Eigenvalue Multiplicities In Abelian Cayley Graphs

By

Easton Singer

Advised By

Salil Vadhan

A thesis presented to

Harvard University

in partial fulfillment of the requirements for the degree of

Bachelor of Arts with Honors

In Mathematics

Acknowledgments

I am deeply grateful to my advisors, Professor Salil Vadhan and Jake Ruotolo, both of whom graciously lent me countless hours of their time to mentor me over the last many months. Their comments and suggestions made this document what it is today. The original work that is presented here is the product of a research project that I undertook starting in the Spring 2025 semester, which has continued to date.

I would also like to acknowledge the Math Department and the theoretical computer science group at Harvard for providing the backbone of my academic pursuit over the last 4 years, and for enabling me to produce this document which represents a great milestone in my journey as a mathematician and a scholar.

Contents

1 Introduction

Over the last several decades, graphs have become ubiquitous in the fields of combinatorics and theoretical computer science. Many natural combinatorial objects can be viewed as graphs with the appropriate modifiers added (directed or undirected, simple or not, weighted or unweighted, etc). For example, Markov chains can be thought of as random walks on directed graphs, and logical circuits can be thought of as directed acyclic graphs augmented with some additional information about "gates."

However, the real power in working with graphs comes from the surprising relationship between their combinatorial aspects and other natural models: here, we will consider the "spectral" or "algebraic" perspective. Spectral graph theorists are largely interested in finding relationships between the combinatorial properties of a graph and the *spectrum*, or set of eigenvalues, of the graph's adjacency matrix. Perhaps the most famous example of such a theorem says that under very general conditions, the number of steps it takes for a random walk on a graph to closely approximate its stationary distribution is governed by the second-largest singular value of a scaled version of the adjacency matrix.

While a large portion of the spectral graph theory literature has focused on the eigen*values* of matrices like the adjacency matrix, some recent breakthroughs have been concerned instead with the *multiplicity* of eigenvalues of these associated graph matrices. Bounding the multiplicity of eigenvalues in very general classes of graphs can be extremely delicate, but even bounds on the order of $n/\text{poly}(\log\log n)$ (where n is the dimension of the matrix) are interesting (see, for example, [MRS23]). In a sequence of breakthroughs in [JTYZZ21] and [BB24], eigenvalue multiplicity bounds were used to derive new results in the study of equiangular lines.

This document is motivated by another line of works that look at eigenvalue multiplicities in a special class of graphs called *Cayley graphs*, which arise from group theory. In [ML08] and [HSZZ23], representation theory is a crucial tool for understanding eigenvalue multiplicity. In [DJRVZ24], a recent paper coauthored by my thesis mentors, similar techniques are employed to obtain a new approximation algorithm for the Sparsest Cut problem on abelian Cayley graphs.

When working over finite abelian groups, representation theory morphs into the theory of characters, which is roughly where our journey begins. In Section 2, we present the necessary background on graphs and spectral graph theory at large, and then move to a discussion of abelian Cayley graphs and how characters relate to their eigenvalues. The section concludes with several detailed examples.

In Section 3, we follow [DJRVZ24] and reproduce many of their proofs that bound eigenvalue multiplicities in abelian Cayley graphs. The main theorem in this section is an upper bound on the multiplicity of the second-largest adjacency matrix eigenvalue that is exponential in the degree of the graph, which applies for all abelian Cayley graphs. After this section, we shift into an exploration of the question of eigenvalue multiplicity in Cayley graphs over cyclic groups.

Section 4 deals with the case of cyclic groups of prime order, in which we prove the following theorem, simultaneously bounding all eigenvalue multiplicities by the degree of the graph:

Theorem 1.1. *Let p be a prime, and let $\Gamma = \mathbb{Z}/p\mathbb{Z}$. Given a non-empty multiset S of nonzero elements of Γ, consider the Cayley graph $G = Cay(\Gamma, S)$. We consider the natural identification of $\Gamma \setminus \{0\}$ with $\mathbb{F}_p^*$, the multiplicative group modulo p. If H is the largest subgroup of $\mathbb{F}_p^*$ such that S is a union of cosets of H, then the multiplicity of every eigenvalue except for λ_1[1] is exactly $|H|$.*

In Section 5 and Section 6, we look at the next simplest cases (algebraically speaking) of $\mathbb{Z}/(p^2)\mathbb{Z}$ and $\mathbb{Z}/(pq)\mathbb{Z}$, respectively. In the former regime, we achieve the tight upper bound of d on the multiplicity of the second-largest eigenvalue. In the latter, this upper bound no longer holds! However, we still obtain the result for "almost all" Cayley graphs.

Theorem 1.2. *Let p be an odd prime, and let $n = p^2$ and $\Gamma = \mathbb{Z}/n\mathbb{Z}$. Given a symmetric multiset S of nonzero elements of Γ, consider the Cayley graph $G = Cay(\Gamma, S)$. If G is connected, then the multiplicity of $\lambda_2(G)$ is at most $|S| = d$.*

Theorem 1.3. *Let $p < q$ be distinct odd primes, and let $n = pq$ and $\Gamma = \mathbb{Z}/n\mathbb{Z}$. Given a symmetric multiset S of nonzero elements of Γ, consider the Cayley graph $G = Cay(\Gamma, S)$. Suppose that G is connected. Then, if S is not **organized** (defined in Definition 6.5), the multiplicity of every eigenvalue of G is at most $|S| = d$.*

Conjecture 1.4. *Let $p < q$ be distinct primes, and let $n = pq$ and $\Gamma = \mathbb{Z}/n\mathbb{Z}$. Given a symmetric multiset S of nonzero elements of Γ, consider the Cayley graph $G = Cay(\Gamma, S)$. If G is connected, then the multiplicity of $\lambda_2(G)$ is at most $\frac{p}{p-1} \cdot d$.*

[1] As long as S is nonempty, the graph G will be connected, so the multiplicity of λ_1 is 1.

2 Graph Theory

2.1 Notation and Background

Definition 2.1. *A **(directed) graph** G is defined by a vertex set V and an edge set $E \subset V \times V$. Equivalently, we can think of E as being described by a function $w_E : V \times V \to \{0, 1\}$, where $w_E(a, b) = 1 \iff (a, b) \in E$. This motivates the more general definition of a **weighted directed graph** G, which is defined by a vertex set V and a weight function $w : V \times V \to \mathbb{R}^{\geq 0}$. If $w(a, b) = w(b, a)$ for every pair of vertices $a, b \in V$, we say that G is **undirected**. The **out-degree** of a vertex a is $d_{\text{out}}(a) := \sum_{b \in V} w(a, b)$, while its **in-degree** is $d_{\text{in}}(a) := \sum_{b \in V} w(b, a)$.*

Remark 2.2. Often, we will restrict our weights to be non-negative *integers*, rather than non-negative real numbers. This reflects the perspective of being allowed to consider "multiple edges" from a to b, but not "fractional edges."

It is natural to think of the weights of every pair of vertices in $V \times V$ as entries of a matrix.

Definition 2.3. *Given a graph G, its **adjacency matrix** $A(G)$ is the $|V| \times |V|$ real matrix with rows and columns indexed by V and entries $A(G)_{a,b} = w(a, b)$.*

For many applications, it helps to consider a small modification of the adjacency matrix, the Laplacian matrix.

Definition 2.4. *Given a graph G, its **Laplacian** $L(G)$ is the matrix $D_{\text{out}} - A(G)$, where D_{out} is a diagonal matrix with (a, a)-entry equal to the out-degree of vertex a for every $a \in V$.*

Observation 2.5. The Laplacian of any graph has an eigenvalue of 0 with right-eigenvector $\vec{1}$ (the all-ones column vector). To see this, note that if A is the adjacency matrix of a graph, then the ath entry of $A\vec{1}$ is $\sum_{b \in V} w(a, b) \cdot 1 = d_{\text{out}}(a)$, which is also the ath entry of $D_{\text{out}}\vec{1}$.

Definition 2.6. *A **(directed) path** in a graph G is a sequence of vertices $a_1, a_2, \ldots, a_m$ such that $w(a_i, a_{i+1}) > 0$ for all $i = 1, 2, \ldots, m - 1$. If for every pair of vertices v and w in G, there is a directed path starting at v and ending at w, we say that G is **strongly connected**. If G is undirected, we simply write that G is **connected**.*

Definition 2.7. *A graph G is **d-regular** if $d_{\text{out}}(a) = d_{\text{in}}(a) = d$ for every vertex $a \in V$. We say that G is **regular** if there exists a real number d such that G is d-regular.*

Observation 2.8. If G is a d-regular graph, then $L(G) = dI - A(G)$, where I is the $|V| \times |V|$ identity matrix. In particular, if v is an eigenvector of $A(G)$ with eigenvalue α, then $L(G)v = dv - \alpha v = (d - \alpha)v$, so v is also an eigenvector of $L(G)$ with eigenvalue $d - \lambda$. This implies that if $\alpha_1, \ldots, \alpha_n$ are the eigenvalues of $A(G)$ including multiplicity, then $d - \alpha_1, \ldots, d - \alpha_n$ are the eigenvalues of $L(G)$ including multiplicity.

Fact 2.9. *[Spi25, Lemma 3.1.1, Lemma 4.2.1] If G is an undirected d-regular graph, then the eigenvalues of $A(G)$ are all real numbers in the interval $[-d, d]$. Furthermore, d is an eigenvalue of $A(G)$, and it has multiplicity 1 if and only if G is connected. Analogously, in this case, the eigenvalues of $L(G)$ are real numbers in the interval $[0, 2d]$, and the eigenvalue 0 has multiplicity 1 if and only if G is connected. We denote by $0 = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n \leq 2d$ the eigenvalues of $L(G)$ including multiplicity.*

Remark 2.10. In light of the last two statements, we will freely interchange "multiplicity of λ_2" with "multiplicity of the second-largest adjacency matrix eigenvalue" when our graph is connected.

Remark 2.11. We will sometimes write "the eigenvalues of G" to mean "the eigenvalues of $A(G)$."

2.2 Cayley Graphs

We will be especially interested in are Cayley graphs, which comprise a diverse array of regular graphs motivated by algebra.

Definition 2.12. *A **multiset** is a generalization of a set which is allowed to have repeated elements. Given a universe $\mathcal{U}$, a multiset S of elements of $\mathcal{U}$ is a function $\text{mult}_S : \mathcal{U} \to \mathbb{N}$, where we think of $\text{mult}_S(s)$ as the number of copies of s in S. We also use $|S|$ to denote the total size of S (counting multiplicities). If $\mathcal{U}$ has a multiplication operation, we write rS for the multiset defined by $\text{mult}_{rS}(x) := \sum_{s:rs=x} \text{mult}_S(s)$.*

Example 2.13. Our primary example of this will be treating the integers modulo N as both a group under addition and with an associated multiplication operation. If we view $S = \{1, 1, 1, 3, 6, 10, 11, 11\}$ as a multiset of elements of $\mathbb{Z}/12\mathbb{Z}$, then $4S = \{0, 0, 4, 4, 4, 4, 8, 8\}$.

Definition 2.14 (Cayley Graph)**.** *Given a group $(\Gamma, \cdot)$ and a multiset S of elements of Γ, we define $G = Cay(\Gamma, S)$ to be the weighted graph with vertex set Γ and weight function $w(x, y) = \text{mult}_S(x^{-1}y)$.*

Remark 2.15. Every Cayley graph is d-regular, where $d = |S|$.

Remark 2.16. If the elements of S generate Γ, then $Cay(\Gamma, S)$ is strongly connected.

Remark 2.17. If S is **symmetric** (i.e. for every $x \in \Gamma$, we have $\text{mult}_S(x) = \text{mult}_S(x^{-1})$), then $Cay(\Gamma, S)$ is undirected.

