

Ramanujan Coverings of Graphs

Chris Hall*, Doron Puder[†] and William F. Sawin[‡]

June 9, 2015

Abstract

Let G be a finite connected graph, and let ρ be the spectral radius of its universal cover. For example, if G is k -regular then $\rho = 2\sqrt{k-1}$. We show that for every d , there is a d -sheeted covering of G where all the new eigenvalues are bounded from above by ρ . It follows that a bipartite Ramanujan graph has a d -sheeted Ramanujan covering for every d . This generalizes the $d = 2$ case due to Marcus, Spielman and Srivastava [MSS15a].

Every d -sheeted covering of G corresponds to a labeling of the edges of G by elements of the symmetric group S_d . We generalize this notion to labeling the edges by various groups and present a broader scenario where Ramanujan coverings are guaranteed to exist.

An important ingredient of our proof is a new generalization of the matching polynomial of a graph. We define the d -th matching polynomial of G to be the average matching polynomial of all d -coverings of G . We show this polynomial shares many properties with the original matching polynomial. For example, it is real-rooted with all its roots inside $[-\rho, \rho]$.

Inspired by [MSS15a], a crucial component of our proof is the existence of interlacing families of polynomials for complex reflection groups. The core argument of this component is taken from [MSS15c].

*Author Hall was partially supported by Simons Foundation award 245619 and IAS NSF grant DMS-1128155.

[†]Author Puder was supported by the Rothschild fellowship and by the National Science Foundation under agreement No. DMS-1128155.

[‡]Author Sawin was supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1148900.

Contents

1	Introduction	2
2	Background, Preliminary Claims and Outline of the Proofs	8
2.1	Expander and Ramanujan Graphs	8
2.2	The d -Matching Polynomial	10
2.3	Group Labelings of Graphs	13
2.4	Group Representations	13
2.5	Interlacing Families of Polynomials	15
3	Property $(\mathcal{P}1)$ and the Proof of Theorem 1.11	17
3.1	Determinant of Sum of Matrices	18
3.2	Matrix Coefficients	19
3.3	Proof of Theorem 1.11	20
4	Property $(\mathcal{P}2)$ and the Proof of Theorem 1.12	22
4.1	Average Characteristic Polynomial of Sum of Random Matrices	23
4.2	Average Characteristic Polynomial of Random Coverings	25
4.3	Proof of Theorem 1.12	27
5	On Pairs Satisfying $(\mathcal{P}1)$ and $(\mathcal{P}2)$ and Further Applications	28
5.1	Complex Reflection Groups	28
5.2	Pairs Satisfying $(\mathcal{P}1)$	30
5.3	Applications of Theorem 1.10	30
5.4	Ramanujan Topological Coverings of Special Kinds	32
6	Open Questions	32

1 Introduction

Ramanujan Coverings

Let G be a finite, connected, undirected graph on n vertices and let A_G be its adjacency matrix. The eigenvalues of A_G are real and we denote them by

$$\lambda_n \leq \dots \leq \lambda_2 \leq \lambda_1 = \mathfrak{pf}(G),$$

where $\lambda_1 = \mathfrak{pf}(G)$ is the Perron-Frobenius eigenvalue of A_G , referred to as the trivial eigenvalue (for example, $\mathfrak{pf}(G) = k$ for G k -regular). The smallest eigenvalue, λ_n , is at least $-\mathfrak{pf}(G)$, with equality if and only if G is bipartite. Denote by $\lambda(G)$ the largest absolute value of a non-trivial eigenvalue, namely $\lambda(G) = \max(\lambda_2, -\lambda_n)$. It is well known that $\lambda(G)$ provides a good estimate to different expansion properties of G : the smaller $\lambda(G)$ is, the better expanding G is (see [HLW06, Pud15]).

However, $\lambda(G)$ cannot be arbitrarily small. Let $\rho(G)$ be the spectral radius of the universal covering tree of G (for example, $\rho(G) = 2\sqrt{k-1}$ when G is k -regular). It is known that $\lambda(G)$ cannot be much smaller than $\rho(G)$, so graphs with $\lambda(G) \leq \rho(G)$ are considered optimal expanders (we elaborate in Section 2.1 below). Following [LPS88] they are called **Ramanujan graphs**, and the interval $[-\rho(G), \rho(G)]$ called **the Ramanujan interval**. In the bipartite case, $\lambda(G) = |\lambda_n| = \mathbf{pf}(G)$ is large, but G can still expand well in many senses (see Section 2.1), and the optimal scenario is when all other eigenvalues are within the Ramanujan interval, namely, when $\lambda_{n-1}, \lambda_{n-2}, \dots, \lambda_2 \in [-\rho(G), \rho(G)]$. We call a bipartite graph with this property a **bipartite-Ramanujan graph**.

Let H be a topological d -sheeted covering of G (d -covering in short) with covering map $p : H \rightarrow G$. If $f : V(G) \rightarrow \mathbb{R}$ is an eigenfunction of G , then $f \circ p$ is an eigenfunction of H with the same eigenvalue. Thus, out of the dn eigenvalues of H (considered as a multiset), n are induced from G and are referred to as **old eigenvalues**. The other $(d-1)n$ are called the **new eigenvalues** of H .

Definition 1.1. Let H be a topological covering of G . We say that H is a **Ramanujan Covering** of G if all the new eigenvalues of H are in $[-\rho(G), \rho(G)]$. We say H is a **one-sided Ramanujan Covering** if all the new eigenvalues are bounded from above¹ by $\rho(G)$.

The existence of infinitely many k -regular Ramanujan graphs for every $k \geq 3$ is a long-standing open question. Bilu and Linial [BL06] suggest the following approach to solving this conjecture: start with your favorite k -regular Ramanujan graph (e.g. the complete graph on $k+1$ vertices) and construct an infinite tower of Ramanujan 2-coverings. They conjecture that every (regular) graph has a Ramanujan 2-covering. This approach turned out to be very useful in the groundbreaking result of Marcus, Spielman and Srivastava [MSS15a], who proved that every graph has a one-sided Ramanujan 2-covering. This translates, as explained below, to that there are infinitely many k -regular *bipartite* Ramanujan graphs of every degree k .

In this paper, we generalize the result of [MSS15a] to coverings of every degree:

Theorem 1.2. *Every connected, loopless graph has a one-sided Ramanujan d -covering for every d .*

In fact, this result holds also for graphs with loops, as long as they are regular (Proposition 2.3), so the only obstruction is irregular graphs with loops. We stress that throughout this paper, all statements involving graphs hold not only for simple graphs, but also for graphs with multiple edges. Unless otherwise stated, the results also hold for graphs with loops.

A finite graph is bipartite if and only if its spectrum is symmetric around zero. In addition, any covering of a bipartite graph is bipartite. Thus, every one-sided Ramanujan covering of a bipartite graph is, in fact, a (full) Ramanujan covering. Therefore,

Corollary 1.3. *Every connected bipartite graph has a Ramanujan d -covering for every d .*

¹We could also define a one-sided Ramanujan covering as having all its eigenvalues bounded from below by $-\rho(G)$. Every result stated in the paper about these coverings would still hold for the lower-bound case, unless stated otherwise.

In the special case where the base graph is $\bullet \begin{smallmatrix} \text{---} \\ \text{---} \\ \text{---} \end{smallmatrix} \bullet$ (two vertices with k edges connecting them), Theorem 1.2 (and Corollary 1.3) were shown in [MSS15c]. In this regard, our result generalizes the 2-coverings result from [MSS15a] as well as the newer result from [MSS15c]. Corollary 1.3 also yields the existence of many more *simple* bipartite Ramanujan graphs than was known before (see Corollary 2.2).

Generalized Matching Polynomials

An important ingredient in our proof of Theorem 1.2 is a new family of polynomials associated to a given graph. These polynomials generalize the well-known matching polynomial of a graph defined by Heilmann and Lieb [HL72]: let m_i to be the number of matchings in G with i edges, and set $m_0 = 1$. The matching polynomial of G is

$$\sum_{i=0}^{\lfloor n/2 \rfloor} (-1)^i m_i x^{n-2i} \in \mathbb{Z}[x].$$

Definition 1.4. Let $d \in \mathbb{Z}_{\geq 1}$. The **d -matching polynomial** of a finite graph G , denoted $\mathcal{M}_{d,G}$, is the average of the matching polynomials of all d -coverings of G .

In Definition 1.6 below we give the precise definition for the family of d -coverings of G we have in mind. Of course, $\mathcal{M}_{1,G}$ is the usual matching polynomial of G (a graph is the only 1-covering of itself). Note that these generalized matching polynomials of G are monic, but their other coefficients need not be integer valued. However, they seem to share many of the nice properties of the usual matching polynomial. For our cause, the following property is crucial:

Theorem 1.5. *Let G be a finite, connected², loopless graph. For every $d \in \mathbb{Z}_{\geq 1}$, the polynomial $\mathcal{M}_{d,G}$ is real rooted with all its roots contained in the Ramanujan interval $[-\rho(G), \rho(G)]$.*

This result for $d = 1$ goes back to [HL72]. (They showed that $\mathcal{M}_{1,G}$ satisfies this statement whenever G is regular. Apparently, the result concerning $\mathcal{M}_{1,G}$ for irregular G was first noticed in [MSS15a], even though some of the original proofs of [HL72] work in the irregular case as well.)

Below (Theorem 1.11) we show that these generalized matching polynomials are equal to the average of a family of characteristic polynomials. This, again, generalizes a well-known property of $\mathcal{M}_{1,G}$. We also give a precise formula for $\mathcal{M}_{d,G}$ (Proposition 2.7).

Covering Graphs by General Matrix Groups

As mentioned above, we suppose that G is undirected, yet we regard it as an oriented graph. More precisely, we choose an orientation for each edge in G , and we write $E^+(G)$ for the $E^+(G)$

²Connectivity here is required only because of the way $\rho(G)$ was defined. The real-rootedness holds for any finite graph. In the general case, the d -matching polynomial is the product of the d -matching polynomials of the different connected components, and $\rho(G)$ can be defined as the maximum of $\rho(G_i)$ over the different components G_i of G .

resulting set of oriented edges and $E^-(G)$ for the edges with the opposite orientation. Finally, if e is an edge in $E^\pm(G)$, then we write $-e$ for the corresponding edge in $E^\mp(G)$ with the opposite orientation, and we identify $E(G)$ with the disjoint union $E^+(G) \sqcup E^-(G)$. We let $h(e)$ and $t(e)$ denote the head vertex and tail vertex of $e \in E(G)$, respectively. We say that G is an *oriented undirected graph*.

Throughout this paper, the family of d -coverings of the graph G is defined via the following natural model, introduced in [AL02] and [Fri03]. The vertex set of every d -covering H is $\{v_i \mid v \in V(G), 1 \leq i \leq d\}$. Its edges are defined via a function $\sigma : E(G) \rightarrow S_d$ satisfying $\sigma(-e) = \sigma(e)^{-1}$ (occasionally, we denote $\sigma(e)$ by σ_e): for every $e \in E^+(G)$ we introduce in H the d edges connecting $h(e)_i$ to $t(e)_{\sigma_e(i)}$ for $1 \leq i \leq d$.

Definition 1.6. Denote by $\mathcal{C}_{d,G}$ the probability space consisting of all d -coverings $\{\sigma : E(G) \rightarrow S_d \mid \sigma(-e) = \sigma(e)^{-1}\}$, endowed with uniform distribution.

Let $H \in \mathcal{C}_{d,G}$ correspond to $\sigma : E(G) \rightarrow S_d$ and let $f : V(H) \rightarrow \mathbb{C}$ be an eigenfunction of H with eigenvalue μ . For every $v \in V(G)$, let f_v be the transpose of the vector $(f(v_1), f(v_2), \dots, f(v_d))$. Considering the permutations σ_e as permutation matrices, the collection of vectors $\{f_v\}_{v \in V(G)}$ satisfies the following equation for every $v \in V(G)$:

$$\sum_{e: t(e)=v} \sigma_e f_{h(e)} = \mu \cdot f_v \quad (1.1)$$

(note that every loop at v appears twice in the summation, once in each orientation.) Conversely, every function $f : V(G) \rightarrow \mathbb{C}^d$ satisfying (1.1) for some fixed μ and every $v \in V(G)$, is an eigenfunction of H with eigenvalue μ .

This way of presenting coverings of G and their spectra suggests the following natural generalization: instead of picking the matrices σ_e from the group of permutation matrices, one can label the edges of G by matrices from any fixed subgroup of $\text{GL}_d(\mathbb{C})$. Since the same group Γ may be embedded in several different ways in $\text{GL}_d(\mathbb{C})$, or even in $\text{GL}_d(\mathbb{C})$ for varying d , the right notion here is that of group representations. Namely, a group Γ together with finite dimensional representation π , which is simply a homomorphism $\pi : \Gamma \rightarrow \text{GL}_d(\mathbb{C})$ (in this case we say that π is d -dimensional).

Given a pair (Γ, π) , where π is d -dimensional, we define a **Γ -labeling** to be a function $\gamma : E(G) \rightarrow \Gamma$ satisfying $\gamma(-e) = \gamma(e)^{-1}$. The **π -spectrum** of the Γ -labeling γ is defined accordingly as the values μ satisfying

$$\sum_{e: t(e)=v} \pi(\gamma(e)) f_{h(e)} = \mu \cdot f_v \quad \forall v \in V(G)$$

for some $0 \neq f : V(G) \rightarrow \mathbb{C}^d$. More concretely, it is the spectrum of the $nd \times nd$ matrix $\mathbf{A}_{\gamma, \pi}$ obtained from A_G , the adjacency matrix of G , as follows: for every $u, v \in V(G)$, replace the (u, v) entry in A_G by the $d \times d$ block $\sum_{e: u \rightarrow v} \pi(\gamma(e))$ (the sum is over all edges from u to v , and is a zero $d \times d$ block if there are no such edges). It is easy to see that whenever π is a unitary representation (this is the case, for example, whenever Γ is finite), the spectrum of $A_{\gamma, \pi}$ is real (see Claim 2.9).

Definition 1.7. Let Γ be a group. A Γ -labeling of the graph G is a function $\gamma : E(G) \rightarrow \Gamma$ satisfying $\gamma(-e) = \gamma(e)^{-1}$. If Γ is finite, we let $\mathcal{C}_{\Gamma, G}$ be the probability space of all Γ -labelings of G endowed with uniform distribution. More generally, if Γ is compact, we let $\mathcal{C}_{\Gamma, G}$ be the probability space of all Γ -labelings of G endowed with Haar measure on $\Gamma^{E^+(G)}$.

Let $\pi : \Gamma \rightarrow \text{GL}_d(\mathbb{C})$ a unitary representation of Γ . For any Γ -labeling γ we say that $A_{\gamma, \pi}$ is a (Γ, π) -**covering** of the graph G . The *spectrum* of this (Γ, π) -covering is a multiset denoted $\text{spec}(A_{\gamma, \pi})$. The (Γ, π) -covering $A_{\gamma, \pi}$ is said to be **Ramanujan** if $\text{spec}(A_{\gamma, \pi}) \subseteq [-\rho(G), \rho(G)]$, and **one-sided Ramanujan** if all the eigenvalues of $A_{\gamma, \pi}$ are at most $\rho(G)$.

