$A_{v_j}(\Gamma) = \frac{\det \mathscr{L}(\Gamma)}{2\,\mathrm{wt}(e)}.$$

However, the row in $\mathscr{L}(\Gamma)$ corresponding to v_j consists of all zeroes except in the column corresponding to v, which contains $2\,\mathrm{wt}(e)$. Thus, $\det \mathscr{L}(G)$ is given by $2\,\mathrm{wt}(e) \det \mathscr{L}_j^j(\Gamma)$, where $\det \mathscr{L}_j^j(\Gamma)$ is the minor of $\mathscr{L}(\Gamma)$ corresponding to removing the j-th row and column. Thus,

$$\det \mathscr{L}_v^v(\Gamma) = A_v(\Gamma)$$

as desired. ■

4. The ratio formula for 2-fold covers

We now state and prove a result about the arborescences of any 2-fold covering graph. Recall in that in Proposition 1.5, we showed that the ratio $\frac{A_{\tilde{v}}(\tilde{\Gamma})}{A_v(\Gamma)}$ is well-defined and independence of the choice of vertex v when G is simple and strongly connected. The study of R-systems occurs almost exclusively in the context of strongly connected simple digraphs, but with the proposition in hand we no longer need to consider the relevant R-system. Thus, we may extend the proposition to any directed multigraph:

Corollary 4.1. (Invariance under rerooting). *Proposition 1.5 extends to arbitrary multigraphs whenever possible. That is, even if G is not simple or strongly connected, we still have that the ratio $\frac{A_{\tilde{v}}(\tilde{\Gamma})}{A_v(\Gamma)}$ is independent of the choice of vertex v and its lift $\tilde{v}$ as long as this ratio is defined (i.e. $A_v(\Gamma) \neq 0$).*

Proof. Suppose Γ is simple but not necessarily strongly connected. We may consider Γ as a subgraph of the complete graph $K_{|V|}$ (any strongly connected graph $K \supseteq \Gamma$ with the same vertex set as G would do as well). Denote the complement of the edge set of G in $K_{|V|}$ as $E_{K\backslash\Gamma}$. By Proposition 1.5, we know that $\frac{A_{\tilde{v}}(\widetilde{K_{|V|}})}{A_v(K_{|V|})}$ is well-defined and independent of the choice of v.

Now note that whenever an edge of Γ has weight 0, any arborescence containing that edge vanishes in the polynomial $A_v(\Gamma)$. Thus, let φ be the evaluation homomorphism that maps the weight of every edge in $E_{K\backslash\Gamma}$ to 0, so that

$$\varphi(A_v(K_{|V|})) = A_v(\Gamma) \tag{5}$$

$$\varphi(A_{\tilde{v}}(\widetilde{K_{|V|}})) = A_{\tilde{v}}(\tilde{\Gamma}) \tag{6}$$

since every arborescence of Γ rooted at v is also an arborescence of $K_{|V|}$ rooted at v, and this set of arborescences is precisely the set of arborescences not containing any edge in $E_{K\backslash V}$ (and similarly for $\widetilde{K_{|V|}}$ and $\tilde{\Gamma}$). Since ratio of the left-hand sides of equations (5) and (6) is invariant under changing root, so is the ratio of the right-hand sides.