Example 2.18 (Cycle)**.** Let $\Gamma = \mathbb{Z}/N\mathbb{Z}$ be the cyclic group of order $N > 0$. We say that $Cay(\Gamma, \{1\})$ is the *directed cycle* on N vertices, while $Cay(\Gamma, \{1, -1\})$ is the *undirected cycle* on N vertices.

Example 2.19 (Hypercube)**.** The *k-dimensional hypercube graph* is the graph with vertex set $\{0,1\}^k$ and edge set given by the edges of the unit hypercube with these vertices in k-space. It can be viewed as a Cayley graph on $\Gamma = (\mathbb{Z}/2\mathbb{Z})^k$ with generating set $S = \{e_1, e_2, \ldots, e_k\}$, where $e_i = (0, 0, \ldots, 1, \ldots, 0, 0)$ has its only 1 in the ith coordinate. In the graph $Cay(\Gamma, S)$, two vertices $(a_1, \ldots, a_k)$ and $(b_1, \ldots, b_k) \in \Gamma$ are connected by an (undirected) edge if and only if there is a unique index i for which $a_i \neq b_i$. Note that S is symmetric, so this is an undirected graph.

2.2.1 Eigenvalues of Abelian Cayley Graphs

One nice reason to study Cayley graphs over abelian groups using spectral graph theory is that we can write down all of their adjacency matrix eigenvalues in terms of the characters of the underlying group.

Definition 2.20. *A **character** χ of a group Γ is a group homomorphism $\chi : \Gamma \to \mathbb{C}^*$.*

Example 2.21. If $\Gamma = \mathbb{Z}/N\mathbb{Z}$ for some positive integer N, then for every $r \in \{0, 1, \ldots, N-1\}$, we have a corresponding character χ_r defined by $\chi_r(k) = e^{2\pi i r k/N}$. This is a well-defined map because $e^{2\pi i} = 1$, so $e^{2\pi i r k_1/N} = e^{2\pi i r k_2/N}$ whenever $k_1 \equiv k_2 \pmod N$. Furthermore, χ_r is a homomorphism because $\chi_r(0) = e^0 = 1$ and $\chi_r(k + \ell) = e^{2\pi i r(k+\ell)/N} = \chi_r(k)\chi_r(\ell)$.

Fact 2.22. *[Con, Theorem 3.5, Corollary 4.2] Let Γ be a finite abelian group. Then, the number of distinct characters of Γ is exactly $|\Gamma|$. Furthermore, if χ and ψ are distinct characters of Γ, then*

$$\sum_{x \in \Gamma} \chi(x)\overline{\psi(x)} = 0.$$

Theorem 2.23. *[Tre16, Lemma 16.16] Let $(\Gamma, +)$ be a finite abelian group, and let S be any multiset of elements of Γ. Then, the eigenvalues of the adjacency matrix of $G := Cay(\Gamma, S)$ are (including multiplicity)*

$$\left\{ \sum_{s \in S} \chi(s) : \chi \text{ is a character of } \Gamma \right\}.$$

Proof. For every character χ of Γ, let v_χ be the vector indexed by Γ with $v_\chi(a) = \chi(a)$. Note that χ takes values in $\mathbb{C}^*$, so $v_\chi \neq \vec{0}$. We claim that v_χ is an eigenvector of $A(G)$. Indeed, for any vertex a, we can calculate

$$\begin{aligned} Av_\chi(a) = \sum_{b \in \Gamma} A(a,b)v_\chi(b) &= \sum_{b \in \Gamma} \text{mult}_S(b-a)\chi(b) \\ &= \chi(a) \cdot \sum_{b \in \Gamma} \text{mult}_S(b-a)\chi(b-a) \\ &= v_\chi(a) \cdot \sum_{b \in \Gamma} \text{mult}_S(b)\chi(b), \end{aligned}$$

where in the last line we simply re-indexed the sum. Thus, v_χ is an eigenvector with eigenvalue $\sum_{b\in\Gamma} \text{mult}_S(b)\chi(b) = \sum_{s\in S}\chi(s)$. By Fact 2.22, there are exactly $|\Gamma|$ distinct v_χ's, and they are pairwise orthogonal, which means they are linearly independent. Thus, the v_χ's form a basis of eigenvectors, and their corresponding eigenvalues are exactly the character sums $\sum_{s\in S}\chi(s)$. □

Corollary 2.24. *Let N be a positive integer, let $\Gamma = \mathbb{Z}/N\mathbb{Z}$, and let S be any multiset of elements of Γ. Then, the eigenvalues of $Cay(\Gamma, S)$ are $\{\alpha_r : r = 0, 1, \ldots, N-1\}$, where*

$$\alpha_r := \sum_{s\in S}\chi_r(s) = \sum_{s\in S} e^{2\pi irs/N} = \sum_{s'\in rS} e^{2\pi is'/N} \qquad (*)$$

We call $\alpha_0 = |S|$ the "trivial eigenvalue" and all other eigenvalues "nontrivial".

Proof. By Fact 2.22, Γ should have exactly N characters, and by Example 2.21, we know of N distinct characters already: the characters χ_r for $r \in \{0, 1, \ldots, N-1\}$. Applying Theorem 2.23 yields the desired result. Note that the last two expressions in $(*)$ are equivalent because $e^{2\pi irs/N}$ depends only on the residue of rs modulo N. □

2.2.2 Examples

Example 2.25 (Eigenvalues of the Cycle)**.** Recall that the undirected cycle on N vertices can be described as the Cayley graph $G = Cay(\mathbb{Z}/N\mathbb{Z}, \{1, -1\})$. For this graph, $\alpha_r = e^{2\pi ir/N} + e^{-2\pi ir/N} = 2\cos(2\pi r/N)$, so the eigenvalues of the undirected cycle are $\{2\cos(2\pi r/N) : r = 0, 1, \ldots, N-1\}$. Note that the trivial eigenvalue is $\alpha_0 = 2$, which corresponds to G being 2-regular.

Example 2.26 (Eigenvalues of the Complete Graph)**.** Again, take the group $\Gamma = \mathbb{Z}/N\mathbb{Z}$, and consider the Cayley graph $G := Cay(\Gamma, \Gamma \setminus \{0\})$. This is known as the *complete graph*, because there is an undirected edge of weight 1 between any two distinct vertices. Its eigenvalues are given by the sums

$$\alpha_r = \sum_{s\in\Gamma\setminus 0}\chi_r(s) = \sum_{s=1}^{N-1} e^{2\pi irs/N}.$$

The trivial eigenvalue is $\alpha_0 = |\Gamma \setminus \{0\}| = N - 1$. In addition, for every $r \neq 0$, we have that $e^{2\pi ir/N} \neq 1$, so we can apply the geometric series summation formula to get

$$\alpha_r = \frac{e^{2\pi ir(1)/N} - e^{2\pi ir(N)/N}}{1 - e^{2\pi ir/N}} = \frac{e^{2\pi ir/N} - 1}{1 - e^{2\pi ir/N}} = -1.$$

Thus, the eigenvalues of G are $N - 1$ (with multiplicity 1) and -1 (with multiplicity $N - 1$).

Example 2.27 (Eigenvalues of the Hypercube)**.** In order to apply Theorem 2.23 to the hypercube, we need to understand the characters of the group $\Gamma = (\mathbb{Z}/2\mathbb{Z})^k$. For every vector $v \in \{0, 1\}^k$, let

$\chi_v(x) = (-1)^{\langle v,x\rangle} = (-1)^{\left(\sum_{i=1}^k v_i x_i\right)}$ be a function on Γ (this is well-defined because $(-1)^2 = 1$). We claim that χ_v is a character for every v and that $\chi_v \neq \chi_w$ when $v \neq w$. To see this, note that $\chi_v(0) = (-1)^0 = 1$ for any v and

$$\chi_v(x+y) = (-1)^{\langle v,x+y\rangle} = (-1)^{\langle v,x\rangle + \langle v,y\rangle} = \chi_v(x)\chi_v(y),$$

so χ_v is a character. Furthermore, if $v \neq w$ are two distinct vectors, they must differ in some coordinate: suppose $v_j \neq w_j$. Then, if $e_j = (0, 0, \ldots, 1, \ldots, 0, 0) \in \Gamma$ is the element with a 1 in the jth coordinate only, we can compute $\chi_v(e_j) = (-1)^{\left(\sum_{i=1}^k v_i e_i\right)} = (-1)^{v_j}$. Similarly, $\chi_w(e_j) = (-1)^{w_j}$, so $\chi_v(e_j) \neq \chi_w(e_j)$, as desired.

By Fact 2.22, since we have $|\Gamma| = 2^k$ distinct characters of the form χ_v already, these must be all of the characters. In particular, the eigenvalues of a Cayley graph over Γ are the sums $\left\{\sum_{s\in S} \chi_v(s) : v \in \{0,1\}^k\right\}$. If we apply this to the hypercube, which is the Cayley graph with $S = \{e_1, e_2, \ldots, e_k\}$ as in Example 2.19, we have

$$\sum_{s\in S} \chi_v(s) = \sum_{i=1}^k \chi_v(e_i) = \sum_{i=1}^k (-1)^{v_i} = \sum_{i=1}^k \Big[1 - 2 \cdot \mathbb{1}\{v_i = 1\}\Big] = k - 2 \cdot (\#\text{ of nonzero entries in } v).$$

Since there are exactly $\binom{k}{\ell}$ vectors $v \in \{0,1\}^k$ with ℓ nonzero entries, we conclude that the eigenvalues of the hypercube are k with multiplicity $\binom{k}{0} = 1$, $k-2$ with multiplicity $\binom{k}{1} = k$, $k-4$ with multiplicity $\binom{k}{2}$, …, ending with the eigenvalue of $-k$ with multiplicity $\binom{k}{k} = 1$.

Example 2.28 (Eigenvalues of the Paley Graph)**.** Let p be a prime number such that $p \equiv 1 \bmod 4$. Let $\Gamma = \mathbb{Z}/p\mathbb{Z}$, and let S be the set of nonzero quadratic residues (i.e. $x \in S$ if and only if $x \neq 0$ and there exists y such that $y^2 \equiv x \bmod p$). The Cayley graph $G = Cay(\Gamma, S)$ is known as the *Paley graph* of order p. Since $1 \in S$, G is clearly connected, because the element 1 alone generates Γ.

Lemma 2.29. *The set S has size $\frac{p-1}{2}$ and is closed under multiplication modulo p.*

Proof. For residues a and b modulo p, the condition that $a^2 \equiv b^2 \pmod p$ is equivalent to the condition that $p \mid a^2 - b^2$ by definition. Since p is prime and $a^2 - b^2 = (a-b)(a+b)$, this is equivalent to the statement that either $p \mid a - b \implies a \equiv b \pmod p$ or $p \mid a + b \implies a \equiv -b \pmod p$. Thus, the distinct nonzero quadratic residues are exactly $1^2 = (-1)^2, 2^2 = (-2)^2, \ldots, (\frac{p-1}{2})^2 = (\frac{p+1}{2})^2$. This proves that $|S| = \frac{p-1}{2}$.

For the last part, if $x_1, x_2 \in S$, this means that there exist integers y_1 and y_2 such that $y_1^2 \equiv x_1 \pmod p$ and $y_2^2 \equiv x_2 \pmod p$. Then, $(y_1 y_2)^2 \equiv x_1 x_2 \pmod p$, so $x_1 x_2 \in S$. □

Remark 2.30. When $p \equiv 1 \bmod 4$, it is well-known that -1 is a quadratic residue mod p. Therefore, by the lemma, S contains an element a if and only if it contains $-a$, so G is undirected.