For example, if G consists of a single vertex with r loops, and π is the *regular* representation³ of Γ , then a (Γ, π) -covering $A_{\gamma, \pi}$ of G is equivalent to the Cayley graph of Γ with respect to the set $\gamma(E(G))$. The Cayley graph is Ramanujan if and only if the corresponding $(\Gamma, \pi - \text{triv})$ -covering is Ramanujan (see Section 2.4).

As another example, the symmetric group S_d has an obvious d -dimensional representation associating to every $\sigma \in S_d$ the corresponding permutation matrix. But S_d also has a $(d-1)$ -dimensional representation, called the *standard* representation, and denoted std (std is, in fact, an irreducible component of the former - see Section 2.4). Every d -covering of G corresponds to a unique (S_d, std) -covering, and, moreover, the new spectrum of the d -covering is precisely the spectrum of the corresponding (S_d, std) -covering (see Claim 2.10). In particular, a Ramanujan d -covering corresponds to a Ramanujan (S_d, std) -covering. The following is, then, a natural generalization of the question concerning ordinary Ramanujan coverings of graphs:

Question 1.8. *For which pairs (Γ, π) of a group Γ with a unitary representation $\pi : \Gamma \rightarrow \text{GL}_d(\mathbb{C})$ is it guaranteed that every connected graph G has a (one-sided / full) Ramanujan (Γ, π) -covering?*

In this paper, we find two group-theoretic properties of pairs (Γ, π) , which, together, guarantee the existence of a *one-sided* Ramanujan covering for every connected graph G . To define the second property we need the notion of *pseudo-reflections*: a matrix $A \in \text{GL}_d(\mathbb{C})$ is called a **pseudo-reflection** if A has finite order and $\text{rank}(A - I) = 1$. Equivalently, A is a pseudo-reflection if it is conjugate to a diagonal matrix of the form

$$\begin{pmatrix} \lambda & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{pmatrix}$$

with $\lambda \neq 1$ and of finite order.

Definition 1.9. Let Γ be a group and $\pi : \Gamma \rightarrow \text{GL}_d(\mathbb{C})$ a unitary representation. We say that:

- (Γ, π) satisfies **(P1)** if Γ is finite or compact and if all exterior powers $\bigwedge^r(\pi)$, $0 \leq r \leq d$, **(P1)**

³Namely, π is a $|\Gamma|$ -dimensional representation, and for every $g \in \Gamma$, the matrix $\pi(g)$ is the permutation matrix describing the action of g on the elements of Γ by left multiplication.

are irreducible and non-isomorphic⁴.

- (Γ, π) satisfies **(P2)** if Γ is finite and if $\pi(\Gamma)$ is a complex reflection group, namely, if **(P2)** it is generated by pseudo-reflections.

(Property **(P2)** may also be generalized to compact groups - see Remark 4.8.)

As we explain in Section 2.4, by showing that (S_d, std) satisfies **(P1)** and **(P2)**, Theorem 1.2 becomes a special case of the following:

Theorem 1.10. *Let Γ be a finite group and $\pi : \Gamma \rightarrow \text{GL}_d(\mathbb{C})$ a representation such that (Γ, π) satisfies **(P1)** and **(P2)**. Then every connected, loopless graph G has a one-sided Ramanujan (Γ, π) -covering.*

The proof of Theorem 1.10 follows the general proof strategy from [MSS15a], and each of the properties **(P1)** and **(P2)** is needed for different parts of this proof. Property **(P1)** is needed in the part where the role of the generalized matching polynomials in the proof arises. To describe it, we denote by $\phi_{\gamma, \pi}$ the characteristic polynomial of the (Γ, π) -covering $A_{\gamma, \pi}$, namely

$$\phi_{\gamma, \pi}(x) \stackrel{\text{def}}{=} \det(xI - A_{\gamma, \pi}) = \prod_{\mu \in \text{spec}(A_{\gamma, \pi})} (x - \mu). \quad (1.2)$$

Along this paper, the default distribution on Γ -labelings of a graph G is the one defined by $\mathcal{C}_{\Gamma, G}$ (see Definition 1.7). Hence, when Γ and G are understood from the context, we use the notation $\mathbb{E}_{\gamma}[\phi_{\gamma, \pi}(x)]$ to denote the expected characteristic polynomial of a random (Γ, π) -covering, the expectation being over the space $\mathcal{C}_{\Gamma, G}$ of Γ -labelings.

Theorem 1.11. *Let the graph G be connected. For every pair (Γ, π) satisfying **(P1)** with π being d -dimensional, the following holds:*

$$\mathbb{E}_{\gamma}[\phi_{\gamma, \pi}(x)] = \mathcal{M}_{d, G}(x).$$

*It particular, as long as **(P1)** holds, $\mathbb{E}_{\gamma}[\phi_{\gamma, \pi}(x)]$ depends only on d and not on (Γ, π) .*

This generalizes an old result from [GG81] for the case $d = 1$, $\Gamma = \mathbb{Z}/2\mathbb{Z}$ and $\pi(\Gamma) = \{\pm 1\}$, which is used in [MSS15a]. Together with Theorem 1.5, we get that whenever (Γ, π) satisfies **(P1)** and G has no loops, the expected characteristic polynomial $\mathbb{E}_{\gamma}[\phi_{\gamma, \pi}(x)]$ has only real-roots, all of which lie inside the Ramanujan interval.

The second part of the proof of Theorem 1.10 shows the role of **(P2)**:

Theorem 1.12. *Let G be a finite, loopless graph. For every pair (Γ, π) satisfying **(P2)**, the following hold:*

- $\mathbb{E}_{\gamma}[\phi_{\gamma, \pi}(x)]$ is real-rooted.

⁴See Section 2.4 for a definition of an exterior power, an irreducible representation and isomorphism of representations.

- *There exists a (Γ, π) -covering $A_{\gamma, \pi}$ with largest eigenvalue at most the largest root of $\mathbb{E}_{\gamma} [\phi_{\gamma, \pi}(x)]$.*

The proof of Theorem 1.12 is based on showing that the family of polynomials $\phi_{\gamma, \pi}(x)$ is interlacing. The core of the argument here is inspired by [MSS15c, Theorem 3.3]. We explain more in Section 2.5 and in Section 4.

The paper is organized as follows: In Section 2 we give more background details, prove some preliminary results and sketch an outline of the remaining proofs. Section 3 is dedicated to property (P1) and the proof of Theorem 1.11, while in Section 4 we study Property (P2) and prove Theorem 1.12. In Section 5, we study groups satisfying the two properties and present further combinatorial applications of Theorem 1.10. We end (Section 6) with a list of open questions arising from the discussion in this paper.

2 Background, Preliminary Claims and Outline of the Proofs

In this section we give more background material, prove some preliminary claims, reduce all the results from Section 1 to the proofs of Theorems 1.11 and 1.12 and give a short outline of these two remaining proofs.

2.1 Expander and Ramanujan Graphs

As in Section 1, let G be a finite connected graph on n vertices and A_G its adjacency matrix. Recall that $\mathbf{pf}(G)$ is the Perron-Frobenius eigenvalue of A_G , that $\lambda_n \leq \dots \leq \lambda_2 \leq \lambda_1 = \mathbf{pf}(G)$ are its entire spectrum, and that $\lambda(G) = \max(\lambda_2, -\lambda_n)$. The graph G is considered to be well-expanding if it is “highly” connected. This can be measured by different combinatorial properties of G , most commonly by its Cheeger constant, by the rate of convergence of a random walk on G , and by how much the number of edges between any two sets of vertices approximates the corresponding number in a random graph (the so-called Expander Mixing Lemma)⁵. All these properties can be measured, at least approximately, by the spectrum of G , and especially by $\lambda(G)$ and the spectral gap $\mathbf{pf}(G) - \lambda(G)$: the smaller $\lambda(G)$ and the bigger the spectral gap is, the better expanding G is⁶. (See [HLW06] and [Pud15, Appendix B] and the references therein.)

Yet, $\lambda(G)$ cannot be arbitrarily small. Let T be the universal covering tree of G . We think of all the finite graphs covered by T as one family. For example, for any $k \geq 2$, all finite k -regular graphs constitute a single such family of graphs: they are all covered by the k -regular tree. It turns out that the spectral radius of T , denoted $\rho(G)$, plays an important role in the theory of expansion of the corresponding family of graphs. This number is the

⁵In this sense, Ramanujan graphs resemble random graphs. The converse is also true in certain regimes of random graphs: see [Pud15] and the references therein.

⁶More precisely, the Cheeger inequality relates the Cheeger constant of a graph with the values of $\lambda_2(G)$.

spectral radius of the adjacency operator A_T acting on $\ell^2(V(T))$ by

$$(A_T f)(v) = \sum_{u \sim v} f(u).$$

For the k -regular tree, this spectral radius is $2\sqrt{k-1}$.

Theorem 2.1. [Gre95, Thm 2.11] *Let T be a tree with finite quotients and ρ its spectral radius. For every $\varepsilon > 0$, there exists $c = c(T, \varepsilon)$, $0 < c < 1$, such that if G is a finite graph with n vertices which is covered by T , then at least cn of its eigenvalues satisfy $\lambda_i \geq \rho - \varepsilon$. In particular, $\lambda(G) \geq \rho - o_n(1)$ (with the $o_n(1)$ term depending only on T).*

The last statement of the theorem, restricted to regular graphs, is due to Alon-Boppana [Nil91]. Thus, graphs G satisfying $\lambda(G) \leq \rho(G)$ are considered to be optimal expanders. Following the terminology from [LPS88], they are named Ramanujan graphs.

Lubotzky [Lub94, Problem 10.7.3] asked whether for every $k \geq 3$ there are infinitely many k -regular Ramanujan graphs⁷. In the regular case, every family has at least one Ramanujan graph (e.g. the complete graph on $k+1$ vertices). Other families may contain no Ramanujan graphs at all. For example, the family of (c, d) -biregular graphs, all covered by the (c, d) -biregular tree, consists entirely of bipartite graphs, so none of them is Ramanujan in the strict sense. Other families with no Ramanujan graphs, not even bipartite-Ramanujan, are shown to exist in [LN98]. In these cases there are certain “bad” eigenvalues outside the Ramanujan interval appearing in every finite graph in the family.

Still, it makes sense to look for optimal expanders under these constraints. These are precisely those graphs where all other eigenvalues lie in the Ramanujan interval. For example, bipartite Ramanujan graphs are optimal expanders in many combinatorial senses within the family of bipartite graphs (e.g. [GP14, Lemma 10]). The strategy of constructing Ramanujan coverings fits this general goal: find any graph in the family which is optimal (has all its values in the Ramanujan interval except for the bad ones) and construct Ramanujan coverings to obtain more optimal graphs in the same family. Of course, (connected) coverings of a graph G are covered by the same tree as G .

Marcus, Spielman and Srivastava have already shown that every graph has a one-sided Ramanujan 2-covering [MSS15a]. Thus, if a family of graphs contains at least one Ramanujan graph (bipartite or not), then it has infinitely many bipartite Ramanujan graphs⁸. They have more recently showed [MSS15c] that for any $k \geq 3$, the graph $\bullet \begin{smallmatrix} \text{---} \text{---} \text{---} \end{smallmatrix} \bullet$ (two vertices with k edges connecting them) has a Ramanujan d -covering for every d . It follows there are k -regular bipartite Ramanujan graphs, not necessarily simple, on $2d$ vertices for every d . Theorem 1.2 proves there is a richer family of bipartite Ramanujan graphs than was known before.

⁷More precisely, Lubotzky’s original definition of Ramanujan graphs included also *bipartite* Ramanujan graphs. Thus, [MSS15a] answered this question to the positive.

⁸Given a Ramanujan graph, its “double cover” - the 2-covering with all permutations being non-identity - is bipartite Ramanujan.

Corollary 2.2. *Every family of graphs (defined by a common universal covering tree) containing a (simple) bipartite Ramanujan graph on n vertices, also contains (simple, respectively) bipartite Ramanujan graphs on nd vertices for every $d \in \mathbb{Z}_{\geq 1}$.*

In particular, there is a simple k -regular, bipartite Ramanujan graph on $2kd$ vertices for every d .

The last statement follows by constructing Ramanujan d -coverings of the full bipartite graph on $2k$ vertices, which is Ramanujan.

As of now, we cannot extend all the results in this papers to graphs with loops (and see Question 6.5). However, we can extend Theorem 1.2 to regular graphs with loops. We now give the short proof of this extension, assuming Theorem 1.2:

Proposition 2.3. *Let G be a regular finite graph, possibly with loops. Then G has a one-sided Ramanujan d -covering for every d .*

We remark that in this proposition the proof does not yield the analogous result for coverings with new spectrum bounded from below by $-\rho(G)$.

Proof. Let G be any finite connected graph with n vertices and m edges. Subdivide each of its edges by introducing a new edges in its middle, to obtain a new, bipartite graph H , with n vertices on one side and m on the other. Clearly, there is a one-to-one correspondence between (isomorphism types of) d -coverings of G and (isomorphism types of) d -coverings of H . It is easy to see that H has eigenvalue 0 with multiplicity (at least) $m - n$. The remaining $2n$ eigenvalues are symmetric around zero, and their squares are the eigenvalues of $A_G + D_G$, where A_G is the adjacency matrix of G and D_G is diagonal with the degrees of the vertices.

If G is k -regular, this means that if μ is an eigenvalue of G , then $\pm\sqrt{\mu + k}$ are eigenvalues of H (and these are precisely all the eigenvalues of H , aside to the $m - n$ zeros). By Theorem 1.2, H has a Ramanujan d -covering \hat{H}_d for every d . Since the spectral radius of the $(k, 2)$ -biregular tree is⁹ $\sqrt{k - 1} + 1$, every eigenvalue μ of the corresponding d -covering \hat{G}_d satisfies $\sqrt{\mu + k} \leq \sqrt{k - 1} + 1$, i.e. $\mu \leq 2\sqrt{k - 1}$. \square

The exact same argument can be used to extend also the statement of Theorem 1.10 to regular graphs with loops: if G is regular, possibly with loops, and (Γ, π) satisfies $(\mathcal{P}1)$ and $(\mathcal{P}2)$, then G has a one-sided Ramanujan (Γ, π) -covering.

2.2 The d -Matching Polynomial

The following is a crucial ingredient in the proof of the main result of [MSS15a]:

Theorem 2.4. *[[HL72] for the regular case, [MSS15a] for the general case] The ordinary matching polynomial $\mathcal{M}_{1,G}$ of every finite connected graph G is real-rooted with all its roots lying in the Ramanujan interval $[-\rho(G), \rho(G)]$.*

⁹In general, the spectral radius of the (c, d) -biregular tree is $\sqrt{c - 1} + \sqrt{d - 1}$.