In the additional case that Γ is not simple, we can augment Γ to a graph Γ^+ by placing a vertex on the midpoint of each edge of Γ—given $e = (u, v)$, we add a vertex v_e with unique ingoing edge (u, v_e) and unique outgoing edge (v_e, v). Set $\nu(u, v_e) = \nu(u, v)$ and $\nu(v_e, v) = 1$. Then Γ^+ is a simple graph, since every edge is either of the form (v_e, v) or (u, v_e), and we know $|E_i(v_e)| = |E_t(v_e)| = 1$. We therefore know that Proposition 1.5 holds for Γ^+. However, note that whenever we root at some vertex $v \in V(\Gamma) \subseteq V(\Gamma^+)$, every arborescence must contain every edge of the form (v_e, v), since this is the only outgoing edge of v_e; similarly, every arborescence of the cover of Γ^+ contains both lifts of every (v_e, v). Since these edges must always be used in both of these constructions, we may freely contract along them without cutting out or merging any arborescences in both Γ^+ and $\tilde{\Gamma^+}$. Contracting along every such edge transforms Γ^+ back into Γ. Therefore, if we let $\varphi : \mathrm{wt}(v_e, v) \mapsto 1$ and $\varphi : \mathrm{wt}(u, v_e) = \mathrm{wt}(u, v)$, we have $A_v(\Gamma) = \varphi(A_v(\Gamma))$ and $A_{\tilde{v}}(\tilde{\Gamma}) = \varphi(A_{\tilde{v}}(\tilde{\Gamma})$. Therefore, Proposition 1.5 result holds for non-simple Γ as well.

∎

We now turn to one the major theorems of this report, which provides a formula for $\frac{A_v(\tilde{\Gamma})}{A_v(\Gamma)}$ when $\tilde{\Gamma}$ is a 2-fold cover:

Theorem 4.2. *Let Γ be an edge-weighted $\mathbb{Z}/2\mathbb{Z}$-volted directed multigraph—that is, a signed graph. For any vertex v of Γ and any lift $\tilde{v}$ of v to the derived graph $\tilde{\Gamma}$ of Γ, we have*

$$A_v(\Gamma) \det \mathscr{L}(\Gamma) = 2A_{\tilde{v}}(\tilde{\Gamma})$$

Equivalently, either $A_v(\Gamma) = 0$ or we have

$$\frac{A_{\tilde{v}}(\tilde{\Gamma})}{A_v(\Gamma)} = \frac{1}{2} \det \mathscr{L}(\Gamma).$$

Proof. We proceed by strong induction, and by relying heavily on the fact that we can usually prove the formula rooting at a specific vertex to proliferate the formula to any vertex. We will also apply the results about negative vector fields from Corollary 3.3.

4.1. **Base cases.** First, note that if some vertex of Γ has outdegree 0, then both sides of the above identity are 0, no matter which root is chosen, since a row of $\det \mathscr{L}(\Gamma)$ will be the zero vector and at least two vertices in $\tilde{\Gamma}$ will have outdegree 0.

Next, suppose that every vertex of Γ has outdegree exactly 1. Choose any $v \in G$. Then there is only one candidate for a negative vector field of Γ, and only one candidate for an arborescence of Γ rooted at v. If Γ has more than one cycle, then Γ is disconnected, so that $A_v(\Gamma) = 0$ and $A_{\tilde{v}}(\Gamma) = 0$. Assume Γ has exactly one cycle. If this cycle does not contain v, then no path from the vertices in the cycle to v exists, so no arborescence rooted at v exists, and $A_v(\Gamma) = A_{\tilde{v}}(\tilde{\Gamma}) = 0$. Now assume that the unique cycle of Γ contains v. If this cycle is positive, then that no negative vector fields exist, and thus $\det \mathscr{L}(\Gamma) = 0$ by Corollary 3.3. Furthermore, there exist two disjoint lifts of this cycle to $\tilde{\Gamma}$, which again means that $\tilde{\Gamma}$ is disconnected, so that $A_{\tilde{v}(\Gamma)} = 0$. Thus, the statement holds in these cases.