Returning to the eigenvalues of our Paley graph G, by Corollary 2.24, we have $\alpha_0 = |S| = \frac{p-1}{2}$ and the nontrivial eigenvalues $\alpha_r = \sum_{s' \in rS} e^{2\pi i s'/p}$ for $r = 1, 2, \ldots, p-1$. However, by the preceding lemma, S is a subgroup of the multiplicative group of nonzero residues modulo p (since it contains 1 and is closed under multiplication). Therefore, if $r \in S$, then $rS = S$, and if $r \notin S$, then $rS = \{1, 2, \ldots, p-1\} \setminus S$. In particular, the eigenvalues α_r for $r \in S$ are all equal, and the eigenvalues α_r for $r \notin S$ are all equal. These eigenvalues can be calculated directly using quadratic Gauss sums, but there is also a neat way to compute them using a common technique in spectral graph theory: the trace method!

The trace method exploits the fact that we can compute the trace of powers of the adjacency matrix of a graph in two different ways. From linear algebra, we know that the trace of a matrix is the sum of its eigenvalues, so $\mathrm{Tr}(A(G)^k) = \sum_{i=0}^{p-1} \alpha_i^k$. On the other hand, there is a graph theoretic interpretation of the entries of $A(G)^k$.

Claim. For every graph G, the (v, v')-entry of the matrix $A(G)^k$ is the sum of the weights of all directed paths of length k from v to v', where the weight of a path is defined as the product of the weights of its edges.

Proof of Claim. We induct on k. For $k = 1$, the statement follows from the definition of the adjacency matrix, and the fact that a directed path of length 1 is just a single directed edge. For the inductive step, suppose the statement holds for $A(G)^{k-1}$. Then, the (v, v')-entry of $A(G)^k = A(G)A(G)^{k-1}$ is $\sum_{u \in \Gamma} A(G)_{vu}(A(G)^{k-1})_{uv'}$. Note that the directed paths of length k from v to v' are in bijection with pairs of an edge from v to u and a directed path of length $k-1$ from u to v', and the weight of the overall path from v to v' is $w(v, u)$ times the weight of the path from u to v'. Therefore, $A(G)_{vu}(A(G)^{k-1})_{uv'}$ computes the sum of the weights of paths of length k from v to v' where the first step goes from v to u, and summing this quantity over u gives the total sum of weights of paths of length k from v to v', as desired. □

Therefore, the trace of $A(G)^k$ is the sum of the weights of all directed paths of length k from v to v as v ranges over the vertices of G.

Returning one last time to our example of the Paley graph, we see that $w(v, v) = 0$ for all vertices v, so $\mathrm{Tr}(A(G)) = 0 = \sum_{i=0}^{p-1} \alpha_i^1$. If we let α' be the value of α_r for $r \in S$ and let α'' be the value of α_r for $r \notin S$, then this means that $0 = \frac{p-1}{2} + \frac{p-1}{2} \cdot \alpha' + \frac{p-1}{2} \cdot \alpha''$. In other words, $\alpha' + \alpha'' = -1$.

Now, we apply the trace method again with $k = 2$! In an undirected graph with no self-loops,

a directed path of length 2 from a vertex to itself must be obtained by following a single edge away and back again. Our edges all have weight 1, so these paths also have weight 1. Thus, $\mathrm{Tr}(A(G)^2) = \sum_{i=0}^{p-1} \alpha_i^2$ will be 2 times the number of edges in G, where the factor of 2 comes from choosing an endpoint to start at. Since G is regular, it has $\frac{1}{2} \cdot p \cdot \frac{p-1}{2}$ edges, so

$$\sum_{i=0}^{p-1} \alpha_i^2 = \left(\frac{p-1}{2}\right)^2 + \frac{p-1}{2} \cdot (\alpha')^2 + \frac{p-1}{2} \cdot (\alpha'')^2 = \frac{p(p-1)}{2}.$$

Simplifying this expression yields $(\alpha')^2 + (\alpha'')^2 = \frac{p+1}{2}$, at which point we can plug in $\alpha'' = -1 - \alpha'$ and solve the resulting quadratic $2(\alpha')^2 + 2\alpha' - \frac{p-1}{2} = 0$. We find that α' and α'' are $\frac{-1+\sqrt{p}}{2}$ and $\frac{-1-\sqrt{p}}{2}$, in some order.

In conclusion, the eigenvalues of the Paley graph of order p are $\frac{p-1}{2}$ with multiplicity 1 and $\frac{-1\pm\sqrt{p}}{2}$, each with multiplicity $\frac{p-1}{2}$.

3 General Bounds on Eigenvalue Multiplicity

As discussed in the introduction, several recent works have considered the multiplicities of the eigenvalues of particular graphs or classes of graphs. In this section, we review some of the known results about eigenvalue multiplicity in abelian Cayley graphs.

3.1 Upper Bounds

The main result of [DJRVZ24] is a faster approximation algorithm for the Sparsest Cut problem on abelian Cayley graphs. It relies on bounding the number of "small" Laplacian eigenvalues of such graphs. While their techniques work in greater generality, we will only concern ourselves with the case of bounding the multiplicity of the second-smallest Laplacian eigenvalue.

Theorem 3.1. *[DJRVZ24, Theorem 1.10] Let $G = Cay(\Gamma, S)$ be a connected and undirected abelian Cayley graph of degree $d := |S|$. If λ_2 is the smallest nonzero eigenvalue of $L(G)$, then the multiplicity of λ_2 as an eigenvalue of $L(G)$ is at most $800 \cdot 2^{30d}$.*

This theorem implies that if $d \ll \log|\Gamma|$ for a family of abelian Cayley graphs, then the multiplicity of λ_2 will be $o(|\Gamma|)$. On the other hand, when $d \geq (\log|\Gamma|)/30$, the theorem is trivial because $800 \cdot 2^{30d} \geq |\Gamma|$, which is the total number of eigenvalues of the graph counting multiplicity.

The main idea of the proof is to relate the eigenvalues of $L(G)$ to a quantity called the *collision probability*, and then to use the structure of abelian Cayley graphs to obtain bounds on expressions involving the collision probability. We follow Section 5 of [DJRVZ24]. I chose to reproduce this proof in full because it demonstrates a very combinatorial approach to an eigenvalue multiplicity question; this will contrast with our later algebraic approaches that exploit the full group structure.

Definition 3.2. *The **lazy random walk** on a weighted directed graph G with weight function w is defined by the following process: at each step, if the walk is currently at a vertex v, then*

1. *with probability $\frac{1}{2}$, remain at vertex v.*
2. *with probability $\frac{1}{2} \cdot \frac{w(v,v')}{d_{out}(v)}$, move to vertex v'.*

Remark 3.3. If $G = Cay(\Gamma, S)$ is a Cayley graph, a step of the lazy random walk on G corresponds to with probability $\frac{1}{2}$, staying at the same vertex, and with probability $\frac{1}{2} \cdot \frac{\text{mult}_S(s)}{|S|}$ moving from vertex x to vertex xs. From this description, one can see that for Cayley graphs, the probability of a lazy random walk starting at x and ending at y after t steps is the same as the probability of a lazy random walk starting at x' and ending at $x'(x^{-1}y)$ after t steps.

Definition 3.4. *For an undirected Cayley graph $G = Cay(\Gamma, S)$, the **t-step lazy collision probability** (denoted CP_t) is the probability that a lazy random walk starting at some vertex x of length $2t$ also ends at vertex x (by the previous remark, this is independent of x).*

Example 3.5. Let $\Gamma = \mathbb{Z}/5\mathbb{Z}$ and $S = \{1, -1\}$, so that $G = Cay(\Gamma, S)$ is the undirected cycle on 5 vertices. Consider CP_2 for this graph. At any vertex $x \in \Gamma$, a single step of the lazy random walk remains at x with probability $\frac{1}{2}$, moves to $x+1$ with probability $\frac{1}{4}$, and moves to $x-1$ with probability $\frac{1}{4}$. Let these steps be denoted X, R, and L, respectively. A walk of length 4 returns to its starting vertex if and only if the walk consists of 4 X moves; 2 X, 1 L, and 1 R move; or 2 L and 2 R moves. Thus,

$$\text{CP}_2 = \left(\frac{1}{2}\right)^4 + \binom{4}{2,1,1} \cdot \left(\frac{1}{2}\right)^2 \cdot \left(\frac{1}{4}\right)^2 + \binom{4}{2,2} \cdot \left(\frac{1}{4}\right)^4 = \frac{1}{16} + \frac{12}{64} + \frac{6}{256} = \frac{35}{128}.$$

This definition of collision probability may seem unmotivated, but we will now see that CP_t can be calculated exactly using the eigenvalues of G.

Lemma 3.6 ([DJRVZ24, Lemma 5.2])**.** *Let $G = Cay(\Gamma, S)$ be an undirected Cayley graph, and define $n = |\Gamma|$ and $d = |S|$. Recall that we write $\lambda_1 \le \lambda_2 \le \cdots \le \lambda_n$ for the eigenvalues of $L(G)$. Then,*

$$\text{CP}_t = \frac{1}{n} \sum_{i=1}^{n} \left(1 - \frac{\lambda_i}{2d}\right)^{2t}.$$

Proof. This is another application of the trace method, which we first introduced to calculate the eigenvalues of the Paley graph. Consider the $n \times n$ matrix $\tilde{W} = \frac{1}{2}I + \frac{1}{2d}A(G)$, where I is the identity matrix. Note that if e_x is the vector indexed by Γ with a 1 in the x-entry and a 0 everywhere else, then $\tilde{W}e_x$ is exactly the probability distribution of where the lazy random walk ends up after one step from x. In particular, the x-entry of $(\tilde{W})^{2t}e_x$ is the probability that a lazy random walk starting at x ends at x after $2t$ steps, which is the definition of CP_t. Since this quantity is independent of x, we conclude that every diagonal entry of $(\tilde{W})^{2t}$ equals CP_t.

On the other hand, by basic linear algebra, if $\alpha_1, \ldots, \alpha_n$ are the eigenvalues of $A(G)$, then the eigenvalues of $\tilde{W}$ are $\frac{1}{2} + \frac{1}{2d}\alpha_i = 1 - \frac{1}{2d}(d - \alpha_i) = 1 - \frac{\lambda_i}{2d}$, where we apply Observation 2.8 in the last step. This implies that the eigenvalues of $(\tilde{W})^{2t}$ are $\left(1 - \frac{\lambda_i}{2d}\right)^{2t}$. In particular,

$$\text{tr}\left((\tilde{W})^{2t}\right) = \sum_{i=1}^{n} \left(1 - \frac{\lambda_i}{2d}\right)^{2t},$$

but the trace of this matrix is also the sum of its diagonal entries, which is just $n \cdot \text{CP}_t$. Dividing by n yields the desired result. □

Suppose now that $G = Cay(\Gamma, S)$ is connected and undirected. Denote the multiplicity of λ_2 by m. To prove Theorem 3.1, we need to relate m to a quantity involving the collision probability, and to also somehow handle this quantity directly using the abelian Cayley graph structure. The specific quantity we'll work with is $\frac{\text{CP}_t}{\text{CP}_{2t}}$ for the appropriate choice of t.