Recall that $\mathcal{M}_{d,G}$, the d -matching polynomial of the graph G , is defined as the average of the matching polynomials in the space of d -coverings $\mathcal{C}_{d,G}$. Every such covering belongs to the same family as G (even when the covering is not connected, each component is covered by the same tree as G). Thus, we obtain:

Corollary 2.5. *All real roots of $\mathcal{M}_{d,G}$ are inside the Ramanujan interval $[-\rho(G), \rho(G)]$.*

Proof. Recall that n denotes the number of vertices of G . The ordinary matching polynomial of every $H \in \mathcal{C}_{d,G}$ is a degree- nd monic polynomial. By Theorem 2.4, it is strictly positive in the interval $(\rho(G), \infty)$, and is either strictly positive or strictly negative in $(-\infty, -\rho(G))$ depending only on the parity of nd . The corollary now follows by the definition of $\mathcal{M}_{d,G}$ as the average of such polynomials. \square

The proof of Theorem 1.5 boils down, then, to showing that $\mathcal{M}_{d,G}$ is real-rooted. For this, we use the full strength of Theorems 1.11 and 1.12. We show (Fact 2.11 below) that for every d , there is a pair (Γ, π) of a group Γ and a d -dimensional representation π which satisfies both $(\mathcal{P}1)$ and $(\mathcal{P}2)$. From theorem 1.11 we obtain that $\mathcal{M}_{d,G}(x) = \mathbb{E}_\gamma[\phi_{\gamma,\pi}(x)]$ (the expectation over $\mathcal{C}_{\Gamma,G}$), and from Theorem 1.12 we obtain that $\mathbb{E}_\gamma[\phi_{\gamma,\pi}(x)]$ is real-rooted. We wonder if there is a more direct proof of the real-rootedness of $\mathcal{M}_{d,G}(x)$ (see Question 6.4).

In the proof of Theorem 1.11, which gives an alternative definition for $\mathcal{M}_{d,G}$, we will use a precise formula for this polynomial which we now develop. Every $H \in \mathcal{C}_{d,G}$, a d -covering of G , has exactly d edges covering any specific edge in G , and, likewise, d vertices covering every vertex of G . Thus, one can think of $\mathcal{M}_{d,G}$ as a generating function of multi-matchings in G : each edge in G can be picked multiple times so that each vertex is covered by at most d edges. We think of such a multi-matching as a function $m : E^+(G) \rightarrow \mathbb{Z}_{\geq 0}$. The weight associated to every multi-matching m is equal to the average number of ordinary matchings projecting to m in a random d -covering of G . Namely, the weight is the average number of matchings in $H \in \mathcal{C}_{d,G}$ with exactly $m(e)$ edges projecting to e , for every $e \in E^+(G)$.

To write an explicit formula, we extend m to all $E(G)$ by $m(-e) = m(e)$. We also denote by $e_{v,1}, \dots, e_{v,\deg(v)}$ the edges in $E(G)$ emanating from a vertex $v \in V(G)$ (in an arbitrary order, loops at v appearing twice, of course), and by $m(v)$ the number of edges covering the vertex $v \in V(G)$. Namely,

$$m(v) = \sum_{i=1}^{\deg(v)} m(e_{v,i}).$$

Finally, we denote by $|m|$ the total number of edges in m (with multiplicity), so $|m| = \sum_{e \in E^+(G)} m(e)$.

Definition 2.6. A **d-multi-matching** of a graph G is a function $m : E(G) \rightarrow \mathbb{Z}_{\geq 0}$ with $m(-e) = m(e)$ for every $e \in E(G)$ and $m(v) \leq d$ for every $v \in V(G)$. We denote the set of d -multi-matchings of G by $\text{MultiMatchings}_d(G)$.

MultiMatch

Proposition 2.7. *Let m be a multi-matching of G . Denote¹⁰*

$$W(m) = \frac{\prod_{v \in V(G)} \binom{d}{m(e_{v,1}), \dots, m(e_{v, \deg(v)})}}{\prod_{e \in E^+(G)} \binom{d}{m(e)}}. \quad (2.1)$$

Then,

$$\mathcal{M}_{d,G}(x) = \sum_{m \in \text{MultiMatchings}_d(G)} (-1)^{|m|} \cdot W(m) \cdot x^{nd-2|m|}. \quad (2.2)$$

Proof. Every matching of a d -covering $H \in \mathcal{C}_{d,G}$ projects to a unique multi-matching m of G covering every vertex of G at most d times. Thus, it is enough to show that $W(m)$ is exactly the average number of ordinary matchings projecting to m in a random $H \in \mathcal{C}_{d,G}$. Every such matching in H contains exactly $m(e)$ edges in the fiber above every $e \in E(G)$. Assume we know, for each $e \in E(G)$, which *vertices* in H are covered by the $m(e)$ edges above it. So there are $m(e)$ specific vertices in the fiber above $h(e)$, and $m(e)$ specific vertices in the fiber above $t(e)$. The probability that a random permutation in S_d matches specific $m(e)$ elements in $\{1, \dots, d\}$ to specific $m(e)$ elements in $\{1, \dots, d\}$ is

$$\frac{m(e)!(d - m(e))!}{d!} = \binom{d}{m(e)}^{-1}.$$

Thus, the denominator of $W(m)$ is equal to the probability that a random d -covering has a matching which projects to m and agrees with the particular choice of vertices. We are done as the numerator is exactly the number of possible choices of vertices. (Recall that since we deal with ordinary matching in H , every vertex is covered by at most one edge, so the set of vertices in the fiber above $v \in V(G)$ which are matched by the pre-image of $e_{v,i}$ is disjoint from those covered by the pre-image of $e_{v,j}$ whenever $i \neq j$.) Finally, we remark that the formula and proof remain valid also for graphs with multiple edges or loops. \square

The proof of Theorem 1.11 in Section 3 will consist of showing that $\mathbb{E}_\gamma[\phi_\gamma(x)]$ is equal to the expression in (2.2).

To summarize, here is what this paper shows about the generalized d -matching polynomial $\mathcal{M}_{d,G}$ of the graph G (see Section 2.4 for more details):


- It can be defined by any of the following:
 1. $\mathbb{E}_{H \in \mathcal{C}_{d,G}}[\mathcal{M}_{1,H}]$ - the average matching polynomial of a random d -covering of G
 2. $\mathbb{E}_{H \in \mathcal{C}_{d+1,G}} \left[\frac{\phi(A_H)}{\phi(A_G)} \right] = \mathbb{E}_{H \in \mathcal{C}_{d+1,G}} \left[\prod_{\mu \in \text{newSpec}(H)} (x - \mu) \right]$ - the average “new part” of the characteristic polynomial of a random $(d+1)$ -covering H of G
 3. $\mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma,G}}[\phi_\gamma, \pi]$ - the average characteristic polynomial of a random (Γ, π) -covering of G whenever (Γ, π) satisfies $(\mathcal{P}1)$ and π is d -dimensional
 4. $\sum_{m \in \text{MultiMatchings}_d(G)} (-1)^{|m|} \cdot W(m) \cdot x^{nd-2|m|}$, with $W(m)$ defined as in (2.1).
- If G has no loops, then $\mathcal{M}_{d,G}$ is real-rooted with all its roots in the Ramanujan interval.

¹⁰We use the notation $\binom{b}{a_1, a_2, \dots, a_k}$ to denote the multinomial coefficient $\frac{b!}{a_1! \dots a_k! (b - \sum a_i)!}$.

2.3 Group Labelings of Graphs

The model $\mathcal{C}_{d,G}$ we use for a random d -covering of a graph G is based on a uniformly random labeling $\gamma : E(G) \rightarrow S_d$. This is generalized in Definition 1.7 to $\mathcal{C}_{\Gamma,G}$, a probability-space of random Γ -labelings of the graph G . There are natural equivalent ways to obtain the same distribution on (isomorphism) types of d -coverings or Γ -labelings. Although the following will not be used in the rest of the paper, we choose to state it here, albeit loosely, for the sake of completeness.

Two d -covering H_1 and H_2 of G are isomorphic if there is a graph isomorphism between them which respects the covering maps. A similar equivalence relation can be given for Γ -labelings. This is the equivalence relation generated, for example, by the equivalence of the following two labelings of the edges incident to some vertex:



(here e_Γ is the identity element of Γ). For example, if the Γ -labelings γ_1 and γ_2 of G are isomorphic, then $\text{spec}(A_{\gamma_1,\pi}) = \text{spec}(A_{\gamma_2,\pi})$ for any finite dimensional representation π of Γ .

Claim 2.8. *Let G be a finite connected graph and Γ a finite/compact group. Let T be a spanning tree of G . The following three probability models yield the same distribution on isomorphism types of Γ -labelings of G :*

1. $\mathcal{C}_{\Gamma,G}$: uniform (Haar) distribution on labelings $\gamma : E^+(G) \rightarrow \Gamma$
2. uniform (Haar) distribution on homomorphisms $\pi_1(G) \rightarrow \Gamma$
3. an arbitrary fixed Γ -labeling of $E^+(T)$ (e.g. with the identity element of Γ) and a uniform (Haar) distribution on labelings of the remaining edges $E^+(G) \setminus E(T)$.

2.4 Group Representations

Let Γ be a group. A (complex, finite-dimensional) representation¹¹ of Γ is any group homomorphism $\pi : \Gamma \rightarrow \text{GL}_d(\mathbb{C})$ for some $d \in \mathbb{Z}_{\geq 1}$; if Γ is a topological group, we also demand π to be continuous. We then say π is a d -dimensional representation. The representation is called *faithful* if π is injective. Two d -dimensional representations π_1 and π_2 are *isomorphic* if they are conjugate to each other in the following sense: there is some $B \in \text{GL}_d(\mathbb{C})$ such that $\pi_2(g) = B^{-1}\pi_1(g)B$ for every $g \in \Gamma$. The *trivial* representation is the constant function $\text{triv} : \Gamma \rightarrow \text{GL}_1(\mathbb{C}) \cong \mathbb{C}^*$ mapping all elements to 1. The direct sum of two representations π_1 and π_2 of dimensions d_1 and d_2 , respectively, is a $(d_1 + d_2)$ -dimensional representation

¹¹A standard reference for the subject of group representations is [FH91].

$\pi_1 \oplus \pi_2 : \Gamma \rightarrow \text{GL}_{d_1+d_2}(\mathbb{C})$ where $(\pi_1 \oplus \pi_2)(g)$ is a block-diagonal matrix, with a $d_1 \times d_1$ block of $\pi_1(g)$ and a $d_2 \times d_2$ block of $\pi_2(g)$. A representation π is called *irreducible* if it is not isomorphic to the direct sum of two representations¹². Otherwise, it is called *reducible*.

The representation π is called *unitary* if its image in $\text{GL}_d(\mathbb{C})$ is conjugate to a subgroup of the unitary group¹³ $U(d) = \{A \in \text{GL}_d(\mathbb{C}) \mid A^{-1} = A^*\}$. In other words, it is isomorphic to a representation $\Gamma \rightarrow U(d)$. All representations of finite groups are unitary (e.g., conjugate π by $B = \left(\frac{1}{|\Gamma|} \sum_{g \in \Gamma} \pi(g)^* \pi(g)\right)^{1/2}$ to obtain a unitary image).

Claim 2.9. *Let π be a unitary representation of Γ and $A_{\gamma,\pi}$ a (Γ, π) -covering of some graph G . Then the spectrum of $A_{\gamma,\pi}$ is real.*

Proof. It is easy to see that $\text{spec}(A_{\gamma,\pi}) = \text{spec}(A_{\gamma,\pi'})$ whenever π and π' are isomorphic. Thus, assume without loss of generality that $\pi(\Gamma) \subseteq U(d)$. Then, by definition, $A_{\gamma,\pi}$ is Hermitian, and the statement follows. \square

The d -dimensional representation π of S_d mapping every $\sigma \in S_d$ to the corresponding permutation matrix is reducible: the 1-dimensional subspace of constant vectors $\langle \mathbf{1} \rangle \leq \mathbb{C}^d$ is invariant under this representation. The action of this representation on the orthogonal complement $\langle \mathbf{1} \rangle^\perp$ is a $(d-1)$ -dimensional irreducible representation of S_d called the **standard** representation and denoted **std**. The action on $\langle \mathbf{1} \rangle$ is isomorphic to the trivial representation. Thus, $\pi \cong \text{std} \oplus \text{triv}$.

Claim 2.10. *If γ is an S_d -labeling of G , then the new spectrum of the d -covering of G associated to γ is equal to the spectrum of $A_{\gamma,\text{std}}$.*

In particular, every (one-sided) Ramanujan d -covering of G corresponds to a unique (one-sided, respectively) Ramanujan (S_d, std) -covering of G .

Proof. For any Γ -labeling γ of the graph G and any two representations π_1 and π_2 , it is clear that $\text{spec}(A_{\gamma,\pi_1 \oplus \pi_2})$ is the disjoint union (as multisets) of $\text{spec}(A_{\gamma,\pi_1})$ and $\text{spec}(A_{\gamma,\pi_2})$. The claim follows as $A_{\gamma,\text{triv}} = A_G$ for any Γ -labeling γ . \square

In this language, Theorem 1.2 says that every graph G has a one-sided Ramanujan (S_d, std) -covering. This theorem will follow from Theorem 1.10 if we show that the pair (S_d, std) satisfies both $(\mathcal{P}1)$ and $(\mathcal{P}2)$. Before showing this, let us recall what exterior powers of representations are.

Let $V = \mathbb{C}^d$. For a d -dimensional representation π of Γ , its r -th exterior power, denoted $\bigwedge^r \pi$, is a $\binom{d}{r}$ -dimensional representation depicting the action of Γ on $\bigwedge^r V$, the r -th exterior power of V . To define $\bigwedge^r V$, consider the tensor power $\bigotimes^r V$ and quotient it by the subspace spanned by $\{v_1 \otimes v_2 \otimes \dots \otimes v_r \mid v_i = v_j \text{ for some } i \neq j\}$. The representative of $v_1 \otimes \dots \otimes v_r$ is denoted $v_1 \wedge \dots \wedge v_r$ and we have $v_{\sigma(1)} \wedge v_{\sigma(2)} \wedge \dots \wedge v_{\sigma(r)} = \text{sgn}(\sigma) \cdot v_1 \wedge v_2 \wedge \dots \wedge v_r$ for any

¹²Equivalently, π is irreducible if it has no non-trivial invariant subspace, namely, no $\{0\} \neq W \leq \mathbb{C}^d$ with $\pi(g)(W) \subseteq W$ for every $g \in \Gamma$.

¹³We use A^* to denote the conjugate-transpose of the matrix A .

permutation $\sigma \in S_r$. The representation $\bigwedge^r \pi$ is defined by the action of Γ on $\bigwedge^r V$, given by

$$g \cdot (v_1 \wedge \dots \wedge v_r) \stackrel{\text{def}}{=} (g.v_1) \wedge \dots \wedge (g.v_r).$$

Fact 2.11. *For every $d \in \mathbb{Z}_{\geq 1}$, the pair (S_{d+1}, std) of the symmetric group S_{d+1} with its standard, d -dimensional representation std satisfies both $(\mathcal{P}1)$ and $(\mathcal{P}2)$.*

Proof. That the exterior powers

$$\bigwedge^0 \text{std} = \text{triv}, \bigwedge^1 \text{std} = \text{std}, \bigwedge^2 \text{std}, \dots, \bigwedge^d \text{std} = \text{sign}$$

of std are all irreducible and non-isomorphic to each other is a classical fact: see, e.g., [FH91, Exercise 4.6]. In fact, $\bigwedge^r \text{std}$ is the irreducible representation corresponding to the Young diagram with $r+1$ rows $(d+1-r, 1, 1, \dots, 1)$. Hence (S_{d+1}, std) satisfies $(\mathcal{P}1)$.