The only remaining case is if Γ has a unique negative cycle that contains v. Then this cycle lifts to a cycle twice as long in $\tilde{\Gamma}$ containing both lifts v' and v'' of v, which is the unique cycle in $\tilde{G}$. Therefore, the edges of $\tilde{G}$, except the unique edge in $E_i(\tilde{v})$, form a unique arborescence $\tilde{T}$ rooted at $\tilde{v}$. Siimilarly, the edges of G, except the unique edge in $E_i(v)$, form a unique arborescence T rooted at v; and Γ forms the unique negative vector field γ of itself. Thus $\mathrm{wt}(T)\,\mathrm{wt}(\gamma) = \mathrm{wt}(\tilde{\Gamma})$. Since $\det \mathscr{L}(\Gamma) = 2\,\mathrm{wt}(\gamma)$ by Corollary 3.3, previous proposition, we conclude that $A_v(\Gamma) \det \mathscr{L}(\Gamma) = 2A_{\tilde{v}}(\tilde{\Gamma})$. This proves the identity when $|E| \le |V|$.

4.2. **Main inductive step.** Now suppose that the identity holds whenever $|E| \le k$ for some $k \ge |V|$, and let G have $k+1$ edges. By the pigeonhole principle, at least one vertex v of G satisfies $|E_i(v)| \ge 2$. Assume further that we can choose such v with $A_v(\Gamma) \ne 0$.

Let e be any edge in $E_i(v)$, and define $E_i^e(v) := E_i(v) \backslash e$. Then both $G\backslash e$ and $G \backslash E_i(v)$ have at most k edges, since $|E_i(v)| \ge 2$. By inductive hypothesis,

$$A_v(\Gamma\backslash e) \det \mathscr{L}(\Gamma\backslash e) = 2A_{\tilde{v}}(\widetilde{\Gamma\backslash e})$$
$$A_v(\Gamma\backslash E_i^e(v)) \det \mathscr{L}(\Gamma\backslash E_i^e(v)) = 2A_{\tilde{v}}(\Gamma\widetilde{\backslash E_i^e(v)}))$$

Without loss of generality, let $\tilde{v} = v'$, and let e', e'' be the lifts of e with sources at v', v'', respectively. Every arborescence of $\tilde{\Gamma}$ rooted at v' contains exactly one edge in $E_i(v'')$. This edge is either e'' or it is not, so we may partition such arborescences into two disjoint classes based on whether they include e''—that is,

$$A_{v'}(\tilde{\Gamma}) = A_{v'}(\tilde{\Gamma}\backslash e'') + A_{v'}(\tilde{\Gamma}\backslash E_i^{e''}(v''))$$

However, note that

$$A_{v'}(\tilde{\Gamma}\backslash e'') = A_{v'}(\tilde{\Gamma}\backslash \{e', e''\}) = A_{v'}(\widetilde{\Gamma\backslash e})$$

No arborescence rooted at v' utilizes any edge with source at v', so we may simply delete the edge e' from Γ as it suits us. Similarly,

$$\begin{aligned} A_{v'}(\tilde{\Gamma}\backslash E_i^{e''}(v'') &= A_{v'}(\tilde{\Gamma}\backslash (E_i^{e''}(v'') \cup E_i^{e'}(v'))) \\ &= A_{v'}(\Gamma\widetilde{\backslash E_i^e(v)}) \end{aligned}$$

Thus,

$$\begin{aligned} A_{\tilde{v}}(\tilde{\Gamma}) &= A_{\tilde{v}}(\widetilde{\Gamma\backslash e}) + A_{\tilde{v}}(\Gamma\widetilde{\backslash E_i^e(v)}) \\ &= \frac{1}{2} A_v(\Gamma\backslash e) \det \mathscr{L}(\Gamma\backslash e) + \frac{1}{2} A_v(\Gamma\backslash E_i^e(v)) \det \mathscr{L}(\Gamma\backslash E_i^e(v)) \end{aligned}$$

Now, note that $A_v(\Gamma\backslash e) = A_v(\Gamma\backslash E_i^e(v)) = A_v(\Gamma)$—again, no arborescence rooted at v utilizes any edge in $E_i(v)$. Thus,

$$2A_{\tilde{v}}(\tilde{\Gamma}) = A_v(\Gamma)\left(\det \mathscr{L}(\Gamma\backslash e) + \det \mathscr{L}(\Gamma\backslash E_i^e(v))\right)$$