Lemma 3.7 ([DJRVZ24, Lemma 5.3]). *For $t = \lfloor d\ln(m)/(3\lambda_2)\rfloor$, the following inequality holds:*

$$\frac{\mathrm{CP}_t}{\mathrm{CP}_{2t}} \geq \frac{1}{2e^{3/2}} \cdot m^{1/3}.$$

Proof. Applying Lemma 3.6 and using the fact that $\lambda_i \in [0, 2d]$, we already have that $\mathrm{CP}_t = \frac{1}{n}\sum_{i=1}^{n}\left(1-\frac{\lambda_i}{2d}\right)^{2t} \geq \frac{1}{n}\cdot m\left(1-\frac{\lambda_2}{2d}\right)^{2t}$, where we pick from the sum only the terms with $\lambda_i = \lambda_2$ (all other terms are non-negative). Additionally,

$$\mathrm{CP}_t = \frac{1}{n}\sum_{i=1}^{n}\left(1-\frac{\lambda_i}{2d}\right)^{2t} \geq \frac{1}{n}\sum_{i=2}^{n}\left(1-\frac{\lambda_i}{2d}\right)^{-2t+4t} \geq \frac{1}{n}\cdot\left(1-\frac{\lambda_2}{2d}\right)^{-2t}\cdot\sum_{i=2}^{n}\left(1-\frac{\lambda_i}{2d}\right)^{4t}.$$

Furthermore, applying Lemma 3.6 to CP_{2t} yields

$$\mathrm{CP}_{2t} = \frac{1}{n}\sum_{i=1}^{n}\left(1-\frac{\lambda_i}{2d}\right)^{4t} = \frac{1}{n}+\frac{1}{n}\sum_{i=2}^{n}\left(1-\frac{\lambda_i}{2d}\right)^{4t} \leq \frac{1}{n}+\left(1-\frac{\lambda_2}{2d}\right)^{2t}\cdot\mathrm{CP}_t.$$

Now, we apply a little trick: for all positive real numbers c and d, either $c+d \leq 2c$ or $c+d \leq 2d$, so either $\frac{1}{c+d} \geq \frac{1}{2c}$ or $\frac{1}{c+d} \geq \frac{1}{2d}$. Therefore, either

$$\frac{\mathrm{CP}_t}{\mathrm{CP}_{2t}} \geq \frac{\mathrm{CP}_t}{2/n} \geq \frac{1}{2}\cdot m\cdot\left(1-\frac{\lambda_2}{2d}\right)^{2t}$$

or

$$\frac{\mathrm{CP}_t}{\mathrm{CP}_{2t}} \geq \frac{\mathrm{CP}_t}{2\left(1-\frac{\lambda_2}{2d}\right)^{2t}\cdot\mathrm{CP}_t} = \frac{1}{2}\cdot\left(1-\frac{\lambda_2}{2d}\right)^{-2t}.$$

To finish, we simply plug in the choice of $t = \lfloor d\ln(m)/(3\lambda_2)\rfloor$. On the one hand, $1-x \leq e^{-x}$ for all $x \in [0,1]$, so $(1-\frac{\lambda_2}{2d})^{-2t} \geq e^{t\lambda_2/d} \geq m^{1/3}$. On the other hand, $1-x \geq e^{-x/(1-x)}$ for all $x \in [0,1]$, so

$$\left(1-\frac{\lambda_2}{2d}\right)^{2t} \geq e^{-t\lambda_2/(2d-\lambda_2)} \geq e^{-(d\ln(m)/(3\lambda_2)+1)\cdot\lambda_2/(2d-\lambda_2)} \geq m^{-d/(3(2d-\lambda_2))}\cdot e^{-\lambda_2/(2d-\lambda_2)}.$$

It is a fact that $\lambda_2 \leq \frac{3}{2}d$ for every regular graph of degree at least 2, so we have $(1-\frac{\lambda_2}{2d})^{2t} \geq m^{-2/3}\cdot e^{-3/2}$. Taking the worse of the two upper bounds above now yields $\frac{\mathrm{CP}_t}{\mathrm{CP}_{2t}} \geq \frac{1}{2e^{3/2}}\cdot m^{1/3}$, as desired. □

Lemma 3.8 ([DJRVZ24, Lemma 5.4]). *If $G = Cay(\Gamma, S)$ is a connected and undirected Cayley graph with $(\Gamma, +)$ an abelian group and $|S| = d$, then for every $t \geq 1$, we have the following bound*

on the ratio of the lazy collision probabilities:

$$\frac{\mathrm{CP}_t}{\mathrm{CP}_{2t}} \leq (2e)^{4d}.$$

Proof. Returning to our original definition of CP_t, we have that CP_t is the probability that a lazy random walk of length $2t$ on G ends at the same place it started. In an abelian Cayley graph, a step along an edge corresponding to the generator $s \in S$ is equivalent to adding s to the current vertex, so to move from x to x is just to select $2t$ "moves" that sum to 0, where a move can either be to stand still (add 0) or move according to a generator (add s). To make notation easier, we let S' be the multiset comprised of the union of two identical copies of S and $2d$ copies of the element 0. In technical language, $\mathrm{mult}_{S'}(0) = 2d + 2\,\mathrm{mult}_S(0)$ and $\mathrm{mult}_{S'}(s) = 2\,\mathrm{mult}_S(s)$ for every $s \neq 0$.

Observe that moving from x to $x + s'$ where s' is selected from S' with probability proportional to its multiplicity is identical to taking a lazy random walk in G. Note that $|S'| = 4d$, and let $[4d] := \{1, 2, \ldots, 4d\}$. Now, we view $S' = \{s_1, \ldots, s_{4d}\}$ as a set with repeated elements so that we can write

$$\mathrm{CP}_t = \frac{1}{(4d)^{2t}} \sum_{(i_1,\ldots,i_{2t})\in[4d]^{2t}} \mathbb{1}\left[s_{i_1} + s_{i_2} + \cdots + s_{i_{2t}} = 0\right].$$

Instead of looking at tuples of indices, we now shift to looking only at how many times each s_i appears in the sum. If $c_1, \ldots, c_{4d}$ are some non-negative integers such that $\sum_{i=1}^{4d} c_i = 2t$ and $\sum_{i=1}^{4d} c_i s_i = 0$, then any tuple with index i appearing c_i times will contribute 1 to the sum above. In other words,

$$\mathrm{CP}_t = \frac{1}{(4d)^{2t}} \sum_{\substack{(c_1,\ldots,c_{4d}) \\ \sum c_i = 2t \\ \sum c_i s_i = 0}} \binom{2t}{c_1, c_2, \ldots, c_{4d}}.$$

Since S was symmetric to begin with, S' is still symmetric, and by construction, we can pair up the $4d$ elements of S' including repetitions such that every element is paired with its inverse (this is the reason for creating 2 identical copies of S, so that we can pair up the elements which are their own inverses). Without loss of generality, suppose $s_1 = -s_2, s_3 = -s_4, \cdots$, and $s_{4d-1} = -s_{4d}$. Now, let μ be the vector of length $4d$ defined by $\mu_i = \lfloor t/2d \rfloor + 1$ if $i \leq 2t - 4d\lfloor t/2d \rfloor$ and $\mu_i = \lfloor t/2d \rfloor$ otherwise. Note that $\sum_{i=1}^{4d} \mu_i = 4d\lfloor t/2d \rfloor + 2t - 4d\lfloor t/2d \rfloor = 2t$. Furthermore, $\mu_1 = \mu_2, \mu_3 = \mu_4, \cdots$ because $2t - 4d\lfloor t/2d \rfloor$ is even, so the point at which the value of μ_i changes is between two pairs. Thus, $\sum_{i=1}^{4d} \mu_i s_i = 0$, since within each pair the sum is canceled out.

Creating this specific vector μ allows us to relate CP_t to CP_{2t}. Given a vector $c = (c_1, \ldots, c_{4d})$ satisfying $\sum c_i = 2t$ and $\sum c_i s_i = 0$, we see that the vector $c + \mu := (c_1 + \mu_1, \ldots, c_{4d} + \mu_{4d})$ satisfies

$\sum(c+\mu)_i = 4t$ and $\sum(c+\mu)_i s_i = 0$. This means that

$$\textsc{cp}_{2t} \leq \frac{1}{(4d)^{4t}} \sum_{\substack{(c_1,\ldots,c_{4d}) \\ \sum c_i = 2t \\ \sum c_i s_i = 0}} \binom{4t}{c_1+\mu_1, c_2+\mu_2, \ldots, c_{4d}+\mu_{4d}}.$$

In particular, it now certainly suffices to show that the ratio of every term in $\textsc{cp}_t$ to its corresponding term in $\textsc{cp}_{2t}$ is bounded; more concretely, we want to show that

$$\frac{1}{(4d)^{2t}} \binom{2t}{c_1, c_2, \ldots, c_{4d}} \leq (2e)^{4d} \cdot \frac{1}{(4d)^{4t}} \binom{4t}{c_1+\mu_1, c_2+\mu_2, \ldots, c_{4d}+\mu_{4d}}$$

for every vector c with $\sum c_i = 2t$ (this is overkill: we only actually need it for those vectors that additionally satisfy $\sum c_i s_i = 0$). After rearranging, the expression that we want to show becomes

$$\frac{(2t)!(c_1+\mu_1)! \cdots (c_{4d}+\mu_{4d})! \cdot (4d)^{2t}}{(4t)! c_1! \cdots c_{4d}!} \leq (2e)^{4d}.$$

Claim. Let $f(c)$ denote the quantity on the left hand side of the inequality above. Then, over the space of all c with $\sum c_i = 2t$, the function $f(c)$ is maximized at $c = \mu$.

Proof of Claim. Assuming $c \neq \mu$, we show that there exists $c' \neq c$ such that $f(c') > f(c)$. If $c \neq \mu$, then since $\sum c_i = \sum \mu_i$, there must be a coordinate j for which $c_j > \mu_j$ and a coordinate k for which $c_k < \mu_k$. Let c' be the vector obtained by taking c, subtracting 1 at the jth coordinate, and adding 1 at the kth coordinate. Clearly, $\sum c'_i$ still equals $2t$, and we have

$$\frac{f(c')}{f(c)} = \frac{c_k + 1 + \mu_k}{c_k + 1} \cdot \frac{c_j}{c_j + \mu_j}.$$

Now, $c_k < \mu_k$ implies $c_k + 1 \leq \mu_k$, so $\frac{c_k+1+\mu_k}{c_k+1} \geq 2$. On the other hand, $c_j > \mu_j$, so $c_j + \mu_j < 2c_j$, which means $\frac{c_j}{c_j+\mu_j} > \frac{1}{2}$. Thus, $f(c')/f(c) > 2 \cdot \frac{1}{2} = 1$, so $f(c') > f(c)$, as desired. □

Finally, by the claim, it suffices to show that $f(\mu) \leq (2e)^{4d}$. By Stirling's approximation, which

says that $2\sqrt{k}(\frac{k}{e})^k \le k! \le 2\sqrt{2k}(\frac{k}{e})^k$ for all positive integers k, we have

$$\begin{aligned} f(\mu) &= \frac{(2t)!(2\mu_1)!\cdots(2\mu_{4d})!\cdot(4d)^{2t}}{(4t)!\mu_1!\cdots\mu_{4d}!} \\ &\le \frac{\sqrt{4t}(2t/e)^{2t}\cdot\sqrt{4\mu_1}(2\mu_1/e)^{2\mu_1}\cdots\sqrt{4\mu_{4d}}(2\mu_{4d}/e)^{2\mu_{4d}}\cdot(4d)^{2t}}{\sqrt{4t}(4t/e)^{4t}\cdot\sqrt{\mu_1}(\mu_1/e)^{\mu_1}\cdots\sqrt{\mu_{4d}}(\mu_{4d}/e)^{\mu_{4d}}} \\ &= 2^{4d}\cdot(4d)^{2t}\cdot\frac{(2t)^{2t}(2\mu_1)^{2\mu_1}\cdots(2\mu_{4d})^{2\mu_{4d}}}{(4t)^{4t}\mu_1^{\mu_1}\cdots\mu_{4d}^{\mu_{4d}}} \\ &= 2^{4d}\cdot(4d)^{2t}\cdot\frac{\mu_1^{\mu_1}\cdots\mu_{4d}^{\mu_{4d}}}{(2t)^{2t}} \\ &\le 2^{4d}\cdot\left(\frac{2d}{t}\right)^{2t}\cdot\left(\frac{t}{2d}+1\right)^{2t}, \end{aligned}$$

where we repeatedly make use of the fact that $\sum\mu_i = 2t$. In the last step, we use the definition of μ to bound each μ_i by $\frac{t}{2d}+1$. Now, the last two terms combine to yield $(1+2d/t)^{2t}$, which using the inequality $1+x\le e^x$ can be bounded by e^{4d}. This completes the proof. □