The symmetric group S_{d+1} is generated by transpositions (permutations with $d-1$ fixed points and a single 2-cycle). The image of a transposition under $\pi \cong \text{triv} \oplus \text{std}$ is a pseudo-reflection (with spectrum $\{-1, 1, 1, \dots, 1\}$). Because the spectrum of $\text{triv}(\sigma)$ is $\{1\}$ for any $\sigma \in S_{d+1}$, we get that $\text{spec}(\text{std}(\sigma)) = \{-1, 1, \dots, 1\}$ (with $d-1$ ones) whenever σ is a transposition, namely, $\text{std}(\sigma)$ is a pseudo-reflection. Thus (S_{d+1}, std) satisfies $(\mathcal{P}2)$. \square

Fact 2.11 shows, then, why Theorem 1.2 follows from Theorem 1.10. It also shows why Theorem 1.5 follows from Theorem 1.11, Theorem 1.12 and Corollary 2.5.

In Section 3 below, we prove Theorem 1.11 and show that whenever the pair (Γ, π) satisfies $(\mathcal{P}1)$, the polynomial $\mathbb{E}_\gamma[\phi_{\gamma, \pi}]$ is equal to $\mathcal{M}_{d, G}$. The crux of this proof is a calculation of $\mathbb{E}_\gamma[\phi_{\gamma, \pi}] = \mathbb{E}_\gamma[\det(xI - A_{\gamma, \pi})]$ by minors of the $d \times d$ blocks, noticing that the determinant of an r -minor of $\pi(g)$ corresponds to an entry (matrix coefficient) of $(\bigwedge^r \pi)(g)$, and using the Peter-Weyl Theorem (Theorem 3.3 below) for matrix coefficients.

2.5 Interlacing Families of Polynomials

Following the technique introduced by Marcus et al in their seminal series of papers (e.g. [MSS15a, MSS15b]), we prove Theorem 1.12 by introducing a family of interlacing polynomials.

Definition 2.12. The polynomials $f, g \in \mathbb{R}[x]$ are *interlacing* if they have the same degree (say, n), their leading coefficient has the same sign, they are real-rooted, and their roots $\alpha_n \leq \dots \leq \alpha_1$ and $\beta_n \leq \dots \leq \beta_1$ satisfy

$$\{\alpha_n, \beta_n\} \leq \{\alpha_{n-1}, \beta_{n-1}\} \leq \dots \leq \{\alpha_2, \beta_2\} \leq \{\alpha_1, \beta_1\}$$

(i.e., $\alpha_{i+1} \leq \beta_i$ and $\beta_{i+1} \leq \alpha_i$ for every i).

This definition can be extended to any set of polynomials: the i -th root of any of them is bigger than (or equal to) the $(i+1)$ -st root of any other. An easy but important property of interlacing polynomials is that any weighted average of them is also real-rooted, with its i -th

root lying between the i -th roots of the polynomials (namely, the i -th root of a convex sum of interlacing polynomials is some convex sum of their i -th roots). Moreover, these weighted average polynomials supply an alternative criterion for interlacing:

Claim 2.13. *[e.g. [MSS15b, Lemma 3.5]] The polynomials $f_1, \dots, f_r \in \mathbb{R}[x]$ are interlacing if and only the average $\lambda_1 f_1 + \dots + \lambda_r f_r$ is real-rooted for every $\lambda_1, \dots, \lambda_r$ with $\lambda_i \geq 0$ and $\sum \lambda_i = 1$.*

Now let Γ be a finite group and $\pi : \Gamma \rightarrow \mathrm{GL}_d(\mathbb{C})$ a representation. In Section 4 below, inspired by [MSS15c, Section 3], we show there are different distributions on (Γ, π) -coverings, so that the corresponding expected characteristic polynomials are interlacing. The main technical result (Theorem 4.2 below) is a generalization of the fact that if $A, B \in \mathrm{GL}_d(\mathbb{C})$ with A Hermitian and B diagonalizable with $\mathrm{rank}(B - I_d) = 1$, then $\phi(A)$ and $\phi(BAB^{-1})$ interlace.

More concretely, we generate a random (Γ, π) -covering of a loopless graph G by generating a random Γ -labeling as follows. If $\gamma_1, \gamma_2 \in E^+(G) \rightarrow \Gamma$ are two Γ -labelings of G , we define their *product* $\gamma_1 \gamma_2$ as the point-wise product, namely $(\gamma_1 \gamma_2)(e) \stackrel{\mathrm{def}}{=} \gamma_1(e) \cdot \gamma_2(e)$ for every $e \in E^+(G)$. Of course, for the reversed orientation we have $(\gamma_1 \gamma_2)(-e) = ((\gamma_1 \gamma_2)(e))^{-1} = \gamma_2(-e) \gamma_1(-e)$.

Assume that X_1, X_2, \dots, X_r are independent random variables, each taking values in the space of Γ -labelings of G , such that any two possible values of X_i differ by a pseudo-reflection on one of the edges and are identical on all other edges. Namely, if γ_1, γ_2 are two Γ -labelings of G in the support of X_i , then the labeling $\gamma_1 \gamma_2^{-1}$ is the identity on every edge except for on one edge e where $\pi(\gamma_1 \gamma_2^{-1}(e))$ is a pseudo-reflection. As we show in Proposition 4.4 below, in this case, the random Γ -labeling of G defined by $Y = X_1 \cdot \dots \cdot X_r$ satisfies that $\mathbb{E}_Y[\phi_{Y, \pi}]$ is real-rooted (recall that $\phi_{\gamma, \pi}$ denotes the characteristic polynomial of $A_{\gamma, \pi}$ - see (1.2)).

Now, assume the possible values of X_i are η_1, \dots, η_t , and for $1 \leq j \leq t$ let $Y_j = X_1 \dots X_{i-1} \eta_j X_{i+1} \dots X_r$ be the random Γ -labeling conditioned on $X_i = \eta_j$. The polynomial $\mathbb{E}_Y[\phi_{Y, \pi}]$ is still real-rooted, by Proposition 4.4, even if we tweak X_i by fixing arbitrary probabilities on η_1, \dots, η_t . Namely, for any probability vector (p_1, \dots, p_t) , the polynomial

$$p_1 \cdot \mathbb{E}_{Y_1}[\phi_{Y_1, \pi}] + p_2 \cdot \mathbb{E}_{Y_2}[\phi_{Y_2, \pi}] + \dots + p_t \cdot \mathbb{E}_{Y_t}[\phi_{Y_t, \pi}]$$

is real-rooted. By Claim 2.13 it follows that the t polynomials $\mathbb{E}_{Y_j}[\phi_{Y_j, \pi}]$, $1 \leq j \leq t$, are interlacing. In particular, at least one of them has its largest root bounded from above by the largest root of their weighted average $\mathbb{E}_Y[\phi_{Y, \pi}]$.

We obtain that for every random Γ -covering Y as above, there is an actual Γ -labeling $\gamma = \gamma_1 \cdot \dots \cdot \gamma_r$ (with γ_i in the support of X_i), so that the maximal root of $\phi_{\gamma, \pi}$ is at most the maximal root of $\mathbb{E}_Y[\phi_{Y, \pi}]$. To see this, use the argument in the previous paragraph to choose γ_1 in the support of X_1 so that the maximal root of $\mathbb{E}_{\gamma_1 X_2 X_3 \dots X_r}[\phi_{\gamma_1 X_2 X_3 \dots X_r, \pi}]$ is at most the maximal root of $\mathbb{E}_{Y=X_1 X_2 \dots X_r}[\phi_{Y, \pi}]$. Then, use the same argument (think of γ_1 as a “Dirac” random labeling), to find $\gamma_2 \in \mathrm{Support}(X_2)$ so that the maximal root of $\mathbb{E}_{\gamma_1 \gamma_2 X_3 \dots X_r}[\phi_{\gamma_1 \gamma_2 X_3 \dots X_r, \pi}]$ is at most the maximal root of $\mathbb{E}_{\gamma_1 X_2 X_3 \dots X_r}[\phi_{\gamma_1 X_2 X_3 \dots X_r, \pi}]$. Continue




Figure 2.1: A tree of interlacing polynomials: the expected characteristic polynomials of the random (Γ, π) -coverings associated to the random labelings in the figure are interlacing.

the same way to end up with a specific Γ -labeling whose largest root is at most the maximal root of $\mathbb{E}_Y [\phi_{Y,\pi}]$. (See Figure 2.1.)

Finally, we show (Section 4.3) that if (Γ, π) satisfies $(\mathcal{P}2)$, namely, if Γ is finite and $\pi(\Gamma)$ is generated by pseudo-reflections, then the uniform distribution $\mathcal{C}_{\Gamma,G}$ can be approximated by distributions Y as above. This will show that the two properties that are satisfied by any such Y , namely, the real-rootedness of the average characteristic polynomial and the existence of an actual Γ -labeling with smallest largest root, are also satisfied by the distribution $\mathcal{C}_{\Gamma,G}$. This is exactly the content of Theorem 1.12.

3 Property $(\mathcal{P}1)$ and the Proof of Theorem 1.11

Recall that G is an undirected oriented graph with n vertices. In this section we assume the pair (Γ, π) satisfies $(\mathcal{P}1)$, namely that Γ is finite, or more generally, compact, and π is a d -dimensional representation such that its exterior powers $\bigwedge^0 \pi, \dots, \bigwedge^d \pi$ are irreducible and non-isomorphic. We need to show that $\mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma,G}} [\phi_{\gamma,\pi}] = \mathcal{M}_{d,G}$.

For every Γ -labeling γ of G , we represent the matrix $A_{\gamma,\pi} \in M_{nd}(\mathbb{C})$ as a sum of $|E(G)|$ matrices as follows. For every $e \in E(G)$, let $\mathbf{A}_{\gamma,\pi}(e) \in M_{nd}(\mathbb{C})$ be the $nd \times nd$ matrix $A_{\gamma,\pi}(e)$ composed of n^2 blocks of size $d \times d$ each. All blocks are zero blocks except for the one corresponding to e , the block $(h(e), t(e))$, in which we put $\pi(\gamma(e))$. Clearly,

$$A_{\gamma,\pi} = \sum_{e \in E(G)} A_{\gamma,\pi}(e).$$

In order to analyze the expected characteristic polynomial of this sum of matrices, we begin with a technical lemma, giving the determinant of a sum of matrices as a formula in terms of the determinants of their minors. We then use this lemma when we complete the proof of Theorem 1.11 in Section 3.3.

3.1 Determinant of Sum of Matrices

Let $A_1, \dots, A_m \in M_d(\mathbb{C})$ be $d \times d$ matrices. The determinant $|A_1 + \dots + A_m|$ can be thought of as a double sum. First, sum over all permutations $\sigma \in S_d$ the term $\text{sgn}(\sigma) \prod_{i=1}^d (A_1 + \dots + A_m)_{i, \sigma(i)}$. Then, for each term and each¹⁴ $i \in [d]$, choose $s_\sigma(i) \in [m]$ which marks which of the m summands is taken in the entry $i, \sigma(i)$. Namely,

$$|A_1 + \dots + A_m| = \sum_{\sigma \in S_d} \text{sgn}(\sigma) \sum_{s_\sigma: [d] \rightarrow [m]} \prod_{i=1}^d [A_{s_\sigma(i)}]_{i, \sigma(i)}. \quad (3.1)$$

The idea of Lemma 3.1 below is to group the terms in this double sum differently: first, for every $j \in [m]$, choose from which rows R_j and from which columns C_j the entry is taken from A_j . Then, go over all permutations σ that respect these constraints, namely the permutations for which $\sigma(R_j) = C_j$. To this aim, we define:

$T(m, d)$

$$T(m, d) = \left\{ \left(\dot{R}, \dot{C} \right) \left| \begin{array}{l} \dot{R} = (R_1, \dots, R_m), \dot{C} = (C_1, \dots, C_m) \text{ are partitions of } [d] \text{ into } m \text{ parts} \\ \text{such that } |R_\ell| = |C_\ell| \text{ for all } \ell \end{array} \right. \right\},$$

and the corresponding permutations:

$\text{Sym}(\dot{R}, \dot{C})$

$$\text{Sym}(\dot{R}, \dot{C}) \stackrel{\text{def}}{=} \{ \sigma \in S_d \mid \sigma(R_\ell) = C_\ell \text{ for all } \ell \}.$$

Finally, for each such pair of partitions $(\dot{R}, \dot{C}) \in T(m, d)$, we need a “relative sign”, denoted $\text{sgn}(\dot{R}, \dot{C})$, which will enable us to calculate the sign of every $\sigma \in \text{Sym}(\dot{R}, \dot{C})$ based solely on the signs of the permutation-matrix σ restricted to the minors (R_ℓ, C_ℓ) . This is the sign of the permutation-matrix obtained by assigning the identity matrix $I_{|R_\ell|}$ to the (R_ℓ, C_ℓ) minor for every ℓ . For example, if $\dot{R} = (\{1, 3, 5\}, \{2, 4\})$ and $\dot{C} = (\{3, 4, 5\}, \{1, 2\})$, then

$\text{Sgn}(\dot{R}, \dot{C})$

$$\text{sgn}(\dot{R}, \dot{C}) = \text{sgn} \begin{pmatrix} 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Lemma 3.1. *If $A_1, \dots, A_m \in M_d(\mathbb{C})$ are $d \times d$ matrices, then*

$$|A_1 + \dots + A_m| = \sum_{(\dot{R}, \dot{C}) \in T(m, d)} \text{sgn}(\dot{R}, \dot{C}) \prod_{\ell=1}^m |A_\ell|_{R_\ell, C_\ell},$$

where for $R, C \subseteq [d]$ with $|R| = |C|$,

$$|A|_{R, C} = \begin{cases} \det((a_{i,j})_{i \in R, j \in C}) & \text{if } |R| = |C| \geq 1 \\ 1 & \text{if } R = C = \emptyset \end{cases}$$

marks the determinant of the (R, C) -minor of A .

¹⁴We use the standard notation of $[d]$ for the set $\{1, \dots, d\}$.