Finally, note that the matrix $\mathscr{L}(\Gamma)$, the matrix $\mathscr{L}(G\backslash e)$, and the matrix $L(G\backslash E_i^e(v))$ are all equal except in the row corresponding to v, and that the sum of the v-th rows of $\mathscr{L}(\Gamma\backslash e)$ and $L(\Gamma\backslash E_i^e(v))$ is equal to the v-th row of $\mathscr{L}(\Gamma)$. Thus, $\det\mathscr{L}(\Gamma) = \det\mathscr{L}(\Gamma\backslash e) + \det L(\Gamma\backslash E_i^e(v))$, so that

$$2A_{\tilde{v}}(\tilde{\Gamma}) = A_v(\Gamma)\det\mathscr{L}(\Gamma)$$

as desired. By Corollary 4.1, we conclude that

$$2A_{\tilde{u}}(\tilde{\Gamma}) = A_u(\Gamma)\det\mathscr{L}(\Gamma)$$

for any choice of $u \in V$.

4.3. **Exceptional cases.** We must choose v to be some vertex with $|E_i(v)| \geq 2$, but what if all such vertices satisfy $A_v(\Gamma) = 0$? Then either $A_u(\Gamma) = 0$ for all $u \in \Gamma$, in which case the theorem is trivially satisfied, or there exists some vertex u with outdegree exactly 1 such that $A_u(\Gamma) \neq 0$.

Suppose that in the latter case we can choose u such that there exist two distinct arborescences T_1 and T_2 rooted at u. Then there must exist some vertex w such that the outgoing edge e of w in T_1 is distinct from the outgoing edge f of w in T_2. Define Γ^+ to be the graph obtained by Γ by adding an auxiliary edge a from u to w, so that Γ^+ has $k+2$ edges, and therefore $\Gamma^+\backslash e$ and $\Gamma^+\backslash E_i^e(w)$ both have at most $k+1$ edges. Since u has outdegree 2 in both $\Gamma^+\backslash e$ and $\Gamma^+\backslash E_i^e(w)$, we may apply the inductive step to conclude

$$A_u(\Gamma^+\backslash e)\det\mathscr{L}(\Gamma^+\backslash e) = 2A_{\tilde{u}}(\widetilde{\Gamma^+\backslash e})$$
$$A_u(\Gamma^+\backslash E_i^e(w))\det\mathscr{L}(\Gamma^+\backslash E_i^e(w)) = 2A_{\tilde{u}}(\Gamma^+\widetilde{\backslash E_i^e(w)})$$

Note that $A_u(\Gamma^+\backslash E_i^e(w)) \neq 0$, since by assumption there exists at least one arborescence T_1 rooted at u using the edge e, so that T_1 remains an arborescence even after removing the edges $E_i^e(v)$. Similarly, $A_u(\Gamma^+\backslash e) \neq 0$. Therefore, we may apply Proposition 4.1 to conclude

$$A_w(\Gamma^+\backslash e)\det\mathscr{L}(\Gamma^+\backslash e) = A_{\tilde{w}}(\widetilde{\Gamma^+\backslash e})$$
$$A_w(\Gamma^+\backslash E_i^e(w))\det\mathscr{L}(\Gamma^+\backslash E_i^e(w)) = A_{\tilde{w}}(\Gamma^+\widetilde{\backslash E_i^e(w)})$$