Proof of Theorem 3.1. Combining Lemma 3.7 and Lemma 3.8, we see that for a connected and undirected abelian Cayley graph, we have $\frac{1}{2e^{3/2}}\cdot m^{1/3} \le \frac{\mathrm{CP}_t}{\mathrm{CP}_{2t}} \le (2e)^{4d}$, where $t = \lfloor d\ln(m)/(3\lambda_2)\rfloor$, so

$$m \le 8e^{9/2}\cdot(2e)^{12d} \le 800\cdot 2^{30d}.$$

□

3.2 Lower Bound Examples

In the setting of general abelian Cayley graphs, it turns out that up to the constant in the exponent, the upper bound on the multiplicity of the second-smallest Laplacian eigenvalue is tight. In other words, there exist abelian Cayley graphs in which the multiplicity of the second-smallest Laplacian eigenvalue is exponential in the degree of the graph, a statement which is made precise in the following proposition:

Proposition 3.9. *[DJRVZ24, Proposition 7.2] There exists a constant $C > 0$ and an infinite family of abelian groups Γ_i and connected undirected Cayley graphs $G_i = Cay(\Gamma_i, S_i)$ such that $|\Gamma_i|$ and $|S_i|$ are unbounded, and the multiplicity of the second-smallest eigenvalue of $L(G_i)$ is at least $2^{C|S_i|}$ for every i.*

On the surface, when combined with Theorem 3.1, this proposition appears to suggest that further improvements to upper bounds on the multiplicity of the second-smallest Laplacian eigenvalue can only come in the form of whittling down the constant in the exponent. However, as far as I am aware, the only families of graphs known to achieve the behavior described in Proposition 3.9 use

$\Gamma_i = (\mathbb{Z}/2\mathbb{Z})^{k_i}$ with appropriately chosen k_i (going to infinity) and S_i. In some sense, $(\mathbb{Z}/2\mathbb{Z})^k$ is the farthest a finite abelian group can be from being cyclic, so we still might hope for a better upper bound in the case of Cayley graphs over cyclic groups. The next proposition, however, suggests we shouldn't be overly optimistic.

Proposition 3.10. *Let q be an odd number, and let $n = 3q$ and $\Gamma = \mathbb{Z}/n\mathbb{Z}$. Let $S = \{1, q+1, 2q+1, -1, q-1, 2q-1\} \subset \Gamma$. Then, one of the eigenvalues of $G = Cay(\Gamma, S)$ has multiplicity $2q = 2n/3$.*

Proof. By Corollary 2.24, the eigenvalues of G are given by $\alpha_r = \zeta_n^r + \zeta_n^{r+rq} + \zeta_n^{r+2rq} + \zeta_n^{-r} + \zeta_n^{-r+rq} + \zeta_n^{-r+2rq}$. When r is divisible by 3, we get $\zeta_n^{rq} = 1$, so the expression simplifies to $\alpha_r = 3\zeta_n^r + 3\zeta_n^{-r} = 6\cos(\frac{2\pi r}{n})$. However, when r is not divisible by 3, we know that $\zeta_n^{rq} \neq 1$, so

$$\alpha_r = \zeta_n^r \cdot \frac{1-\zeta_n^{3rq}}{1-\zeta_n^{rq}} + \zeta_n^{-r} \cdot \frac{1-\zeta_n^{3rq}}{1-\zeta_n^{rq}} = 0.$$

Therefore, the eigenvalue 0 has multiplicity $2n/3$, coming from all choices of r which are not divisible by 3 (since n is odd, $\cos(\frac{2\pi r}{n})$ can never equal 0). □

Clearly, for large q, the eigenvalue 0 is not the second-largest adjacency matrix eigenvalue, but this example still demonstrates that we cannot get a sublinear bound on all eigenvalue multiplicities when working over cyclic groups – this is a constant-degree family of graphs with eigenvalue multiplicity linear in n. Nevertheless, this reinforces the idea that looking at the multiplicity of the second-largest eigenvalue is the "correct" question.

The remainder of this write-up works to improve the exponential dependence on d for the multiplicity of the second-largest adjacency matrix eigenvalue for Cayley graphs over cyclic groups. In light of Corollary 2.24, analyzing eigenvalue multiplicities becomes a problem about sums of roots of unity being equal. This observation allows us to apply facts about cyclotomic fields and prove much stronger upper bounds.

4 Cyclic Groups of Prime Order

We begin with the simplest case, algebraically speaking. In this section, we will prove Theorem 1.1, the main theorem of exact eigenvalue multiplicities for Cayley graphs over $\mathbb{Z}/p\mathbb{Z}$.

Theorem 1.1. *Let p be a prime, and let $\Gamma = \mathbb{Z}/p\mathbb{Z}$. Given a non-empty multiset S of nonzero elements of Γ, consider the Cayley graph $G = Cay(\Gamma, S)$. We consider the natural identification of $\Gamma \setminus \{0\}$ with $\mathbb{F}_p^*$, the multiplicative group modulo p. If H is the largest subgroup of $\mathbb{F}_p^*$ such that S is a union of cosets of H[2], then the multiplicity of every eigenvalue except for λ_1 is exactly $|H|$.*

Proof. Fix p and a Cayley graph $G = Cay(\Gamma, S)$ in the setting of Theorem 1.1. Since S is non-empty, it must contain a nonzero element of Γ, which generates Γ since p is prime. Therefore, G is connected, and the multiplicity of λ_1 is exactly 1. By Corollary 2.24, the eigenvalues of G other than $\alpha_0 = d$ are $\alpha_r = \sum_{s\in S} \zeta_p^{rs}$ for $1 \le r \le p-1$. Since $\zeta_p^p = 1$, we can reduce all exponents modulo p, so $\alpha_r = f_r(\zeta_p)$ where $f_r(x) = \sum_{s\in S} x^{rs \bmod p}$ is a polynomial of degree at most $p-1$. Thus, $\alpha_{r_1} = \alpha_{r_2}$ if and only if $f_{r_1}(\zeta_p) = f_{r_2}(\zeta_p) \iff (f_{r_1} - f_{r_2})(\zeta_p) = 0$.

By Fact A.1, if ζ_p is a root of $f_{r_1} - f_{r_2}$, then so is ζ_p^k for all $1 \le k \le p-1$. Furthermore, $f_r(1) = |S|$ for every r, so 1 is also a root of $f_{r_1} - f_{r_2}$. Therefore, if $(f_{r_1} - f_{r_2})(\zeta_p) = 0$, then $f_{r_1} - f_{r_2}$ has at least p distinct roots, but it is a polynomial of degree at most $p-1$. This means it must be the identically 0 polynomial! In other words, we now know that $\alpha_{r_1} = \alpha_{r_2}$ if and only if $f_{r_1} \equiv f_{r_2}$, which happens if and only if $r_1 S = r_2 S$ as multisets. From this it is already evident that the multiplicities of all nontrivial eigenvalues are equal, and the precise conclusion follows from a closer look at the condition $r_1 S = r_2 S$.

Since multiplication by an element of $\mathbb{F}_p^*$ induces a bijection on Γ, we see that $r_1 S = r_2 S$ if and only if $S = r_1^{-1} r_2 S$. Let $H = \{r : S = rS\} \subseteq \mathbb{F}_p^*$. Observe that H is a subgroup of $\mathbb{F}_p^*$, since $1 \in H$, $t \in H \implies S = tS \implies t^{-1}S = S \implies t^{-1} \in H$, and

$$t, t' \in H \implies S = tS = t(t'S) = (tt')S \implies tt' \in H.$$

Note that the multiplicity of α_{r_1} is the number of r_2's such that $r_1^{-1} r_2 \in H$, which is simply $|r_1 H| = |H|$. Furthermore, if a and b are in the same coset of H, then there exists $t \in H$ such that $b = at$, and then $\text{mult}_S(b) = \text{mult}_{tS}(b) = \text{mult}_S(a)$, which means that the multiplicities of elements in S are constant within cosets of H (in other words, S is a union of cosets of H).

Finally, we claim that H is the largest subgroup of $\mathbb{F}_p^*$ with this property. If S is the union of

[2] I say that S is a union of cosets of H if for every coset aH, the multiplicity in S of every element of aH is the same.

cosets of a subgroup H', then for every $t \in H'$ and every $a \in \Gamma$, we have $\text{mult}_S(a) = \text{mult}_S(at^{-1}) = \text{mult}_{tS}(a)$, so $t \in H$. Thus, $H' \subseteq H$, which implies that H is the largest subgroup for which S can be written as a union of cosets of that subgroup. □

Corollary 4.1. *In the setting of Theorem 1.1, the multiplicity of every eigenvalue of G is at most d.*

Proof. Since G is connected, the multiplicity of λ_1 is $1 \leq d$. For the other eigenvalues, note that when S is a union of cosets of a subgroup $H \subset \mathbb{F}_p^*$, we must have that $|H|$ divides $|S|$. In particular, the largest subgroup H for which S is a union of cosets of a subgroup of H cannot have cardinality exceeding $|S| = d$. □

Three key features of Theorem 1.1 warrant some commentary. First, this theorem is a huge improvement over Theorem 3.1, moving from an exponential dependence on d to a linear dependence. In particular, Theorem 3.1 yielded no information when $d \geq (\log|\Gamma|)/30$, whereas Theorem 1.1 remains nontrivial for *all* degrees.

Second, the theorem simultaneously bounds the multiplicity of *every* eigenvalue by the degree, not just the second-largest adjacency matrix eigenvalue. This is surprising in light of Proposition 3.10, where we saw that such a general statement for every eigenvalue multiplicity could never apply to Cayley graphs over cyclic groups of order pq where p and q are distinct primes.

Finally, the theorem actually characterizes the eigenvalue multiplicities in terms of the structure of the generating set S. For example, it tells us exactly how to construct examples which meet the bound of d tightly, many of which we've already seen!

Example 4.2 (Graphs with Eigenvalue Multiplicity d)**.** The undirected cycle (Example 2.25), the Paley graph of order p when $p \equiv 1 \bmod 4$ (Example 2.28), and the complete graph (Example 2.26) are all examples of this phenomenon.

- The nontrivial eigenvalues of the undirected cycle on p vertices were $\{\alpha_r = 2\cos(2\pi r/p) : r = 1, \ldots, p-1\}$, so we can see that every eigenvalue has multiplicity 2: $\alpha_r = \alpha_{p-r}$. This reflects the fact that $S = \{1, -1\}$ is a subgroup of $\mathbb{F}_p^*$.
- The nontrivial eigenvalues of the Paley graph occurred with multiplicity $\frac{p-1}{2}$ each. This reflects the fact that $S = \{\text{quadratic residues modulo } p\}$ is a subgroup of $\mathbb{F}_p^*$ of order $\frac{p-1}{2}$, which was the statement of Lemma 2.29.
- The nontrivial eigenvalues of the complete graph are all equal to -1, so the multiplicity is $p-1$, which reflects the fact that $S = (\mathbb{Z}/p\mathbb{Z}) \setminus \{0\} = \mathbb{F}_p^*$ is a subgroup of $\mathbb{F}_p^*$ of order $p-1$.

5 Squares of Primes

After understanding the eigenvalue multiplicities of Cayley graphs over $\mathbb{Z}/p\mathbb{Z}$, it makes sense to consider next the case of $\mathbb{Z}/(p^2)\mathbb{Z}$. In this section, we will prove Theorem 1.2, restated here for convenience.