Proof. For every $\sigma \in S_d$ and $s_\sigma : [d] \rightarrow [m]$ as in (3.1), there is a unique pair of partitions $(\dot{R}, \dot{C}) \in T(m, d)$ which respects σ and s_σ , namely for which $\sigma \in \text{Sym}(\dot{R}, \dot{C})$ and $s_\sigma^{-1}(\ell) = R_\ell$ for each ℓ . This is precisely the pair (\dot{R}, \dot{C}) defined by $R_\ell = s_\sigma^{-1}(\ell)$ and $C_\ell = \sigma(R_\ell)$. Therefore,

$$|A_1 + \dots + A_m| = \sum_{(\dot{R}, \dot{C}) \in T(m, d)} \sum_{\sigma \in \text{Sym}(\dot{R}, \dot{C})} \text{sgn}(\sigma) \sum_{\substack{s_\sigma : [d] \rightarrow [m] \\ s_\sigma^{-1}(\ell) = R_\ell \forall \ell}} \prod_{i=1}^d [A_{s_\sigma(i)}]_{i, \sigma(i)}.$$

Now, to determine a permutation $\sigma \in \text{Sym}(\dot{R}, \dot{C})$ is equivalent to determine the permutation σ_ℓ induced by σ on each of the minors (R_ℓ, C_ℓ) . Thus, $\text{Sym}(\dot{R}, \dot{C}) \cong S_{|R_1|} \times \dots \times S_{|R_m|}$ (as sets) by $\sigma \mapsto (\sigma_1, \dots, \sigma_m)$. It is easy to see that the signs of these permutations are related by

$$\text{sgn}(\sigma) = \text{sgn}(\dot{R}, \dot{C}) \cdot \text{sgn}(\sigma_1) \cdot \dots \cdot \text{sgn}(\sigma_m).$$

Thus, if $R_\ell(i)$ is the i -th element in R_ℓ , we get

$$\begin{aligned} |A_1 + \dots + A_m| &= \sum_{(\dot{R}, \dot{C}) \in T(m, d)} \text{sgn}(\dot{R}, \dot{C}) \prod_{\ell=1}^m \sum_{\sigma_\ell \in S_{|R_\ell|}} \text{sgn}(\sigma_\ell) \prod_{i=1}^{|R_\ell|} [A_\ell]_{R_\ell(i), C_\ell(\sigma_\ell(i))} \\ &= \sum_{(\dot{R}, \dot{C}) \in T(m, d)} \text{sgn}(\dot{R}, \dot{C}) \prod_{\ell=1}^m |A_\ell|_{R_\ell, C_\ell}. \end{aligned}$$

□

3.2 Matrix Coefficients

Recall that if π is a d -dimensional representation, then $\bigwedge^r \pi$ is $\binom{d}{r}$ -dimensional, and if $\{v_1, \dots, v_d\}$ is a basis for \mathbb{C}^d , then $\{v_{i_1} \wedge \dots \wedge v_{i_r} \mid 1 \leq i_1 < i_2 < \dots < i_r \leq d\}$ is a basis for $\bigwedge^r(\mathbb{C}^d)$ (see Section 2.4). The following standard claim and classical theorem explain the role in Theorem 1.11 of the conditions on $\bigwedge^r \pi$ as defined in property (P1):

Claim 3.2. (e.g. [KRY09, Theorem 6.6.3]) *If the matrices $\pi(g)$ are given in terms of the basis $V = \{v_1, \dots, v_d\}$ and $(\bigwedge^r \pi)(g)$ in terms of the basis $\{v_{i_1} \wedge \dots \wedge v_{i_r} \mid 1 \leq i_1 < i_2 < \dots < i_r \leq d\}$, then the entry (matrix coefficient) of $(\bigwedge^r \pi)(g)$ in row (i_1, \dots, i_r) and column (j_1, \dots, j_r) is given by the minor-determinant $|\pi(g)|_{\{i_1, \dots, i_r\}, \{j_1, \dots, j_r\}}$.*

Theorem 3.3. [Peter-Weyl, see [Bum04, Chapter 2]] *The matrix coefficients of the irreducible representations of a compact group Γ are an orthogonal basis of $L^2(\Gamma)$. In particular, if $\pi_1 : \Gamma \rightarrow U(d_1)$ and $\pi_2 : \Gamma \rightarrow U(d_2)$ are irreducible non-isomorphic unitary representations of Γ , then*

$$\mathbb{E}_{g \in \Gamma} \left[\pi_1(g)_{i_1, j_1} \cdot \overline{\pi_2(g)_{i_2, j_2}} \right] = 0$$

for every $i_1, j_1 \in [d_1]$ and $i_2, j_2 \in [d_2]$, the expectation taken according to the Haar measure of Γ . Moreover, if $\pi : \Gamma \rightarrow U(d)$ is an irreducible representation, then

$$\mathbb{E}_{g \in \Gamma} \left[\pi(g)_{i_1, j_1} \cdot \overline{\pi(g)_{i_2, j_2}} \right] = \begin{cases} \frac{1}{d} & (i_1, j_1) = (i_2, j_2) \\ 0 & \text{otherwise} \end{cases}.$$

If Γ is finite, the Haar measure is simply the uniform measure, so the expectations in the theorem are given by

$$\frac{1}{|\Gamma|} \sum_{g \in \Gamma} \pi_1(g)_{i_1, j_1} \cdot \overline{\pi_2(g)_{i_2, j_2}}.$$

We now have all the tools needed to prove Theorem 1.11.

3.3 Proof of Theorem 1.11

As π is unitary, we assume without loss of generality that $\pi : \Gamma \rightarrow U(d)$ maps the elements of Γ to unitary matrices, so that for every $e \in E(G)$, $A_{\gamma, \pi}(-e)$ is the conjugate-transpose matrix $A_{\gamma, \pi}(e)^*$.

We analyze the expected characteristic polynomial

$$\mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma, G}} [\phi_{\gamma, \pi}] = \mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma, G}} [\det(xI - A_{\gamma, \pi})] = \mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma, G}} \left[\det \left(xI - \sum_{e \in E(G)} A_{\gamma, \pi}(e) \right) \right]. \quad (3.2)$$

Our goal is to show it is equal to the formula given for $\mathcal{M}_{d, G}$ in Proposition 2.7. We use Lemma 3.1 to rewrite the determinant in the right hand side of (3.2). We now let

$$T = T(|E(G)| + 1, nd) = \left\{ (\dot{R}, \dot{C}) \left| \begin{array}{l} \dot{R} \text{ and } \dot{C} \text{ are partitions of } [nd] \text{ to } |E(G)| + 1 \text{ parts} \\ \text{indexed by } \{x\} \cup E(G), \\ \text{with } R_x = C_x \text{ and } |R_e| = |C_e| \text{ for all } e \in E(G) \end{array} \right. \right\}.$$

By Lemma 3.1,

$$\phi_{\gamma, \pi} = \sum_{(\dot{R}, \dot{C}) \in T} \text{sgn}(\dot{R}, \dot{C}) \cdot x^{|R_x|} \prod_{e \in E(G)} (-1)^{|R_e|} |A_{\gamma, \pi}(e)|_{R_e, C_e},$$

and taking expected values gives

$$\begin{aligned} \mathbb{E}_{\gamma} [\phi_{\gamma, \pi}] &= \sum_{(\dot{R}, \dot{C}) \in T} \text{sgn}(\dot{R}, \dot{C}) \cdot x^{|R_x|} \mathbb{E}_{\gamma} \left[\prod_{e \in E(G)} (-1)^{|R_e|} |A_{\gamma, \pi}(e)|_{R_e, C_e} \right] \\ &= \sum_{(\dot{R}, \dot{C}) \in T} \text{sgn}(\dot{R}, \dot{C}) \cdot x^{|R_x|} (-1)^{nd - |R_x|} \prod_{e \in E^+(G)} \mathbb{E}_{\gamma} [|A_{\gamma, \pi}(e)|_{R_e, C_e} \cdot |A_{\gamma, \pi}(-e)|_{R_{-e}, C_{-e}}], \end{aligned} \quad (3.3)$$

since the $A_{\gamma,\pi}(e)$ are independent except for the pairs $A_{\gamma,\pi}(e)$ and $A_{\gamma,\pi}(-e)$.

Since $A_{\gamma,\pi}(-e) = A_{\gamma,\pi}(e)^*$, the term inside the expectation in the right hand side of (3.3) is equal to

$$\mathbb{E}_\gamma [|A_{\gamma,\pi}(e)|_{R_e, C_e} \cdot |A_{\gamma,\pi}(e)^*|_{R_{-e}, C_{-e}}] = \mathbb{E}_\gamma [|A_{\gamma,\pi}(e)|_{R_e, C_e} \cdot \overline{|A_{\gamma,\pi}(e)|_{C_{-e}, R_{-e}}}].$$

Clearly, this term is zero, unless the minors we choose for e and $-e$ are inside the $d \times d$ blocks corresponding to e and $-e$, respectively. That is, if B_v denotes the set of d indices of rows and columns corresponding to the vertex $v \in V(G)$, then this term is zero unless $R_e, C_{-e} \subseteq B_{h(e)}$ and $C_e, R_{-e} \subseteq B_{t(e)}$. If this is the case, we can think of R_e, C_e, R_{-e}, C_{-e} as subsets of $[d]$, so Claim 3.2 yields this term is

$$\mathbb{E}_\gamma \left[\left(\left(\bigwedge^{|R_e|} \pi \right) (\gamma(e)) \right)_{R_e, C_e} \cdot \overline{\left(\left(\bigwedge^{|R_{-e}|} \pi \right) (\gamma(e)) \right)_{C_{-e}, R_{-e}}} \right],$$

where we identify an r -subset of $[d]$ with a basis element of $\bigwedge^r(\mathbb{C}^d)$ in the obvious way. Finally, by Peter-Weyl Theorem (Theorem 3.3) and our assumptions on the exterior powers $\bigwedge^r \pi$ for $0 \leq r \leq d$, this expectation is zero unless $|R_e| = |R_{-e}|$, $R_e = C_{-e}$, and $C_e = R_{-e}$. If all these equalities hold, the expectation is $\binom{d}{|R_e|}^{-1}$.

Define $T^{\text{sym}} \subseteq T$ to be the subset of T containing the partitions for which the expectation in (3.3) is not zero. Namely,

$$T^{\text{sym}} = \left\{ (\dot{R}, \dot{C}) \left| \begin{array}{l} \dot{R} \text{ and } \dot{C} \text{ are partitions of } [nd] \text{ to } |E(G)| + 1 \text{ parts} \\ \text{indexed by } \{x\} \cup E(G), \\ \text{with } R_x = C_x, \text{ and for all } e \in E^+(G) \\ |R_e| = |C_e|, C_{-e} = R_e, R_{-e} = C_e, R_e \subseteq B_{h(e)} \text{ and } R_{-e} \subseteq B_{t(e)} \end{array} \right. \right\}.$$

Our discussion shows that

$$\mathbb{E}_\gamma [\phi_{\gamma,\pi}] = \sum_{(\dot{R}, \dot{C}) \in T^{\text{sym}}} \text{sgn}(\dot{R}, \dot{C}) \cdot x^{|R_x|} (-1)^{nd - |R_x|} \prod_{e \in E^+(G)} \frac{1}{\binom{d}{|R_e|}}.$$

Now, notice that because $|R_{-e}| = |R_e|$, we get that $nd - |R_x| = \sum_{e \in E(G)} |R_e|$ is even, so $(-1)^{nd - |R_x|} = 1$ for every $(\dot{R}, \dot{C}) \in T^{\text{sym}}$. Because of the conditions $C_{-e} = R_e$ and $R_{-e} = C_e$ on the partitions in T^{sym} , the permutation matrix defining $\text{sgn}(\dot{R}, \dot{C})$ is symmetric. Thus, the corresponding permutation is an involution, with exactly $|R_x|$ fixed points and $\frac{nd - |R_x|}{2}$ 2-cycles¹⁵, so $\text{sgn}(\dot{R}, \dot{C}) = (-1)^{(nd - |R_x|)/2}$. Hence,

$$\mathbb{E}_\gamma [\phi_{\gamma,\pi}] = \sum_{(\dot{R}, \dot{C}) \in T^{\text{sym}}} (-1)^{(nd - |R_x|)/2} \cdot x^{|R_x|} \prod_{e \in E^+(G)} \frac{1}{\binom{d}{|R_e|}}.$$

¹⁵In particular, if $(\dot{R}, \dot{C}) \in T^{\text{sym}}$ then $R_e \cap C_e = \emptyset$, even for loops.

Recall the definition of a d -multi-matching given in Definition 2.6: this is a function $m : E(G) \rightarrow \mathbb{Z}_{\geq 0}$ with $m(-e) = m(e)$ such that $m(v) \leq d$ for every $v \in V(G)$, where $m(v)$ is the sum of m on all oriented edges originating from v . The map $\eta(\dot{R}, \dot{C})$ given by $e \mapsto |R_e|$ is a d -multi-matching for every $(\dot{R}, \dot{C}) \in T^{\text{sym}}$, since for every $v \in V(G)$,

$$\eta(\dot{R}, \dot{C})(v) = \sum_{e: h(e)=v} |R_e| \leq d$$

as \dot{R} is a partition and $R_e \subseteq B_{h(e)}$ if $h(e) = v$.

Finally, for every $(\dot{R}, \dot{C}) \in T^{\text{sym}}$, \dot{C} is completely determined by \dot{R} . Denote by $e_{v,1}, \dots, e_{v,\deg(v)}$ the oriented edges emanating from v . Then, for every d -multi-matching m , the number of partitions $(\dot{R}, \dot{C}) \in T^{\text{sym}}$ associated to m is exactly

$$\prod_{v \in V(G)} \binom{d}{m(e_{v,1}), \dots, m(e_{v,\deg(v)})}.$$

We obtain

$$\begin{aligned} \mathbb{E}_\gamma[\phi_{\gamma,\pi}] &= \sum_m \sum_{(\dot{R}, \dot{C}) \in T^{\text{sym}} : \eta(\dot{R}, \dot{C})=m} (-1)^{(nd-|R_x|)/2} \cdot x^{|R_x|} \prod_{e \in E^+(G)} \frac{1}{\binom{d}{|R_e|}} \\ &= \sum_m (-1)^{|m|} x^{nd-2|m|} \frac{\prod_{v \in V(G)} \binom{d}{m(e_{v,1}), \dots, m(e_{v,\deg(v)})}}{\prod_{e \in E^+(G)} \binom{d}{m(e)}}, \end{aligned}$$

where the summation is over all d -multi-matchings m of G , and $|m| = \sum_{e \in E^+(G)} m(e)$. This is precisely the formula for $\mathcal{M}_{d,G}$ from Proposition 2.7, so the proof of Theorem 1.11 is complete. \square

4 Property ($\mathcal{P}2$) and the Proof of Theorem 1.12

The main goal of this Section is to show that the expected characteristic polynomial of certain distributions of random (Γ, π) -coverings is real rooted, and in particular that this is true for the uniform distribution. We follow the outline depicted in Section 2.5. The main component of the proof is Theorem 4.2, showing that for certain distributions of Hermitian (self-adjoint) matrices, the expected characteristic polynomial is real-rooted. This theorem imitates and generalizes the argument of [MSS15c, Thm 3.3]. We repeat the argument in Section 4.1, because we need the more general statement, but we refer the interested reader to [MSS15c, Section 3] for some more elaborated concepts and notions. Theorem 4.2 is a

generalization of the fact that the characteristic polynomials¹⁶ $\phi(A)$ and $\phi(BAB^*)$ interlace whenever $A \in M_d(\mathbb{C})$ is Hermitian and $B \in U(d)$ satisfies $\text{rank}(B - I_d) = 1$.

4.1 Average Characteristic Polynomial of Sum of Random Matrices

Definition 4.1. We say that the random variable W taking values in $U(d)$ is **U(1)-like** if every two different possible values B_1 and B_2 satisfy $\text{rank}(B_1 B_2^{-1} - I_d) = 1$.

It is not hard to see that W is $U(1)$ -like if and only if it takes values in some $P\Lambda Q$ where $P, Q \in U(d)$ and $\Lambda \leq U(d)$ is the subgroup of diagonal matrices

$$\left\{ \begin{pmatrix} \lambda & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{pmatrix} \middle| |\lambda| = 1 \right\}.$$

Theorem 4.2. Let $m \in \mathbb{Z}_{\geq 1}$, let $\ell(1), \dots, \ell(m) \in \mathbb{Z}_{\geq 0}$, and let $\mathcal{W} = \{W_{i,j}\}_{1 \leq i \leq m, 1 \leq j \leq \ell(i)}$ be a set of independent $U(1)$ -like random variables taking values in $U(d)$. If $A_1, \dots, A_m \in M_d(\mathbb{C})$ are Hermitian matrices, then

$$P_{\mathcal{W}}(A_1, \dots, A_m) \stackrel{\text{def}}{=} \mathbb{E}_{\mathcal{W}} \left[\phi \left(W_{1,1} \dots W_{1,\ell(1)} A_1 W_{1,\ell(1)}^* \dots W_{1,1}^* + \dots + W_{m,1} \dots W_{m,\ell(m)} A_m W_{m,\ell(m)}^* \dots W_{m,1}^* \right) \right]$$

is real rooted.