Since e and the edges of $E_i^e(w)$ are elements of $E_i(w)$, we can apply the same arguments as we did in the original inductive step to show that $A_w(\Gamma^+\backslash e) = A_w(\Gamma^+\backslash E_i^e(w)) = A_w(\Gamma^+)$, that $A_{\tilde{w}}(\widetilde{\Gamma^+\backslash e}) + A_{\tilde{w}}(\Gamma^+\widetilde{\backslash E_i^e(w)}) = A_{\tilde{w}}(\tilde{\Gamma^+})$, and ultimately t

$$A_w(\Gamma^+)\det\mathscr{L}(\Gamma^+) = 2A_{\tilde{w}}(\tilde{\Gamma^+})$$

Since $A_w(\Gamma^+) \neq 0$—the auxiliary edge a ensures that any arborescence rooted at u may be modified into an arborescence rooted at w—we may reroot to conclude

$$A_u(\Gamma^+)\det\mathscr{L}(\Gamma^+) = 2A_{\tilde{u}}(\tilde{\Gamma^+})$$

Note that every arborescence $T \in \mathcal{T}_u(\Gamma)$ lifts uniquely to an arborescence $T^+ \in \mathcal{T}_u(\Gamma^+)$ not containing a, and conversely that every arborescence $T^+ \in \mathcal{T}_u(\Gamma^+)$ not containing a descends uniquely to an arborescence $T \in \mathcal{T}_u(\Gamma)$. We therefore perform the same trick that we did in the proof of Corollary 4.1. Let φ be the evaluation homomorphism mapping $\mathrm{wt}(a) \mapsto 0$. Then have $\varphi(A_u(\Gamma^+)) = A_u(\Gamma)$, $\varphi(\mathscr{L}(\Gamma^+)) = \mathscr{L}(\Gamma)$, and $\varphi(A_{\tilde{u}}(\tilde{\Gamma^+})) = A_{\tilde{u}}(\tilde{\Gamma})$. Since φ is a homomorphism, we conclude

$$A_u(\Gamma)\det\mathscr{L}(\Gamma) = 2A_{\tilde{u}}(\tilde{\Gamma})$$

Thus, the formula is proven.

4.4. **Rooted tree case.** Finally, we consider the case where

(1) No vertices with outdegree ≥ 2 root an arborescence;

(2) There exists at least one arborescence rooted at some vertex; *and*

(3) All vertices with outdegree 1 root no more than 1 arborescence?

In this case, Γ must have a structure similar to a rooted tree. Let u be a vertex with outdegree 1 that roots exactly one arborescence T. Without loss of generality, u is the only vertex of outdegree 1—we may contract along the unique outgoing edge e of any other such vertex u' to yield a graph with fewer edges otherwise, since

- Every arborescence of Γ rooted at u passes through e, so that that $A_u(\Gamma) = \mathrm{wt}(e)A_u(\Gamma/e)$;
- Every arborescence of $\tilde{\Gamma}$ passes through both lifts of e, so that $A_{\tilde{u}}(\tilde{\Gamma}) = \mathrm{wt}(e)^2 A_u(\widetilde{\Gamma/e})$; and
- $\det \mathscr{L}(\Gamma) = \mathrm{wt}(e) \det \mathscr{L}(\Gamma/e)$ via expansion by minors along the row corresponding to u'.

Therefore, the unique outgoing edge of u must be a loop, since otherwise the terminal vertex of this edge roots an arborescence, violating condition 1 above since all vertices other than u have outdegree ≥ 2. We may treat Γ as a Hasse diagram for the poset defined by T, with u the unique minimal element. Every other vertex v of Γ has exactly one edge belonging to the arborescence T, and all other edges of v must point to some $v \geq u$. Otherwise, a non-cyclic path from v to u distinct from the one given by T would exist, violating the uniqueness of the arborescence T.

Take any vertex $w \neq u$. Let e be the edge of w belonging to T. Define Γ^+ to be the graph obtained from Γ by adding an auxiliary edge a from u to w. Then we apply the same trick with the arborescences of the cover to conclude that $\Gamma^+\backslash e$ and $\Gamma^+\backslash E_i^e(w)$ satisfy the formula when rooted at u. Note that $A_u(\Gamma^+\backslash E_i^e(w))$ is never zero, since T remains an arborescence in $\Gamma^+\backslash E_i^e(w)$, so we may apply Corollary 4.1 to conclude that $A_v(\Gamma^+\backslash E_i^e(w)) \det \mathscr{L}(\Gamma^+\backslash E_i^e(w)) = A_{\tilde{w}}(\widetilde{\Gamma^+\backslash E_i^e(w)})$.