Theorem 1.2. *Let p be an odd prime, and let $n = p^2$ and $\Gamma = \mathbb{Z}/n\mathbb{Z}$. Given a symmetric multiset S of nonzero elements of Γ, consider the Cayley graph $G = Cay(\Gamma, S)$. If G is connected, then the multiplicity of $\lambda_2(G)$ is at most $|S| = d$.*

5.1 Setup

Let p, n, Γ, S, and $G = Cay(\Gamma, S)$ a connected graph be fixed in the setting of the theorem. Recall that the eigenvalues of G are given by Corollary 2.24 as

$$\alpha_r = \sum_{s \in S} \zeta_n^{rs} = \sum_{s' \in rS} \zeta_n^{s'}. \tag{$*$}$$

In the case of $\mathbb{Z}/p\mathbb{Z}$, this expression allowed us to show that $\alpha_{r_1} = \alpha_{r_2} \iff r_1 S = r_2 S$, but this will no longer be true due to additional linear dependencies among the powers of ζ_n. Instead, we present the following lemma, which describes a necessary and sufficient condition on the multisets $r_1 S$ and $r_2 S$ to obtain $\alpha_{r_1} = \alpha_{r_2}$.

Lemma 5.1. *Let $r_1, r_2 \in \{0, 1, \ldots, n-1\}$ be any two elements. Then, the corresponding eigenvalues α_{r_1} and α_{r_2} are equal if and only if there exist offsets $v_0, v_1, \ldots, v_{p-1} \in \mathbb{Z}$ such that for all $m \in \Gamma$,*

$$\text{mult}_{r_1 S}(m) = \text{mult}_{r_2 S}(m) + v_{m \bmod p}.$$

Remark 5.2. If they exist, these values will satisfy $\sum_{t=0}^{p-1} v_t = 0$ because $|r_1 S| = |r_2 S|$.

The proof of Lemma 5.1 relies on a key fact about elements that form a basis for $\mathbb{Q}[\zeta_n]$ over $\mathbb{Q}$.

Fact 5.3. *Let $M' \subset \Gamma$ be any subset which contains exactly one representative of every residue modulo p, and let $M = \Gamma \setminus M'$. Then, the set*

$$B := \left\{ \zeta_n^m : m \in M \right\}$$

forms a basis for $\mathbb{Q}[\zeta_n]$ as a $\mathbb{Q}$-vector space. Furthermore, the following holds for every $m' \in M'$:

$$\zeta_n^{m'} = -\sum_{j=1}^{p-1} \zeta_n^{m' + jp}, \tag{1}$$

Proof of Fact 5.3. Since $|B| = |M| = p^2 - p$ and $[\mathbb{Q}[\zeta_n] : \mathbb{Q}] = \varphi(p^2) = p^2 - p$ by Fact A.1, it suffices to show that B generates the whole vector space $\mathbb{Q}[\zeta_n]$. We see that for every residue $t \in \{0, 1, \ldots, p-1\}$, we have the identity

$$\zeta_n^t + \zeta_n^{t+p} + \cdots + \zeta_n^{t+(p-1)p} = \zeta_n^t \cdot (1 + \zeta_p + \zeta_p^2 + \cdots + \zeta_p^{p-1}) = 0.$$

Thus, since M' contains exactly one representative of every residue modulo p, for every $m' \in M'$, we can express

$$\zeta_n^{m'} = -\sum_{j=1}^{p-1} \zeta_n^{m'+jp},$$

where $m' + jp \in M$ for each j in the summation. In particular, every $\zeta_n^{m'}$ can be expressed as a $\mathbb{Q}$-linear combination of elements from B. Thus, every ζ_n^k for $0 \le k < n$ is in the span of B, but these elements generate all of $\mathbb{Q}[\zeta_n]$, so B spans $\mathbb{Q}[\zeta_n]$, as desired. □

Proof of Lemma 5.1. For the if direction, suppose $v_0, \ldots, v_{p-1}$ are integers with the property that $\text{mult}_{r_1 S}(m) = \text{mult}_{r_2 S}(m) + v_{m \bmod p}$ for all $m \in \Gamma$. Then, by $(*)$, we see that

$$\begin{aligned}
\alpha_{r_1} &= \sum_{m \in \Gamma} \text{mult}_{r_1 S}(m) \cdot \zeta_n^m \\
&= \sum_{m \in \Gamma} (\text{mult}_{r_2 S}(m) + v_{m \bmod p}) \cdot \zeta_n^m \\
&= \alpha_{r_2} + \sum_{m \in \Gamma} v_{m \bmod p} \cdot \zeta_n^m \\
&= \alpha_{r_2} + \sum_{t=0}^{p-1} v_t \cdot \left(\zeta_n^t + \zeta_n^{t+p} + \cdots + \zeta_n^{t+(p-1)p} \right) = \alpha_{r_2},
\end{aligned}$$

as desired.

For the only if direction, let $M' = \{0, 1, \ldots, p-1\} \subset \Gamma$ (any choice as in Fact 5.3 would work) and $M = \Gamma \setminus M'$. Returning to our eigenvalue expressions, we have $\alpha_r \in \mathbb{Q}[\zeta_n]$ for every r, so by Fact 5.3, we can express each α_r uniquely as a $\mathbb{Q}$-linear combination of elements of $B = \{\zeta_n^m : m \in M\}$. More specifically, by equations $(*)$ and (1), we can determine exactly that the coefficient of ζ_n^m in the expression of α_r (with respect to the basis B) is $\text{mult}_{rS}(m) - \text{mult}_{rS}(\tilde{m})$, where $\tilde{m}$ is the unique element of M' such that $\tilde{m} \equiv m \pmod p$. By the linear independence of B, we get that $\alpha_{r_1} = \alpha_{r_2}$ if and only if their representations as $\mathbb{Q}$-linear combinations of the elements of B are the same, which happens if and only if for all $m \in M$,

$$\text{mult}_{r_1 S}(m) - \text{mult}_{r_1 S}(\tilde{m}) = \text{mult}_{r_2 S}(m) - \text{mult}_{r_2 S}(\tilde{m}).$$

Thus, if $t \in \{0, 1, \ldots, p-1\}$ represents any residue modulo p, we can set

$$v_t = \text{mult}_{r_1 S}(t) - \text{mult}_{r_2 S}(t),$$

and a rearrangement of the equation above shows that $\text{mult}_{r_1 S}(m) = \text{mult}_{r_2 S}(m) + v_{m \bmod p}$ for all m, as desired. □

At this point, we are ready to prove the theorem by casework on the set S.

Definition 5.4. *Given a multiset S of elements from Γ, we say that S is **organized** if for all $1 \le t \le p-1$ (nonzero residues modulo p), we have*

$$\text{mult}_S(t) = \text{mult}_S(t+p) = \cdots = \text{mult}_S(t + (p-1)p).$$

In other words, the multiplicity of every element of Γ not divisible by p depends only on its residue modulo p.

5.2 The Disorganized Case

Proof of Theorem 1.2 (Disorganized Case). Suppose that S is not organized. We will actually prove the stronger statement that *every* eigenvalue of G has multiplicity at most d in this case.

Lemma 5.5. *If S is not organized, then if $p \mid r_1$ and $p \nmid r_2$, we must have $\alpha_{r_1} \neq \alpha_{r_2}$.*

Proof of Lemma 5.5. We prove the contrapositive, assuming there are $p \mid r_1$ and $p \nmid r_2$ such that $\alpha_{r_1} = \alpha_{r_2}$. By Lemma 5.1, there exist offsets $v_0, \ldots, v_{p-1} \in \mathbb{Z}$ such that for every $m \in \Gamma$, we can write $\text{mult}_{r_1 S}(m) = \text{mult}_{r_2 S}(m) + v_{m \bmod p}$. Since $p \nmid r_2$, the action on Γ of multiplication by r_2 is a bijection, so $\text{mult}_{r_2 S}(m) = \text{mult}_S(r_2^{-1} m)$, where r_2^{-1} is the multiplicative inverse of r_2 modulo p^2. Furthermore, for every $m \not\equiv 0 \bmod p$, we have $\text{mult}_{r_1 S}(m) = 0$ because p divides every element of $r_1 S$. Thus, for every $m \not\equiv 0 \bmod p$, we conclude that

$$\text{mult}_S(r_2^{-1} m) = -v_{m \bmod p}.$$

Varying m over all elements of Γ not divisible by p, we see that $r_2^{-1} m$ also varies over all elements of Γ not divisible by p. Furthermore, for every element $\ell \in \Gamma$ not divisible by p, the value $\text{mult}_S(\ell) = -v_{r_2 \ell \bmod p}$ depends only on the residue of $\ell \bmod p$. Therefore, S is organized, which gives the contrapositive result. □

By Lemma 5.5, we only have to consider eigenvalue multiplicities within the sets $\{\alpha_0, \alpha_p, \ldots, \alpha_{(p-1)p}\}$ and $\{\alpha_r : p \nmid r\}$; there cannot be any overlap. For the first set, we appeal to Theorem 1.1 (the

prime case). Note that

$$\alpha_{kp} = \sum_{s\in S} \zeta_{p^2}^{ksp} = \sum_{s\in S} \zeta_p^{ks} = \sum_{s'\in S'} \zeta_p^{ks'},$$

where $S' = \{s \bmod p : s \in S\}$ is a multiset of elements of $\mathbb{Z}/p\mathbb{Z}$. In other words, $\{\alpha_0, \alpha_p, \ldots, \alpha_{(p-1)p}\}$ are just the eigenvalues of $Cay(\mathbb{Z}/p\mathbb{Z}, S')$, which have multiplicity at most $|S'| = |S| = d$ by Theorem 1.1, as desired[3].

For the set $\{\alpha_r : p \nmid r\}$, we have to do a bit more work. By our assumption that S is not organized, there exist elements $\ell, \ell' \in \Gamma$ not divisible by p such that $\ell \equiv \ell' \pmod p$ and $\text{mult}_S(\ell) > \text{mult}_S(\ell')$. Now, consider any r_1 not divisible by p, and look at the multiplicity of α_{r_1}. By Lemma 5.1, if $\alpha_{r_1} = \alpha_{r_2}$, then there exist offsets $v_0, \ldots, v_{p-1} \in \mathbb{Z}$ such that for all $m \in \Gamma$, $\text{mult}_S(r_1^{-1}m) = \text{mult}_S(r_2^{-1}m) + v_{m \bmod p}$. Plugging in $m = r_1\ell$ and $r_1\ell'$, respectively, we see that

$$\text{mult}_S(\ell) = \text{mult}_S(r_2^{-1}r_1\ell) + v_{r_1\ell \bmod p} > \text{mult}_S(\ell') = \text{mult}_S(r_2^{-1}r_1\ell') + v_{r_1\ell' \bmod p}.$$

But $r_1\ell \equiv r_1\ell' \pmod p$, so this implies that $\text{mult}_S(r_2^{-1}r_1\ell) > \text{mult}_S(r_2^{-1}r_1\ell') \geq 0$. Finally, as r_2 ranges over the elements of Γ not divisible by p, so does $r_2^{-1}r_1\ell$. Therefore, the multiplicity of α_{r_1} is at most the number of r_2's such that $\text{mult}_S(r_2^{-1}r_1\ell) \geq 1$, which is clearly bounded by $|S| = d$. □

5.3 The Organized Case

Proof of Theorem 1.2 (Organized Case). In this case, let $c_t := \text{mult}_S(t)$ for every $1 \leq t \leq p-1$. By the definition of being organized, we can write

$$\begin{aligned}\alpha_r = \sum_{s\in S} \zeta_n^{rs} &= \sum_{j=1}^{p-1} \text{mult}_S(jp) \cdot \zeta_n^{rjp} + \sum_{t=1}^{p-1} c_t \cdot (\zeta_n^{rt} + \zeta_n^{r(t+p)} + \cdots + \zeta_n^{r(t+(p-1)p)}) \\ &= \sum_{j=1}^{p-1} \text{mult}_S(jp) \cdot \zeta_p^{rj} + \sum_{t=1}^{p-1} c_t \cdot \zeta_n^{rt} \cdot (1 + \zeta_p^r + \zeta_p^{2r} + \cdots + \zeta_p^{(p-1)r}).\end{aligned}$$

If r is not divisible by p, we have $1 + \zeta_p^r + \cdots + \zeta_p^{(p-1)r} = 0$, so

$$|\alpha_r| = \left|\sum_{j=1}^{p-1} \text{mult}_S(jp) \cdot \zeta_p^{rj}\right| \leq \sum_{j=1}^{p-1} \text{mult}_S(jp) =: C. \tag{2}$$

[3]Note that in order for G to be connected, there must be at least one element in S that is not divisible by p, which means there is at least one nonzero element in S'.