Note that the characteristic polynomial of a Hermitian matrix is in $\mathbb{R}[x]$, and so $P_{\mathcal{W}}(A_1, \dots, A_m)$ is in $\mathbb{R}[x]$ since it is an average of such polynomials.

Lemma 4.3. In the notation of Theorem 4.2, assume that $P_{\mathcal{W}}(A_1, \dots, A_m)$ is real rooted whenever $A_1, \dots, A_m \in M_d(\mathbb{C})$ are Hermitian. Then, for every $v \in \mathbb{C}^d$, the roots $\alpha_d \leq \dots \leq \alpha_1$ of $P_{\mathcal{W}}(A_1, \dots, A_i + vv^*, \dots, A_m)$ and the roots $\beta_d \leq \dots \leq \beta_1$ of $P_{\mathcal{W}}(A_1, \dots, A_m)$ satisfy

$$\beta_d \leq \alpha_d \leq \beta_{d-1} \leq \alpha_{d-1} \leq \dots \leq \beta_1 \leq \alpha_1.$$

In other words, the polynomials $P_{\mathcal{W}}(A_1, \dots, A_i + vv^*, \dots, A_m)$ and $P_{\mathcal{W}}(A_1, \dots, A_m)$ interlace in a strong sense.

Proof. Denote $Q(A_1, \dots, A_m) = \left[\frac{\partial}{\partial s} P_{\mathcal{W}}(A_1, A_2, \dots, A_i + s vv^*, \dots, A_m) \right]_{s=0}$. Note that¹⁷

$$P_{\mathcal{W}}(A_1, \dots, A_i + \mu vv^*, \dots, A_m) = P_{\mathcal{W}}(A_1, \dots, A_m) + \mu \cdot Q(A_1, \dots, A_m) \quad (4.1)$$

¹⁶Recall that for any matrix A , we denote its characteristic polynomial by $\phi(A) = \det(xI - A)$.

¹⁷This property is called “Rank-1 linearity” in [MSS15c].

for every $\mu \in \mathbb{C}$. To see this, it is enough to show (4.1) in the case \mathcal{W} is constant (namely, $W_{i,j}$ is constant for every i, j), the general statement will then follow by linearity of expectation and of derivative. For a constant \mathcal{W} , we only need to show that for any Hermitian $A \in M_d(\mathbb{C})$, the characteristic polynomial $\phi(A + svv^*)$ is linear in s . By conjugating A and vv^* by some unitary matrix, we may assume $v^* = (\alpha, 0, 0, \dots, 0)$, and then the claim is clear by developing the determinant of $xI - (A + svv^*)$ by, say, the first row.

Note that $Q(A_1, \dots, A_m)$ it is a polynomial of degree $d - 1$ with a negative leading coefficient while $P_{\mathcal{W}}(A_1, \dots, A_m)$ is monic of degree d . For every $\mu \in \mathbb{R}_{\geq 0}$, $A_i + \mu vv^*$ is Hermitian, so by our assumption the left hand side of (4.1) is real-rooted. Namely, $P_{\mathcal{W}}(A_1, \dots, A_m) + \mu \cdot Q(A_1, \dots, A_m)$ is real rooted for every $\mu \in \mathbb{R}_{\geq 0}$. Equivalently, $(1 - \lambda) \cdot P_{\mathcal{W}}(A_1, \dots, A_m) + \lambda \cdot Q(A_1, \dots, A_m)$ is real rooted for every $\lambda \in [0, 1]$, and by continuity also for $\lambda = 1$. Recall the roots of $P_{\mathcal{W}}(A_1, \dots, A_m)$ are denoted $\beta_d \leq \dots \leq \beta_1$. A standard argument from the theory of interlacing polynomials (see, e.g. [Fis06]) shows that the real roots $\vartheta_{d-1} \leq \dots \leq \vartheta_1$ of $Q(A_1, \dots, A_m)$ satisfy

$$\beta_d \leq \vartheta_{d-1} \leq \beta_{d-1} \leq \dots \leq \beta_2 \leq \vartheta_1 \leq \beta_1.$$

Moreover, the i -th root of $(1 - \lambda) \cdot P_{\mathcal{W}}(A_1, \dots, A_m) + \lambda \cdot Q(A_1, \dots, A_m)$ moves continuously from β_i to ϑ_{i-1} as λ moves from 0 to 1 (the first root moves from β_1 to ∞). In particular, $\lambda = \frac{1}{2}$ corresponds (up to scaling) to $P_{\mathcal{W}}(A_1, \dots, A_i + vv^*, \dots, A_m)$, the roots $\alpha_d \leq \dots \leq \alpha_1$ satisfy

$$\beta_d \leq \alpha_d \leq \vartheta_{d-1} \leq \beta_{d-1} \leq \alpha_{d-1} \leq \vartheta_{d-1} \leq \dots \leq \beta_2 \leq \alpha_2 \leq \vartheta_1 \leq \beta_1 \leq \alpha_1,$$

and the lemma is proven. \square

Proof. [of Theorem 4.2] We prove by induction on the number of $W_{i,j}$ (namely, induction on $\ell(\mathcal{W}) = \ell(1) + \dots + \ell(m)$). The statement is clear for $\ell(\mathcal{W}) = 0$. Given \mathcal{W} with $\ell(\mathcal{W}) > 0$, assume without loss of generality that $\ell(1) > 0$, and denote by \mathcal{W}' the set of random variables $\mathcal{W} \setminus \{W_{1,\ell(1)}\}$. We assume, by the induction hypothesis, that $P_{\mathcal{W}'}(A_1, \dots, A_m)$ is real-rooted for every Hermitian matrices A_1, \dots, A_m .

For clarity we assume $W_{1,\ell(1)}$ takes only finitely many values, but the same argument works in the general case. Let $B_1, \dots, B_t \in U(d)$ be the possible values of $W_{1,\ell(1)}$, obtained with probabilities p_1, \dots, p_t . We need to show that

$$P_{\mathcal{W}}(A_1, \dots, A_m) = p_1 \cdot P_{\mathcal{W}'}(B_1 A_1 B_1^*, A_2, \dots, A_m) + \dots + p_t \cdot P_{\mathcal{W}'}(B_t A_1 B_t^*, A_2, \dots, A_m)$$

is real rooted for all Hermitian matrices $A_1, \dots, A_m \in M_d(\mathbb{C})$. By Claim 2.13, it is enough to show the t polynomials $P_{\mathcal{W}'}(B_j A_1 B_j^*, A_2, \dots, A_m)$ ($1 \leq j \leq t$) are interlacing. By definition, this is equivalent to showing that any two of them are interlacing. Thus, it is enough to show that if $B, C \in U(d)$ satisfy $\text{rank}(BC^{-1} - I_d) = 1$, then the polynomials $P_{\mathcal{W}'}(BA_1B^*, A_2, \dots, A_m)$ and $P_{\mathcal{W}'}(CA_1C^*, A_2, \dots, A_m)$ interlace.

By replacing A_1 with CA_1C^* and writing $D = BC^{-1}$ we need to prove that $P_{\mathcal{W}'}(DA_1D^*, A_2, \dots, A_m)$ and $P_{\mathcal{W}'}(A_1, A_2, \dots, A_m)$ interlace (now $\text{rank}(D - I_d) = 1$). We

claim that $DA_1D^* - A_1$ is a rank-2 trace-0 Hermitian matrix: by unitary conjugation we

may assume $D = \begin{pmatrix} \lambda & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{pmatrix} \in \Lambda$, and direct calculation then shows that

$$DA_1D^* - A_1 = \begin{pmatrix} 0 & \lambda a_{1,2} & \cdots & \lambda a_{1,d} \\ \bar{\lambda} a_{2,1} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \bar{\lambda} a_{d,1} & 0 & \cdots & 0 \end{pmatrix}.$$

Let $\pm\nu$ ($\nu \in \mathbb{R}_{\geq 0}$) be the non-zero eigenvalues of $DA_1D^* - A_1$. By spectral decomposition, $DA_1D^* - A_1 = uu^* - vv^*$ for some vectors $u, v \in \mathbb{C}^d$ of length $\sqrt{\nu}$. Consider also $P_{\mathcal{W}'}(A_1 - vv^*, A_2, \dots, A_m)$ and denote

$$\begin{aligned} \alpha_d \leq \dots \leq \alpha_1 & \quad \text{the roots of } P_{\mathcal{W}'}(A_1, A_2, \dots, A_m) \\ \beta_d \leq \dots \leq \beta_1 & \quad \text{the roots of } P_{\mathcal{W}'}(A_1 - vv^*, A_2, \dots, A_m) \\ \gamma_d \leq \dots \leq \gamma_1 & \quad \text{the roots of } P_{\mathcal{W}'}(DA_1D^*, A_2, \dots, A_m). \end{aligned}$$

The assumptions of Lemma 4.3 are satisfied for \mathcal{W}' by the induction hypothesis. We can apply this lemma on $P_{\mathcal{W}'}(A_1 - vv^*, A_2, \dots, A_m)$ and $P_{\mathcal{W}'}(A_1, A_2, \dots, A_m)$ to obtain that

$$\beta_d \leq \alpha_d \leq \beta_{d-1} \leq \alpha_{d-1} \leq \dots \leq \beta_2 \leq \alpha_2 \leq \beta_1 \leq \alpha_1,$$

and on

$$P_{\mathcal{W}'}(A_1 - vv^*, A_2, \dots, A_m) \text{ and } P_{\mathcal{W}'}(DA_1D^*, A_2, \dots, A_m) = P_{\mathcal{W}'}(A_1 - vv^* + uu^*, A_2, \dots, A_m)$$

to obtain that

$$\beta_d \leq \gamma_d \leq \beta_{d-1} \leq \gamma_{d-1} \leq \dots \leq \beta_2 \leq \gamma_2 \leq \beta_1 \leq \gamma_1.$$

It follows that $P_{\mathcal{W}'}(DA_1D^*, A_2, \dots, A_m)$ and $P_{\mathcal{W}'}(A_1, A_2, \dots, A_m)$ interlace. This completes the proof. \square

4.2 Average Characteristic Polynomial of Random Coverings

Let G be a finite graph *without loops*, Γ a group and $\pi : \Gamma \rightarrow \text{GL}_d(\mathbb{C})$ a unitary representation. We now deduce from Theorem 4.2 that for certain distributions of (Γ, π) -coverings of G , the average characteristic polynomial is real rooted. Recall that $\phi_{\gamma, \pi}$ denotes the characteristic polynomial of $A_{\gamma, \pi}$ — see (1.2).

Proposition 4.4. *Let X_1, \dots, X_r be independent random variables, each taking values in the space of Γ -labelings of G , such that any two possible values γ_1 and γ_2 of X_i agree on all edges in $E^+(G)$ but one, and on that edge $\text{rank}(\pi(\gamma_1\gamma_2^{-1}(e)) - I_d) = 1$. Then $\mathbb{E}_{X_1, \dots, X_r}[\phi_{X_1, \dots, X_r, \pi}]$ is real-rooted.*

Proof. As we noted in the proof of Claim 2.9, for any Γ -labeling γ , we have $\phi_{\gamma,\pi} = \phi_{\gamma,\pi'}$ whenever π and π' are isomorphic, so we assume without loss of generality that $\pi : \Gamma \rightarrow U(d)$. For every Γ -labeling γ , the matrix $A_{\gamma,\pi}$ is a $nd \times nd$ matrix composed of n^2 blocks of size $d \times d$. The blocks are indexed by ordered pairs of vertices of G . Similarly to a notation we used in Section 3, for any $e \in E^+(G)$, we let $A_{\gamma,\pi}^\pm(e) \in M_{nd}$ be the matrix with zero blocks except for the blocks corresponding to e and to $-e$. In the block $(h(e), t(e))$ we have $\pi(\gamma(e))$ and in the block $(t(e), h(e))$ we have $\pi(\gamma(-e)) = \pi(\gamma(e))^*$. It is clear that $A_{\gamma,\pi}^\pm(e)$ is Hermitian and that

$$A_{\gamma,\pi} = \sum_{e \in E^+(G)} A_{\gamma,\pi}^\pm(e).$$

For every random Γ -labeling X of G and $e \in E^+(G)$, denote by $W_e(X)$ the following random matrix in $U(nd)$:

$$W_e(X) = \begin{pmatrix} I_d & & & \\ & \ddots & & \\ & & \pi(X(e)) & \\ & & & \ddots \\ & & & & I_d \end{pmatrix}$$

where all non-diagonal $d \times d$ blocks are zeros, the $(h(e), h(e))$ block is $\pi(X(e))$, and the remaining diagonal blocks are I_d . Also let $\mathbf{1} : E(G) \rightarrow \Gamma$ be the trivial labeling which labels all edges by the identity element of Γ . With these notations, we have¹⁸

$$A_{X_1, \dots, X_r, \pi}^\pm(e) = W_e(X_1) \cdot \dots \cdot W_e(X_r) A_{\mathbf{1}, \pi}^\pm(e) W_e(X_r)^* \cdot \dots \cdot W_e(X_1)^*,$$

and

$$A_{X_1, \dots, X_r, \pi} = \sum_{e \in E^+(G)} A_{X_1, \dots, X_r, \pi}^\pm(e).$$

By assumption, the random Γ -labeling X_i is constant on all edges except for on one edge e . Thus, $\{X_i(e)\}_{e \in E^+(G)}$ is a set of independent variables. Moreover, the set $\{W_e(X_i)\}_{e \in E^+(G), 1 \leq i \leq r}$ are a set of independent variables taking values in $U(nd)$, and every $W_e(X_i)$ is $U(1)$ -like by the assumption on the values of X_i . The proposition now follows by applying Theorem 4.2. \square

As explained in Section 2.5, we deduce:

Corollary 4.5. *In the notation of Proposition 4.4, there is a Γ -labelings $\gamma = \gamma_1 \cdot \dots \cdot \gamma_r$ of G , γ_i in the support of X_i , so that the largest root of $\phi_{\gamma,\pi}$ is at most the largest root of $\mathbb{E}_{X_1, \dots, X_r}[\phi_{Y,\pi}]$.*

¹⁸This formula is exactly the place this proof breaks for loops.

4.3 Proof of Theorem 1.12

We finally have all the tools needed to prove Theorem 1.12. Let G be a finite, loopless graph, and let (Γ, π) satisfy property $(\mathcal{P}2)$, namely, Γ is a finite group, $\pi : G \rightarrow \text{GL}_d(\mathbb{C})$ a representation and $\pi(\Gamma)$ a complex reflection group (i.e. generated by pseudo-reflections). Assume that $\Gamma = \langle g_1, \dots, g_s \rangle$ (Γ is generated by g_1, \dots, g_s), and that $\pi(g_i)$ is a pseudo-reflection for all i , so $\text{rank}(\pi(g_i) - I_d) = 1$. We first show that a certain random walk on Γ , where in each step we use only one of the g_i 's, converges to the uniform distribution:

Claim 4.6. *Define a random walk $\{a_n\}_{n=0}^\infty$ on Γ as follows: $a_0 = \mathbf{1}_\Gamma$ (the identity element of Γ), and for $n \geq 1$*

$$a_n = \begin{cases} g_{n \bmod s} \cdot a_{n-1} & \text{with probability } \frac{1}{3} \\ g_{n \bmod s}^{-1} \cdot a_{n-1} & \text{with probability } \frac{1}{3} \\ a_{n-1} & \text{with probability } \frac{1}{3} \end{cases}.$$

Then a_n converges to the uniform distribution on Γ as $n \rightarrow \infty$.