If any edge in $E_i^e(w)$ does not point towards w, it points to some vertex $w' > w$, so that w' roots an arborescence by modifying T to pass through a and the edge $(w, w') \in E_i^e(w)$. Since w' has outdegree ≥ 2 and roots at least one arborescence, we conclude that the desired identity also holds on $\Gamma^+\backslash e$ when rooting at w instead. In this case, we know the formula holds for $\Gamma^+\backslash e$ and $\Gamma^+\backslash E_i^e(w)$ when rooting at w, so now we may apply the same logic as the inductive step to conclude that the formula holds for Γ^+ when rooting at w, and therefore when rooting at u. Setting $\mathrm{wt}(a) = 0$ then shows that the formula holds for Γ rooting at u.

If this process goes through for at least one vertex $w \neq u$, then we are done. Otherwise, we conclude that edge set of Γ consists only of the tree T plus loops, in which case we may prove the formula directly. Without loss of generality, all loops are negative, since positive loops do not contribute to either the negative vector fields of Γ nor the arborescences of $\tilde{\Gamma}$. Then every arborescence of $\tilde{\Gamma}$ contains at least one lift of every edge in T, but this is the only condition on the arborescences—as long as the lift of the same negative loop is not used twice, there can be no cycles. For every loop besides the one at u, there are two choices of lifts. Thus, for each negative vector field $\gamma \subseteq \Gamma$, we obtain $2^{\#C(\gamma)-1}$ arborescences of $\tilde{\Gamma}$—one factor of two for each loop of γ other than the one at u. Since this process uniquely describes all arborescences of $\tilde{\Gamma}$, we have

$$\begin{aligned} 2A_{\tilde{u}}(\tilde{\Gamma}) &= \mathrm{wt}(T) \sum_{\gamma \subseteq \Gamma} 2^{\#C(\gamma)} \mathrm{wt}(\gamma) \\ &= A_u(\Gamma) \det \mathscr{L}(\Gamma) \end{aligned}$$

This exhausts all possible exceptions to the inductive step, completing the proof.

■

5. Generalization to other covers

The preceding proof unfortunately does not generalize to k-fold covers for $k > 2$—the main inductive step fails because there are too many lifts of an outgoing edge of v, even if we disregard the outgoing edges of $\tilde{v}$. In this section, we attempt a more algebraic approach that generalizes to higher covers.

5.1. **Cyclic Prime Covers.** We now restrict ourselves in the case when the voltage group G is a cyclic group of prime order: let $G = \mathbb{Z}/p\mathbb{Z} = \{1, g, g^2, \cdots, g^{p-1}\}$, where p is a prime number..

For a graph Γ and its cyclic p-cover $\tilde{\Gamma}$, we have the following conjecture as a generalization of Theorem 3.3.

Conjecture 5.1. *Let $\Gamma = (V, E)$ be a weighted $\mathbb{Z}/p\mathbb{Z}$-voltage directed multigraph, and let $\tilde{\Gamma}$ be its derived graph. Let K be the reduced group algebra of $\mathbb{Z}/p\mathbb{Z}$ over $\mathbb{Q}$, so that $K \cong \mathbb{Q}(\zeta_p)$, where ζ_p is a primitive p-th root of unity. Take any vertex v in Γ, and any lift of of v, $\tilde{v} \in \tilde{\Gamma}$. Then*

$$A_{\tilde{v}}(\tilde{\Gamma}) = \frac{1}{p} A_v(\Gamma) N_{K/\mathbb{Q}}(\det \mathscr{L}(\Gamma))$$

where $N_{K/\mathbb{Q}}(\det \mathscr{L}(\Gamma))$ denotes the field norm of $\det \mathscr{L}(\Gamma)$. Equivalently,

$$A_{\tilde{v}}(\tilde{\Gamma}) = \frac{1}{p} A_v(\Gamma) \prod_{i=1}^{p-1} [\sigma_i(\det \mathscr{L}(\Gamma)]$$

where σ_i is the field automorphism that maps ζ_p to ζ_p^i.