On the other hand, if $r = kp$, then we instead have

$$\alpha_{kp} = \sum_{j=1}^{p-1} \text{mult}_S(jp) \cdot 1 + \sum_{t=1}^{p-1} c_t \cdot \zeta_p^{kt} \cdot p = C + p \cdot \sum_{t=1}^{p-1} c_t \cdot \zeta_p^{kt}.$$

Let $\tilde{S}$ be the multiset of elements of $\mathbb{Z}/p\mathbb{Z}$ that contains each residue $t \in \{1, \ldots, p-1\}$ with multiplicity c_t (note that $\tilde{S}$ is symmetric because S is symmetric, and $\tilde{S}$ is non-empty because we required our original graph to be connected, which means some c_t is positive). Then the values $(\alpha_{kp} - C)/p$ as k ranges from 0 to $p-1$ are just the eigenvalues of $Cay(\mathbb{Z}/p\mathbb{Z}, \tilde{S})$!

If the second-largest adjacency matrix eigenvalue of $\tilde{G} = Cay(\mathbb{Z}/p\mathbb{Z}, \tilde{S})$ is $\beta > 0$, then there exists k such that $\alpha_{kp} = C + p\beta$ is strictly greater than C, which by (2) implies $\alpha_{kp} > \alpha_r$ for every r such that $p \nmid r$. This means that α_{kp} is the second-largest eigenvalue of G. Furthermore, the multiplicity of α_{kp} equals the multiplicity of the second-largest adjacency matrix eigenvalue of $\tilde{G}$, which is bounded by $|\tilde{S}| = \sum_{t=1}^{p-1} c_t \leq |S| = d$ by Theorem 1.1, as desired.

The last case is when the second-largest adjacency matrix eigenvalue of $\tilde{G}$ is nonpositive, which can only happen if all of the c_t's are 0 or all of them are positive[4]. The c_t's cannot all be 0 because G is connected. Furthermore, $|S| \geq \sum_{t=1}^{p-1} pc_t$, so when $\sum_{t=1}^{p-1} c_t \geq p$, the statement of Theorem 1.2 is trivial because $|S| \geq n$. The only remaining case is when $c_1 = c_2 = \cdots = c_{p-1} = 1$, in which case $\tilde{G}$ is the complete graph on p vertices and $|S| = C + p(p-1)$.

Clearly, the only way for an eigenvalue's multiplicity to exceed $|S| \geq p(p-1)$ would be for the eigenvalue to appear both as α_{kp} for some k and as α_r for some r with $p \nmid r$. We know the eigenvalues of the complete graph, so we have $\alpha_{kp} = C - p$ is an integer for every $k \neq 0$. On the other hand, when $p \nmid r$, we have $\alpha_r = \sum_{j=1}^{p-1} \text{mult}_S(jp) \cdot \zeta_p^{rj}$, which by Corollary A.3 can only be an integer if $\text{mult}_S(jp)$ is constant as j varies.

If $\text{mult}_S(jp) = 0$ for all $j \neq 0$, then $C = 0$ and $\alpha_r = 0$ for all r not divisible by p. Hence, $C - p \neq 0$, so no eigenvalue of the form α_{kp} is equal to an eigenvalue of the form α_r with $p \nmid r$. If $\text{mult}_S(jp) \geq 1$ for all $j \neq 0$, then $C \geq p - 1$, so $|S| \geq p^2 - 1 = n - 1$. This automatically bounds the multiplicity of a nontrivial eigenvalue, so we are done. □

This completes the proof of Theorem 1.2, which we have now established in both the disorganized and organized cases.

[4]This is shown by slightly modifying the proof of the well-known fact that the second-largest adjacency matrix eigenvalue of a simple graph is nonpositive if and only if the graph is complete or complete multi-partite.

6 Products of 2 Distinct Primes

In this final section, we will prove Theorem 1.3. The structure of this section mirrors that of Section 5, but the analysis will be more involved. We begin by restating the theorem for convenience.

Theorem 1.3. *Let $p < q$ be distinct odd primes, and let $n = pq$ and $\Gamma = \mathbb{Z}/n\mathbb{Z}$. Given a symmetric multiset S of nonzero elements of Γ, consider the Cayley graph $G = Cay(\Gamma, S)$. Suppose that G is connected. Then, if S is not **organized** (defined in Definition 6.5), the multiplicity of every eigenvalue of G is at most $|S| = d$.*

6.1 Setup

Given a multiset S of elements from Γ, let $S \bmod p$ be the multiset consisting of the elements $(s \bmod p) \in \mathbb{Z}/p\mathbb{Z}$ where s ranges over the elements of S. Define $S \bmod q$ analogously.

Let p^{-1} denote the smallest positive integer such that $pp^{-1} \equiv 1 \pmod q$, and similarly let q^{-1} denote the smallest positive integer such that $qq^{-1} \equiv 1 \pmod p$. For integers i and j, let $f(i,j)$ denote the unique number t such that $0 \le t < pq$, $t \equiv i \pmod p$, and $t \equiv j \pmod q$. Note that $f(i,j) = (pp^{-1}j + qq^{-1}i) \bmod pq$, since $pp^{-1}j + qq^{-1}i \equiv 0 + 1 \cdot i \equiv i \bmod p$ and it is also equivalent to $j \bmod q$. Depending on context, we also use $f(i,j)$ to refer to the corresponding element of Γ.

Proposition 6.1. *For every two integers i and j, we have*

$$\sum_{j'=0}^{q-1} \zeta_n^{f(i,j')} = \sum_{i'=0}^{p-1} \zeta_n^{f(i',j)} = 0.$$

Proof. We can freely replace an exponent of ζ_n by anything with the same residue modulo $n = pq$ without changing the value of the exponential, so we have

$$\sum_{j'=0}^{q-1} \zeta_n^{f(i,j')} = \sum_{j'=0}^{q-1} \zeta_n^{qq^{-1}i + pp^{-1}j'} = \zeta_n^{qq^{-1}i} \cdot \sum_{j'=0}^{q-1} \zeta_q^{p^{-1}j'} = 0,$$

since in the last summation, the exponent $p^{-1}j'$ ranges over every residue modulo q exactly once. An analogous calculation shows that $\sum_{i'=0}^{p-1} \zeta_n^{f(i',j)} = 0$ as well, as desired. □

With these preliminaries, we now fix p, q, n, Γ, S, and $G = Cay(\Gamma, S)$ a connected graph in the context of Theorem 1.3. We begin by expressing an arbitrary eigenvalue α_r in terms of the basis from Lemma A.2, and then prove a lemma like Lemma 5.1 that characterizes the equivalence of two eigenvalues by a property of their corresponding translates of S.

Fact 6.2. *For every r with $0 \le r < n$, the following equation holds:*

$$\alpha_r = \sum_{i=1}^{p-1}\sum_{j=1}^{q-1}[\mathrm{mult}_{rS}(f(i,j)) - \mathrm{mult}_{rS}(f(i,0)) - \mathrm{mult}_{rS}(f(0,j)) + \mathrm{mult}_{rS}(f(0,0))] \cdot \zeta_n^{f(i,j)} \qquad (\star)$$

Proof. By Proposition 6.1, we know that for every i and j, we have

$$\zeta_n^{f(i,0)} = -\sum_{j'=1}^{q-1} \zeta_n^{f(i,j')} \quad \text{and} \quad \zeta_n^{f(0,j)} = -\sum_{i'=1}^{p-1} \zeta_n^{f(i',j)}.$$

Combining these yields $\zeta_n^{f(0,0)} = \sum_{i'=1}^{p-1}\sum_{j'=1}^{q-1} \zeta_n^{f(i',j')}$. Now that we know how to write every power of ζ_n as a linear combination of the elements of the set $\{\zeta_n^m : \gcd(m,n) = 1\} = \{\zeta_n^{f(i,j)} : 1 \le i \le p-1, 1 \le j \le q-1\}$, we can finally re-express each eigenvalue α_r as a linear combination of the same elements as desired:

$$\begin{aligned}\alpha_r &= \sum_{i=0}^{p-1}\sum_{j=0}^{q-1} \mathrm{mult}_{rS}(f(i,j)) \cdot \zeta_n^{f(i,j)} \\ &= \sum_{i=1}^{p-1}\sum_{j=1}^{p-1} [\mathrm{mult}_{rS}(f(i,j)) - \mathrm{mult}_{rS}(f(i,0)) - \mathrm{mult}_{rS}(f(0,j)) + \mathrm{mult}_{rS}(f(0,0))] \cdot \zeta_n^{f(i,j)}.\end{aligned}$$

□

Lemma 6.3. *Let $r_1, r_2 \in \{0, 1, \ldots, n-1\}$ be any two elements. Then, the corresponding eigenvalues α_{r_1} and α_{r_2} are equal if and only if there exist offsets $v_0, v_1, \ldots, v_{p-1}, w_0, w_1, \ldots, w_{q-1} \in \mathbb{Z}$ such that for all $m \in \Gamma$,*

$$\mathrm{mult}_{r_1 S}(m) = \mathrm{mult}_{r_2 S}(m) + v_{m \bmod p} + w_{m \bmod q}.$$

Proof. For the if direction, if integers v_i and w_j exist and satisfy this property, then by Corollary 2.24,

$$\begin{aligned}\alpha_{r_1} &= \sum_{m \in \Gamma} \mathrm{mult}_{r_1 S}(m) \cdot \zeta_n^m \\ &= \sum_{m \in \Gamma} (\mathrm{mult}_{r_2 S}(m) + v_{m \bmod p} + w_{m \bmod q}) \cdot \zeta_n^m \\ &= \alpha_{r_2} + \sum_{m \in \Gamma} (v_{m \bmod p} \cdot \zeta_n^m + w_{m \bmod q} \cdot \zeta_n^m) \\ &= \alpha_{r_2} + \sum_{i=0}^{p-1}\sum_{j=0}^{q-1} v_i \cdot \zeta_n^{f(i,j)} + \sum_{j=0}^{q-1}\sum_{i=0}^{p-1} w_j \cdot \zeta_n^{f(i,j)}\end{aligned}$$

$$= \alpha_{r_2} + \sum_{i=0}^{p-1} v_i \cdot \sum_{j=0}^{q-1} \zeta_n^{f(i,j)} + \sum_{j=0}^{q-1} w_j \cdot \sum_{i=0}^{p-1} \zeta_n^{f(i,j)}$$
$$= \alpha_{r_2}$$

by Proposition 6.1 applied to the inner sums in the second-to-last line.