Proof. Consider a_n as an element of the group-ring $\mathbb{C}[\Gamma]$ so that the coefficient of g is $\text{Prob}[a_n = g]$. Then for $n \geq 1$,

$$a_{s \cdot n} = \left(\frac{1}{3} \mathbf{1}_\Gamma + \frac{1}{3} g_s + \frac{1}{3} g_s^{-1} \right) \cdots \left(\frac{1}{3} \mathbf{1}_\Gamma + \frac{1}{3} g_1 + \frac{1}{3} g_1^{-1} \right) \cdot a_{s \cdot (n-1)}.$$

The s -step random walk $\{a_{s \cdot n}\}_{n=0}^\infty$ is defined by the distribution $h = \left(\frac{1}{3} \mathbf{1}_\Gamma + \frac{1}{3} g_s + \frac{1}{3} g_s^{-1} \right) \cdots \left(\frac{1}{3} \mathbf{1}_\Gamma + \frac{1}{3} g_1 + \frac{1}{3} g_1^{-1} \right)$, which is symmetric with $\langle \text{supp}(h) \rangle = \Gamma$. Moreover, this random walk is lazy (in every step it has a positive probability of staying at the same element). Thus, it converges to the only stationary distribution of this Markov chain: the uniform distribution. The same argument applies to $\{a_{s \cdot n + i}\}_{n=0}^\infty$ for any modulo $1 \leq i \leq s-1$. \square

Now define random Γ -labelings $\{Z_n\}_{n=1}^\infty$ of G as follows: let $\varepsilon = |E^+(G)|$ and enumerate in an arbitrary way the edges of G , so $E^+(G) = \{e_1, \dots, e_\varepsilon\}$. For $i \geq 1$ and $1 \leq j \leq \varepsilon$ define $X_{i,j}$ to be the random Γ -labeling of G which labels every edge besides e_j with the identity element $\mathbf{1}_\Gamma$, and

$$X_{i,j}(e_j) = \begin{cases} g_{i \bmod s} & \text{with probability } \frac{1}{3} \\ g_{i \bmod s}^{-1} & \text{with probability } \frac{1}{3} \\ \mathbf{1}_\Gamma & \text{with probability } \frac{1}{3} \end{cases}.$$

Now define $Y_i = X_{i,1} \cdots X_{i,\varepsilon}$ and $Z_n = Y_1 Y_2 \cdots Y_n$. By definition, each random Γ -labeling $X_{i,j}$ is constant on every edge except one, and on the remaining edge every two values differ by a pseudo-reflection. Proposition 4.4 yields, therefore, that $\mathbb{E}_{Z_n}[\phi_{Z_n, \pi}]$ is real-rooted. By Claim 4.6, the random Γ -coverings Z_n converge, as $n \rightarrow \infty$, to the uniform distribution $\mathcal{C}_{\Gamma, G}$ of all Γ -labelings of G . Since the map $Z \rightarrow \mathbb{E}_Z[\phi_{Z, \pi}]$ is a continuous map from the space of

distributions of Γ -labelings of G to $\mathbb{R}[x]$, we get that $\mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma, G}} [\phi_{\gamma, \pi}]$ is real-rooted, which is the first statement of Theorem 1.12.

Finally, by Corollary 4.5, for every n , there is a Γ -labeling γ_n of G so that the largest root of $\phi_{\gamma, \pi}$ is at most the largest root of $\mathbb{E}_{Z_n} [\phi_{Z_n, \pi}]$. Because the set of Γ -labeling of G is finite, the γ_n have a limit point γ_0 . As the largest root of $\mathbb{E}_{Z_n} [\phi_{Z_n, \pi}]$ converges, as $n \rightarrow \infty$, to the largest root of $\mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma, G}} [\phi_{\gamma, \pi}]$, the largest root of $\phi_{\gamma_0, \pi}$ is at most the largest root of $\mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma, G}} [\phi_{\gamma, \pi}]$. This completes the proof of Theorem 1.12.

Remark 4.7. When $\Gamma = S_d$ is the symmetric group, [MSS15c, Lemma 3.5] gives a specific sequence of $2^{d-1} - 1$ random $U(1)$ -like permutations (“random swaps” in their terminology) the product of which is the uniform distribution on S_d .

Remark 4.8. We stated Property $(\mathcal{P}2)$ and Theorem 1.12 for finite groups only. However, it seems the result can be generalized to compact groups with unitary representations. The condition that $\pi(\Gamma)$ is a complex reflection group should be then something like “there is a set of subgroups in $\pi(\Gamma)$, each one of which conjugate to a subgroup of Λ , which together generate a dense subgroup of Γ ”.

5 On Pairs Satisfying $(\mathcal{P}1)$ and $(\mathcal{P}2)$ and Further Applications

In this Section we say a few words about pairs (Γ, π) of a group and a representation satisfying properties $(\mathcal{P}1)$ and/or $(\mathcal{P}2)$, and elaborate on the combinatorial applications of Theorem 1.10 (in addition to the existence of one-sided Ramanujan d -coverings as stated in Theorem 1.2). We begin with $(\mathcal{P}2)$, where a complete classification is known.

5.1 Complex Reflection Groups

Recall that the pair (Γ, π) satisfies $(\mathcal{P}2)$ if Γ is finite and $\pi(\Gamma)$ is a complex reflection group, namely generated by pseudo-reflections (elements $A \in \mathrm{GL}_d(\mathbb{C})$ of finite order with $\mathrm{rank}(A - I_d) = 1$). If π is not faithful (not injective), it factors through the faithful $\bar{\pi}: \Gamma/\ker \pi \rightarrow \mathrm{GL}_d(\mathbb{C})$, and (Γ, π) satisfies $(\mathcal{P}2)$ if and only if $(\Gamma/\ker \pi, \bar{\pi})$ does. In addition, if π is faithful but reducible and (Γ, π) satisfies $(\mathcal{P}2)$, then necessarily there are pairs (Γ_1, π_1) and (Γ_2, π_2) satisfying $(\mathcal{P}2)$ with $\Gamma \cong \Gamma_1 \times \Gamma_2$ and $\pi \cong (\pi_1, 1) \oplus (1, \pi_2)$.

Hence, the classification of pairs satisfying $(\mathcal{P}2)$ boils down to classifying finite, irreducible complex reflection groups: finite-order matrix groups inside $\mathrm{GL}_d(\mathbb{C})$ which are generated by pseudo-reflections and have no invariant proper subspaces of \mathbb{C}^d . This classification was established in 1954 by Shephard and Todd:

Theorem 5.1. [ST54] *Any finite irreducible complex reflection group W is one the following:*

1. $W \leq \mathrm{GL}_d(\mathbb{C})$ is isomorphic to S_{d+1} ($d \geq 2$), via the standard representation of S_{d+1} (see the paragraph preceding Claim 2.10).

2. $W = G(m, p, d)$ with $m, d \in \mathbb{Z}_{\geq 2}$, $p \in \mathbb{Z}_{\geq 1}$ and $p|m$. This is a generalization of signed permutations groups: the group $G(m, p, d) \leq \mathrm{GL}_d(\mathbb{C})$ consists of monomial matrices (matrices with exactly one non-zero entry in every row and every column), the non-zero entries are m -th roots of unity (not necessarily primitive), and their product is a $\frac{m}{p}$ -th root of unity. This a group of order $\frac{d! \cdot m^d}{p}$. For example,

$$\begin{pmatrix} 0 & \zeta & 0 \\ 0 & 0 & \zeta^{-1} \\ \zeta^4 & 0 & 0 \end{pmatrix}$$

where $\zeta = e^{\frac{2\pi i}{6}}$, is an element of $G(6, 2, 3)$.

3. $W = \mathbb{Z}/m\mathbb{Z} \leq \mathrm{GL}_1(\mathbb{C})$, $m \in \mathbb{Z}_{\geq 2}$ is the cyclic group of order m whose elements are m -th roots of unity (can be denoted $G(m, 1, 1)$).
4. W is one of 34 exceptional finite irreducible complex reflection groups of different dimensions d , $2 \leq d \leq 8$.

We remark that a finite complex reflection group conjugated to a subgroup of $\mathrm{GL}_d(\mathbb{R})$ is, by definition, a finite Coxeter group. All groups listed in the theorem are finite complex reflection groups, and all irreducible except for $G(2, 2, 2)$.

Theorem 5.2. [Steinberg, [GM06, Thm 4.6]] If (Γ, π) satisfies (P2) and π is irreducible, then (Γ, π) satisfies (P1) as well.

Namely, if Γ is a finite group, $\pi: \Gamma \rightarrow \mathrm{GL}_d(\mathbb{C})$ an irreducible representation and $\pi(\Gamma)$ is a complex reflection group, then the exterior powers $\bigwedge^r \pi$, $0 \leq r \leq d$, are irreducible and non-isomorphic.

Evidently, if π is reducible, the pair (Γ, π) does not satisfy (P1). Thus,

Corollary 5.3. The pairs (Γ, π) satisfying both (P1) and (P2) are precisely the irreducible finite complex reflection groups¹⁹.

Finally, consider two pairs (Γ_1, π_1) and (Γ_2, π_2) , and a third pair (Γ, π) constructed as their direct product: $\Gamma \cong \Gamma_1 \times \Gamma_2$ and $\pi \cong (\pi_1, 1) \oplus (1, \pi_2)$. Constructing a (Γ, π) -covering of a graph is equivalent to constructing independent coverings, one for (Γ_1, π_1) and one for (Γ_2, π_2) . We conclude:

Corollary 5.4. Let (Γ, π) satisfy (P2), and let G be a finite graph with no loops. Then,

- If d_1, \dots, d_r are the dimensions of the irreducible components of π , then

$$\mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma, G}} [\phi_{\gamma, \pi}] = \mathcal{M}_{d_1, G} \cdot \dots \cdot \mathcal{M}_{d_r, G}.$$

- $\mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma, G}} [\phi_{\gamma, \pi}]$ is real rooted and there is some labeling with smaller largest root.
- There is a one-sided Ramanujan (Γ, π) -covering of G .

¹⁹To be precise, this is true for faithful representations. If π factors through $\bar{\pi}: \Gamma/\ker \pi \rightarrow \mathrm{GL}_d(\mathbb{C})$, then (Γ, π) satisfies (P1) and (P2) if and only if so does $(\Gamma/\ker \pi, \bar{\pi})$.

5.2 Pairs Satisfying $(\mathcal{P}1)$

The list in Theorem 5.1 does not exhaust all pairs (Γ, π) (with π faithful) satisfying $(\mathcal{P}1)$. Even when restricting to finite groups, there are pairs satisfying $(\mathcal{P}1)$ but not $(\mathcal{P}2)$. A handful of such examples arises from the observation that $(\mathcal{P}1)$ is preserved by passing to bigger groups:

Claim 5.5. *Let Γ be a group, $\pi: \Gamma \rightarrow \mathrm{GL}_d(\mathbb{C})$ a representation and $\Lambda \leq \Gamma$ a subgroup. If $(\Lambda, \pi|_\Lambda)$ satisfies $(\mathcal{P}1)$ then so does (Γ, π) .*

Proof. It is clear that if $\bigwedge^r \pi$ cannot have an invariant proper subspace if $(\bigwedge^r \pi)|_\Lambda$ has none. An isomorphism of $\bigwedge^r \pi$ and $\bigwedge^{d-r} \pi$ induces an isomorphism on the same representation restricted to Λ . \square

For example, we can increase $\mathrm{std}(S_{d+1})$ by adding some scalar matrix of finite order m as an extra generator, and obtain a d -dimensional faithful representation of $S_{d+1} \times \mathbb{Z}/m\mathbb{Z}$ which satisfies $(\mathcal{P}1)$.

There are also pairs with Γ finite which do not contain any complex reflection group. For instance, consider the index-2 subgroup Γ of $G(2, 1, 3)$ where we restrict to *even* permutation 3×3 matrices with ± 1 signing of every non-zero entry. The natural 3-dimensional representation of this group satisfies $(\mathcal{P}1)$, but does not contain any complex reflection group. We are not aware of a full classification of pairs (Γ, π) satisfying $(\mathcal{P}1)$, even when Γ is finite.

There are some interesting examples of pairs (Γ, π) satisfying $(\mathcal{P}1)$ where Γ is infinite and compact. For example, the standard representation π of the orthogonal group $O(d)$ or of the unitary group $U(d)$, is such (by, e.g., Claim 5.5 and the fact one can identify $\mathrm{std}(S_{d+1})$ as a subgroup of $O(d)$ or of $U(d)$).

Corollary 5.6. *Let $\Gamma = O(d)$ or $\Gamma = U(d)$, and let π be the standard d -dimensional representation. Then, for every finite graph G ,*

$$\mathbb{E}_{\gamma \in \mathcal{C}_{\Gamma, G}} [\phi_{\gamma, \pi}] = \mathcal{M}_{d, G}.$$

5.3 Applications of Theorem 1.10

In this section we elaborate the combinatorial consequences of Theorem 1.10 stating that if (Γ, π) satisfies both $(\mathcal{P}1)$ and $(\mathcal{P}2)$, then there is a one-sided Ramanujan (Γ, π) -covering of G whenever G is finite with no loops. Corollary 5.3 tells us exactly what pairs satisfy the conditions of the theorem. The most interesting consequence, based on the pair (S_{d+1}, std) , was already stated as Theorem 1.2: every G as above has a one-sided Ramanujan d -covering for every d .

Another interesting application stems from one-dimensional representations (item (3) in Theorem 5.1):

Corollary 5.7. *For every $m \in \mathbb{Z}_{\geq 2}$ and every loopless²⁰ finite graph G , there is a labeling of the oriented edges of G by m -th roots of unity (with $\gamma(-e) = \gamma(e)^{-1}$, as usual), so that the resulting spectrum is one-sided Ramanujan.*

²⁰In this special case it is actually possible to prove the result even for graphs with loops.

Of course, the result for m follows from the result for d whenever $1 \neq d|m$. For $m = 2$ this is the main result of [MSS15a]. As this corollary deals only with one-dimensional representations, the original proof of [MSS15a] can be relatively easily adapted to show it. This was noticed also by [LPV14].

Recall that all irreducible representations of abelian groups are one-dimensional. Therefore, given an abelian group Γ and a finite graph G , there is a Γ -labeling of G which yields a one-sided Ramanujan (Γ, π) -covering for any irreducible representation π of Γ . However, this certainly does not guarantee the existence of a single Γ -labeling which is “Ramanujan” for all irreducible representations simultaneously. If true, this would guarantee the existence of a one-sided Ramanujan (Γ, \mathcal{R}) -covering of G , where \mathcal{R} is the regular representation. As mentioned in the paragraph following Definition 1.7, in the special case where G is a bouquet of one vertex with r loops, a (Γ, \mathcal{R}) -covering is a Cayley graph of Γ with respect to r group elements. But it is well-known (and easy to prove) that with r fixed, large abelian groups do not have Ramanujan Cayley graphs with r generators (in fact, the spectral gap tends to 0 as the size of Γ grows).