5.2. **Change of Basis.** Another possible direction that could give us the desired arborescence ratio is via a change of basis of the Laplacian of $\tilde{\Gamma}$. Here we consider any regular k cover of a graph Γ. Throughout this section, we fix the basis vector of the Laplacian matrix (viewed as endomorphism of a vector space) to be $(v_1^{(1)}, v_2^{(1)}, \cdots, v_n^{(1)}, v_1^{(2)}, \cdots, v_n^{(2)}, \cdots, v_1^{(n)}, \cdots, v_n^{(n)})$, where $\{v_1, \cdots, v_n\}$ is the set of vertices of graph G.

Let $\Gamma = (V, E)$ be a graph, $\tilde{\Gamma}$ a regular k-fold cover of Γ, let $L(\tilde{\Gamma})$ denote the regular Laplacian of $\tilde{\Gamma}$. Suppose $|V| = E$ and let I_n be the $n \times n$ identity matrix.

Lemma 5.2. *Let S be a $kn \times kn$ matrix consists of $k \times k$ blocks:*

$$S := \begin{pmatrix} I_n & 0 & 0 & \cdots & 0 & 0 \\ -I_n & I_n & 0 & \cdots & 0 & 0 \\ 0 & -I_n & I_n & \cdots & 0 & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & \cdots & -I_n & I_n \end{pmatrix}$$

Let $L' = S\mathscr{L}(\Gamma)S^{-1}$, then L' is of the form

$$L' = \begin{pmatrix} \mathscr{L}(G) & F \\ 0 & D \end{pmatrix}$$

where $\mathscr{L}(\Gamma)$ is the Laplacian of the graph Γ and D is some $n(k-1)$ by $n(k-1)$ matrix. In case of $k = 2$, the matrix D is just the signed Laplacian of G.

Proof. Notice that the conjugation by matrix S changes the basis of the Laplacian matrix to

$$\left(\sum_{i=1}^{n} v_1^{(i)}, \cdots, \sum_{i=1}^{n} v_n^{(i)}, \sum_{i=2}^{n} v_1^{(i)}, \cdots, \sum_{i=2}^{n} v_n^{(i)}, \sum_{i=3}^{n} v_1^{(i)}, \cdots, \sum_{i=3}^{n} v_n^{(i)}, \cdots, \cdots, \cdots, v_1^{(n)}, \cdots, v_n^{(n)}\right).$$

∎

Conjecture 5.3. *Let D be the matrix derived as in lemma 4.5, then for any $v \in \Gamma$ and its lift $\tilde{v}$ in $\tilde{\Gamma}$, we have*

$$\mathcal{A}_{\tilde{v}}(\tilde{\Gamma}) = \mathcal{A}_v(\Gamma) \det D$$

Proof. This conjecture has been turned into a theorem, whose proof will be updated later. ∎

References

[Cha82] Seth Chaiken. A combinatorial proof of the all minors matrix tree theorem. *Siam Journal of Algebraic and Discrete Methods*, 3(3):35–42, 1982.

[GP17] Pavel Galashin and Pavlo Pylyavskyy. R-systems. *arXiv e-prints*, page arXiv:1709.00578, Sep 2017.

[GT75] Jonathan Gross and Thomas Tucker. Generating all graph coverings by permutation voltage assignments. *Discrete Mathematics*, pages 273–283, Aug 1975.

[RT13] Victor Reiner and Dennis Tseng. Critical groups of covering, voltage, and signed graphs. *arXiv e-prints*, page arXiv:1301.2977, Jan 2013.

[SW86] Dennis Stanton and Dennis White. *Constructive Combinatorics*. Springer-Verlag, 1986.