An even easier counter-example stems from cases when $\text{rank}(\Gamma) > \text{rank}(\pi_1(G))$ (here $\text{rank}(\Gamma)$ marks the minimal size of a generating set of Γ). In this case, there is no surjective homomorphism $\pi_1(G) \rightarrow \Gamma$, so every (Γ, \mathcal{R}) -covering is necessarily disconnected, and the spectral gap is zero (see Claim 2.8).

Still, in the special case where $\Gamma = \mathbb{Z}/3\mathbb{Z}$ is the cyclic group of order 3, there are only two non-trivial representations π_1 and π_2 , and one is the complex-conjugate of the other. Hence, $\phi_{\gamma, \pi_2} = \phi_{\gamma, \pi_1}$ for any Γ -labeling γ , and so the spectra are identical, and we get, as noticed by [LPV14, CV15]²¹:

Corollary 5.8. *Every finite graph G has a one-sided Ramanujan 3-covering, where the permutation above every edge is cyclic.*

From the third infinite family of complex reflection groups (item (2) in Theorem 5.1), we do not get any significant combinatorial implications. If $\Gamma = G(m, p, d)$, Theorem 1.10 guarantees that every graph has a one-sided Ramanujan “signed d -covering”: a d -covering of G where every oriented edge is then labeled by an m -th root of unity, and such that the product of roots in every fiber of edges is an $\frac{m}{p}$ -th root of unity. But Corollary 5.7 shows that *every* d -covering of G can be edge-labeled by m -th roots of unity so that the resulting spectrum is one-sided Ramanujan. If $p < m$, we can label by $\frac{m}{p}$ -th roots, so applying Theorem 1.10 on Γ yields nothing new. If $p = m$, the statement of the theorem cannot be (easily) derived from former results: we get that G has a d -covering with edges labeled by m -th roots of unity, so that the product of the labels on every fiber is 1, and the resulting spectrum is one-sided Ramanujan.

²¹Interestingly, it is also shown in [CV15] that every graph has a one-sided Ramanujan 4-covering with cyclic permutations. This does not follow from the results in the current paper.

5.4 Ramanujan Topological Coverings of Special Kinds

Every group action of Γ on a finite set X yields a representation π of dimension $|X|$. In this case, π can be taken to map Γ into permutation matrices, hence (Γ, π) -coverings of a graph G correspond to topological $|X|$ -coverings of G (with permutations restricted to the set $\pi(\Gamma)$). For instance, the natural action of S_d on $\{1, \dots, d\}$ yields the set of all d -coverings from Theorem 1.2. The action of $\mathbb{Z}/3\mathbb{Z}$ by cyclic shifts on a set of size 3 yields the regular representation of this group and the coverings in Corollary 5.8. In general, the regular representation of a group is always of this kind.

It is interesting to consider the set \mathcal{A} of all possible pairs (Γ, π) where Γ is a finite group and π an action-representation which guarantees (one-sided) Ramanujan coverings of every graph. Of course, the action must be transitive: otherwise, the coverings are never connected. Observe this set is closed under two “operations”:

1. If $\Lambda \leq \Gamma$ and $(\Lambda, \pi|_\Lambda)$ is in \mathcal{A} , then so is (Γ, π) .
2. The set \mathcal{A} is closed under towers of coverings: a Ramanujan covering of a Ramanujan covering is a Ramanujan covering of the original graph. In algebraic terms this corresponds to wreath products of groups. Namely, if (Γ, π) and (Λ, ρ) are both in \mathcal{A} with respect to actions on the sets X and Y , respectively, then so is the pair $(\Gamma \text{wr}_X \Lambda, \psi)$, where

$$\Gamma \text{wr}_Y \Lambda = \left(\prod_{y \in Y} \Gamma_y \right) \rtimes \Lambda$$

is the restricted wreath product (Γ_y is a copy of Γ for every $y \in Y$, and Λ acts by permuting the copies according to its action on Y), and ψ is based on the action of $\Gamma \text{wr}_Y \Lambda$ on the set $X \times Y$ by

$$(\{g_y\}, \ell) \cdot (x, y) = (g_y \cdot x, \ell \cdot y).$$

In this language, for example, a tower of 2-coverings, as considered by Bilu-Linial and Marcus-Spielman-Srivastava, corresponds to a pair (Γ, π) with Γ a nested wreath product of $\mathbb{Z}/2\mathbb{Z}$.

6 Open Questions

We finish with some open questions arising naturally from the discussion in this paper.

Question 6.1. Irreducible representations and one-sided Ramanujan coverings: Which pairs (Γ, π) of a finite (compact) group with an irreducible finite-dimensional representation guarantee the existence of one-sided Ramanujan (Γ, π) -coverings for every finite graph? Can the statement of Theorem 1.10 be extended to more pairs? Does (P1) suffice? In fact, we are not aware of a single example of a pair (Γ, π) with π irreducible and non-trivial and a finite graph G so that there is no (one-sided) Ramanujan (Γ, π) -covering of G .

Question 6.2. Full Ramanujan coverings: The previous question can be asked for full (two-sided) Ramanujan coverings as well. The difference is that in this case nothing is known for general graphs. The case $(\mathbb{Z}/2\mathbb{Z}, \pi)$ with π the non-trivial one-dimensional representation is the Bilu-Linial Conjecture [BL06]. We conjecture this can be generalized and if (Γ, π) satisfies (P1) and (P2), then every finite graph has a (full) Ramanujan (Γ, π) -covering.

In particular, proving this conjecture, even the stricter form of Bilu-Linial, would yield the existence of infinitely many k -regular, non-bipartite, Ramanujan graphs for every $k \geq 3$.

Question 6.3. Regular representations and Cayley graphs: We find special interest in $(\Gamma, \mathcal{R} - \text{triv})$, where Γ is finite and \mathcal{R} is its regular representation, because (Γ, \mathcal{R}) -coverings of graphs generalize the notion of Cayley graphs (these are (Γ, \mathcal{R}) -coverings of bouquets). As remarked in Section 5.3, Theorem 1.10 cannot hold in general in this case, even if all irreducible representations of Γ satisfy the conditions of the theorem. However, if for every irreducible π and a random Γ -labeling of the graph G , the (γ, π) -covering is Ramanujan with very high probability, it is possible to find a Γ -labeling that works for all irreducible representations of Γ simultaneously. Is it possible that, given Γ , if G is “large” enough and has good expansions properties (e.g. if G is Ramanujan), then G necessarily has a (one-sided) Ramanujan (Γ, \mathcal{R}) -covering? A result in this direction is given in [BL06], where it is shown that if G is a good expander, then a random 2-covering of G has a large spectral gap with high probability.

Question 6.4. The d -matching polynomial: In the current paper, we defined $\mathcal{M}_{d,G}$, the d -matching polynomial of the graph G and showed it has some properties which parallel those of the classical matching polynomial, $\mathcal{M}_{1,G}$. But $\mathcal{M}_{1,G}$ has many more interesting properties (a good reference is [God93]). What parts of this Theory can be generalized to $\mathcal{M}_{d,G}$?

In particular, it would be desirable to find a proof to the real-rootedness of $\mathcal{M}_{d,G}$, which is more direct than the one given in this paper. Such a proof may work just as well for graph with loops.

Question 6.5. Dealing with loops: Some of the results of this paper hold for any finite graph, even with loops (e.g. Theorem 1.11). We conjecture that, in fact, all the results hold for graphs with loops. In particular, we conjecture that any finite graph G with loops should have a one-sided Ramanujan d -covering (Theorem 1.2), that $\mathcal{M}_{d,G}$ is real-rooted for every d (Theorem 1.5) and that if (Γ, π) satisfies (P1) and (P2), then G has a one-sided Ramanujan (Γ, π) -covering (Theorem 1.10). (And see Question 6.6.)

If true, this would yield, for example, that if A is a uniformly random permutation matrix, or Haar-random orthogonal or unitary matrix in $U(d)$, then $\mathbb{E}[\phi(A + A^*)]$ is real-rooted.

Question 6.6. Another interlacing family of characteristic polynomials: The one argument in this paper that breaks for loops is in the proof of Proposition 4.4. The problem is that if e is a loop, then $\pi(\gamma(e))$ and $\pi(\gamma(e))$ lie in the same $d \times d$ block of $A_{\gamma, \pi}$. One way to extend the arguments for loops is to prove the following parallel of Theorem 4.2, which we believe should hold:

For a matrix A denote $A^{\mathcal{HERM}} \stackrel{\text{def}}{=} A + A^*$. If $\mathcal{W} = \{W_{i,j}\}_{1 \leq i \leq m, 1 \leq j \leq \ell(i)}$ is defined as in Theorem 4.2, then

$$\mathbb{E}_{\mathcal{W}} \left[\phi \left([W_{1,1} \dots W_{1,\ell(1)} A_1 + \dots + W_{m,1} \dots W_{m,\ell(m)} A_m]^{\mathcal{HERM}} \right) \right]$$

is real-rooted for every $A_1, \dots, A_m \in M_d(\mathbb{C})$.

If true, this statement generalizes the fact that the characteristic polynomials $\phi(A + A^*)$ and $\phi(BA + (BA)^*)$ interlace whenever $A, B \in GL_d(\mathbb{C})$ with B a pseudo-reflection.

Question 6.7. Half-loops and random perfect matchings: A common model for generating a random k -regular graph on $2n$ vertices is by considering the union of k independent random perfect matchings of the vertices. The technique of [MSS15c, Section 3] can be used to show the expected characteristic polynomial is real-rooted. In the language of the current paper, this model can be thought of as a random $2n$ -covering of a bouquet of one vertex with k “half-loops”: every half-loop contributes 1 to the degree of the vertex and to the corresponding diagonal entry of the adjacency matrix. In a random cover, every half-edge is labeled by a random involution in S_{2n} , consisting of n 2-cycles. The arguments in Section 4 of the current paper can show this, namely Theorem 1.12, can be extended to arbitrary graphs with half-loops (but not with ordinary loops).

It is also shown in [MSS15c], that for the bouquet case, the expected characteristic polynomial is one-sided Ramanujan, namely, its second eigenvalue is at most $2\sqrt{k-1}$. This does not follow from the arguments in this paper. We wonder if there is a parallel to Theorem 1.11 if the base graph G contains half-loops.

Acknowledgments

We would like to thank Miklós Abért, Péter Csikvári, Nati Linial and Ori Parzanchevski for valuable discussions regarding some of the themes of this paper. We also thank Daniel Spielman for sharing with us an early version of [MSS15c].

References

- [AL02] A. Amit and N. Linial, *Random graph coverings I: General theory and graph connectivity*, *Combinatorica* **22** (2002), no. 1, 1–18.
- [BL06] Y. Bilu and N. Linial, *Lifts, discrepancy and nearly optimal spectral gap*, *Combinatorica* **26** (2006), no. 5, 495–519.
- [Bum04] D. Bump, *Lie groups*, Graduate Texts in Mathematics, vol. 225, Springer, 2004.
- [CV15] K. Chandrasekaran and A. Velingker, *Towards constructing Ramanujan graphs using shift lifts*, arXiv preprint arXiv:1502.07410, 2015.

- [FH91] W. Fulton and J. Harris, *Representation theory*, vol. 129, Springer Science & Business Media, 1991.
- [Fis06] S. Fisk, *Polynomials, roots, and interlacing*, arXiv preprint math/0612833, 2006.
- [Fri03] J. Friedman, *Relative expanders or weakly relatively Ramanujan graphs*, Duke Mathematical Journal **118** (2003), no. 1, 19–35.
- [GG81] C. D. Godsil and I. Gutman, *On the matching polynomial of a graph*, Algebraic Methods in Graph Theory (L. Lovász and V.T. Sós, eds.), Colloquia Mathematica Societatis János Bolyai, vol. 25, János Bolyai Mathematical Society, 1981, pp. 241–249.
- [GM06] M. Geck and G. Malle, *Reflection groups*, Handbook of Algebra (W. Hazewinkel, ed.), vol. 4, Elsevier, 2006, pp. 337–383.
- [God93] C.D. Godsil, *Algebraic combinatorics*, CRC Press, 1993.
- [GP14] K. Golubev and O. Parzanchevski, *Spectrum and combinatorics of Ramanujan triangle complexes*, arXiv preprint arXiv:1406.6666, 27 pages, 2014.
- [Gre95] Y. Greenberg, *On the spectrum of graphs and their universal coverings*, (in Hebrew), Ph.D. thesis, Hebrew University, 1995.
- [HL72] O.J. Heilmann and E. H. Lieb, *Theory of monomer-dimer systems*, Communications in Mathematical Physics **25** (1972), no. 3, 190–232.
- [HLW06] S. Hoory, N. Linial, and A. Wigderson, *Expander graphs and their applications*, Bulletin of the American Mathematical Society **43** (2006), no. 4, 439–562.
- [KRY09] J.P.S. Kung, G.C. Rota, and C. H. Yan, *Combinatorics: the rota way*, Cambridge University Press, 2009.
- [LN98] A. Lubotzky and T. Nagnibeda, *Not every uniform tree covers Ramanujan graphs*, Journal of Combinatorial Theory, Series B **74** (1998), no. 2, 202–212.
- [LPS88] A. Lubotzky, R. Phillips, and P. Sarnak, *Ramanujan graphs*, Combinatorica **8** (1988), no. 3, 261–277.
- [LPV14] S. Liu, N. Peyerimhoff, and A. Vdovina, *Signatures, lifts, and eigenvalues of graphs*, arXiv preprint arXiv:1412.6841, 2014.
- [Lub94] A. Lubotzky, *Discrete groups, expanding graphs and invariant measures*, Progress in Mathematics, vol. 125, Birkhauser, 1994.
- [MSS15a] A. Marcus, D. A. Spielman, and N. Srivastava, *Interlacing families I: Bipartite Ramanujan graphs of all degrees*, Annals of Mathematics **182** (2015), no. 1, 307–325.

- [MSS15b] ———, *Interlacing families II: Mixed characteristic polynomials and the kadison-singer problem*, Annals of Mathematics **182** (2015), no. 1, 327–350.
- [MSS15c] ———, *Interlacing families IV: Bipartite Ramanujan graphs of all sizes*, arXiv preprint arXiv:1505.08010v1, 21 pages, 2015.
- [Nil91] A. Nilli, *On the second eigenvalue of a graph*, Discrete Mathematics **91** (1991), no. 2, 207–210.
- [Pud15] D. Puder, *Expansion of random graphs: New proofs, new results*, Inventiones Mathematicae (2015), to appear. arXiv:1212.5216.
- [ST54] G. C. Shephard and J. A. Todd, *Finite unitary reflection groups*, Canad. J. Math **6** (1954), no. 2, 274–301.

Chris Hall,
 Department of Mathematics,
 University of Wyoming
 Laramie, WY 82071 USA
 chall14@uwyo.edu

Doron Puder,
 School of Mathematics,
 Institute for Advanced Study,
 Einstein Drive, Princeton, NJ 08540 USA
 doronpuder@gmail.com

William F. Sawin,
 Department of Mathematics,
 Princeton University
 Fine Hall, Washington Road
 Princeton NJ 08544-1000 USA
 wsawin@math.princeton.edu