For the only if direction, suppose $\alpha_{r_1} = \alpha_{r_2}$. By Lemma A.2, if we write α_{r_1} and α_{r_2} as linear combinations of the elements of the set $\{\zeta_n^m : \gcd(m,n) = 1\}$, then the coefficients must match up in the two expressions. Note that $\gcd(m,n) = 1$ if and only if $p \nmid m$ and $q \nmid m$. By equation $(\star)$ from Fact 6.2, this means that if $\alpha_{r_1} = \alpha_{r_2}$, we must have $M_{ij} - M_{i0} - M_{0j} + M_{00} = 0$ for every i and j such that $1 \le i \le p-1$ and $1 \le j \le q-1$, where $M_{ab} = \text{mult}_{r_1 S}(f(a,b)) - \text{mult}_{r_2 S}(f(a,b))$. Note that the formula also holds trivially by cancellation when $i = 0$, $j = 0$, or both. Let $v_i = M_{i0}$ and let $w_j = M_{0j} - M_{00}$. Then, for every $m \in \Gamma$, if we let $i = m \bmod p$ and $j = m \bmod q$ so that $m = f(i,j)$, we see that

$$\text{mult}_{r_1 S}(m) - \text{mult}_{r_2 S}(m) = M_{ij} = M_{i0} + M_{0j} - M_{00} = v_i + w_j = v_{m \bmod p} + w_{m \bmod q},$$

as desired. □

Corollary 6.4. *If $\gcd(r_1, n) = \gcd(r_2, n) = 1$, then $\alpha_{r_1} = \alpha_{r_2}$ if and only if $\alpha_1 = \alpha_{r_1^{-1} r_2}$.*

Proof. By the lemma, $\alpha_{r_1} = \alpha_{r_2}$ if and only if there exist v_i's and w_j's such that $\text{mult}_{r_1 S}(m) = \text{mult}_{r_2 S}(m) + v_{m \bmod p} + w_{m \bmod q}$ for all $m \in \Gamma$. Writing $m = r_1 m'$ and using the fact that r_1 and r_2 are relatively prime to n, we see that this is equivalent to having $\text{mult}_S(m') = \text{mult}_{r_1^{-1} r_2 S}(m') + v_{r_1 m' \bmod p} + w_{r_1 m' \bmod q}$ for all $m' \in \Gamma$. But $v_{r_1 m' \bmod p}$ only depends on $m' \bmod p$ and $w_{r_1 m' \bmod q}$ only depends on $m' \bmod q$, so this is equivalent to saying $\alpha_1 = \alpha_{r_1^{-1} r_2}$. □

Like in Section 5, we will need a new concept of an "organized" generating set.

Definition 6.5. *Given a multiset S of elements from Γ, we say that S is **organized** if for every i and j such that $1 \le i \le p-1$ and $1 \le j \le q-1$, we have*

$$\text{mult}_S(f(i,j)) = \text{mult}_S(f(i,1)) + \text{mult}_S(f(1,j)) - \text{mult}_S(f(1,1)).$$

Proposition 6.6 (Alternate Definition of Organized)**.** *A multiset S of elements of Γ is organized if and only if there exist integers $a_1, \ldots, a_{p-1}$ and $b_1, \ldots, b_{q-1}$ such that $\text{mult}_S(f(i,j)) = a_i + b_j$ for all $1 \le i \le p-1$ and $1 \le j \le q-1$.*

Proof. If $\text{mult}_S(f(i,j)) = a_i + b_j$ for every i and j, then $\text{mult}_S(f(i,j)) = a_i + b_j = (a_i + b_1) + (a_1 + b_j) - (a_1 + b_1)$, so S is organized. If S is organized, just take $a_i = \text{mult}_S(f(i,1))$ and $b_j = \text{mult}_S(f(1,j)) - \text{mult}_S(f(1,1))$ □

6.2 The Disorganized Case

Lemma 6.7. *If S is not organized, then if $\gcd(r_1, n) > 1$ and $\gcd(r_2, n) = 1$, we must have $\alpha_{r_1} \neq \alpha_{r_2}$.*

Proof. We prove the contrapositive. Suppose r_1 and r_2 are such that $\gcd(r_1, n) > 1$ and $\gcd(r_2, n) = 1$ but $\alpha_{r_1} = \alpha_{r_2}$. By assumption, r_1 is divisible by at least one of p and q, so WLOG assume $p \mid r_1$. By Lemma 6.3, there exist integers $v_0, \ldots, v_{p-1}, w_0, \ldots, w_{q-1}$ such that for every $m \in \Gamma$, we can write $\text{mult}_{r_1 S}(m) = \text{mult}_{r_2 S}(m) + v_{m \bmod p} + w_{m \bmod q}$. Since $\gcd(r_2, n) = 1$, multiplication by r_2 acts as a bijection on Γ, so $\text{mult}_{r_2 S}(m) = \text{mult}_S(r_2^{-1} m)$, where r_2^{-1} is the multiplicative inverse modulo n. On the other hand, $r_1 S$ only consists of elements that are divisible by p. Therefore, for every i and j such that $1 \leq i \leq p-1$ and $1 \leq j \leq q-1$, we have

$$0 = \text{mult}_{r_1 S}(f(i,j)) = \text{mult}_S(r_2^{-1} f(i,j)) + v_i + w_j.$$

By the Chinese Remainder Theorem, we have $r_2^{-1} f(i,j) \equiv f(r_2^{-1} i, r_2^{-1} j) \bmod n$, so by re-indexing we can rewrite the last line as $\text{mult}_S(f(i,j)) = -v_{r_2 i} - w_{r_2 j}$, which still holds whenever $1 \leq i \leq p-1$ and $1 \leq j \leq q-1$. By Proposition 6.6, this means that S is organized, as desired. □

By the lemma, we need only consider eigenvalue multiplicities separately within the sets $\{\alpha_r : \gcd(r, n) > 1\}$ and $\{\alpha_r : \gcd(r, n) = 1\}$.

Proposition 6.8. *Every eigenvalue in the set $\{\alpha_r : \gcd(r, n) > 1\}$ has multiplicity at most d.*

Proof. By $(*)$, if $r = kp$, we have

$$\alpha_r = \sum_{s \in S} \zeta_n^{kps} = \sum_{s \in S} \zeta_q^{ks} = \sum_{s \in S} \zeta_q^{k(s \bmod q)}.$$

Therefore, the eigenvalues α_{kp} are the eigenvalues of the Cayley graph $G'_q := Cay(\mathbb{Z}/q\mathbb{Z}, S \bmod q)$. Similarly, the eigenvalues $\alpha_{\ell q}$ are the eigenvalues of the Cayley graph $G'_p := Cay(\mathbb{Z}/p\mathbb{Z}, S \bmod p)$. Note the overlap at $\alpha_0 = d$ corresponding to the trivial eigenvalue of both G'_p and G'_q, which is only counted once in our original graph. Since we assume that our graph G is connected, each of $S \bmod q$ and $S \bmod p$ must contain at least one non-zero element, so by Theorem 1.1, the multiplicity of every eigenvalue of G'_p is at most $|S \bmod p| = |S| = d$, and the multiplicity of every eigenvalue of G'_q is similarly at most d. If the eigenvalues of G'_p and G'_q are distinct except for the trivial eigenvalue, then we're done. But what if there is a nontrivial overlap?

Note that every eigenvalue of G'_p is an element of $\mathbb{Z}[\zeta_p]$ and every eigenvalue of G'_q is an element of $\mathbb{Z}[\zeta_q]$. Furthermore, $\mathbb{Z}[\zeta_p] \cap \mathbb{Z}[\zeta_q] = \mathbb{Z}$, so if there is an eigenvalue that is shared between G'_p and G'_q, it must be an integer. Other than the trivial eigenvalue d, if G'_q has an eigenvalue that is an integer,

it must be the case that $S \bmod q$ contains each of $1, 2, \ldots, q-1$ with the same multiplicity by Corollary A.3. If this multiplicity is larger than 1, then $|S| \geq 2(q-1) \geq p+q-1 = |\{r : \gcd(r,n) > 1\}|$, so we automatically achieve the desired result.

Finally, suppose $S \bmod q$ contains each of $1, 2, \ldots, q-1$ once. Let a be the multiplicity of 0 in $S \bmod q$, and b be the multiplicity of 0 in $S \bmod p$. In order for G'_p to have an eigenvalue that is an integer, we know that $S \bmod p$ has to contain each of $1, 2, \ldots, p-1$ with the same multiplicity, so let this multiplicity be x. In this situation, we have that the nontrivial eigenvalues of G'_q are all $a-1$ and the nontrivial eigenvalues of G'_p are all $b-x$. In addition, by counting the elements of S in two ways (by their residues mod p or mod q), we see that $x(p-1)+b = q-1+a$. Thus, to have an eigenvalue overlap, we must have $a-1 = b-x \implies b-a = x-1$, but we must also then have $x(p-1)+b-a+1 = q \implies xp = q$. This is impossible, since $p \neq q$ and q is prime. □

Proposition 6.9. *Every eigenvalue in the set $E := \{\alpha_r : \gcd(r,n) = 1\}$ has multiplicity at most d within E.*

Proof. By Corollary 6.4, if $\gcd(r,n) = 1$, then the multiplicity of α_r within E is the same as the multiplicity of α_1 within E. Therefore, it suffices to prove that α_1 has multiplicity at most d within E. Let $\Gamma^\times \cong \mathbb{F}_p^* \times \mathbb{F}_q^*$ be the multiplicative group consisting of the residues modulo n that are relatively prime to n, and let $H = \{r \in \Gamma^\times : \alpha_1 = \alpha_r\}$. Applying Corollary 6.4 again, we see that H must be a subgroup of $\Gamma^\times$. It remains to show that $|H| \leq |S|$.

Let $\bullet \bmod p$ and $\bullet \bmod q$ be the projections of $\Gamma^\times$ onto its first and second factors, respectively. We see that $\Phi_1 := H \bmod p$ is a subgroup of $\mathbb{F}_p^*$ and $\Phi_2 := H \bmod q$ is a subgroup of $\mathbb{F}_q^*$, and moreover H is a subgroup of $\Phi := \Phi_1 \times \Phi_2$. Let $D_{ij} = \mathrm{mult}_S(f(i,j)) - \mathrm{mult}_S(f(i,0)) - \mathrm{mult}_S(f(0,j))$, which is the coefficient of $\zeta_n^{f(i,j)}$ in Equation $(\star)$ of Fact 6.2 for α_1, since we assume that S does not contain 0. By Lemma A.2, we see that if $r \in H$, then $D_{r^{-1}i,r^{-1}j} = D_{ij}$ for every $1 \leq i \leq p-1$ and $1 \leq j \leq q-1$, since r is relatively prime to n. In other words, $D_{ij} = D_{i'j'}$ whenever $f(i,j)$ and $f(i',j')$ are in the same coset of H. If cH is a coset of H in $\Gamma^\times$, let D_{cH} denote the value of D_{ij} for i and j such that $f(i,j) \in cH$.

By Goursat's Lemma, there are subgroups N_1 of Φ_1 and N_2 of Φ_2 such that H is a union of cosets of $N_1 \times N_2$ and the image of H in $\Phi_1/N_1 \times \Phi_2/N_2$ is the graph of an isomorphism between Φ_1/N_1 and Φ_2/N_2. For a coset $iN_1 \subset \mathbb{F}_p^*$ or a coset $jN_2 \subset \mathbb{F}_q^*$, define

$$i^* = \arg\min_{i' \in iN_1} \mathrm{mult}_S(f(i',0)) \quad \text{and}$$
$$j^* = \arg\min_{j' \in jN_2} \mathrm{mult}_S(f(0,j')).$$

Notice that $f(i^*, j^*)$ is in the same coset of H as $f(i,j)$, so $D_{i^*j^*} = D_{ij}$. Also define the following non-negative integer quantity:

$$Q := \sum_{iN_1 \subset \mathbb{F}_p^*} \sum_{jN_2 \subset \mathbb{F}_q^*} \text{mult}_S(f(i^*, j^*)),$$

where the sums are taken over distinct cosets of N_1 and N_2, respectively.

Lemma 6.10. *The following three statements hold:*

1. *Q is divisible by k, where $k := |\Phi_1/N_1| = |\Phi_2/N_2|$.*
2. *If $Q = 0$, then S is organized.*
3. *$|S| \geq Q \cdot |N_1| \cdot |N_2|$.*

Proof. We begin with statement (1). Rewriting $\text{mult}_S(f(i^*, j^*)) = D_{i^*j^*} + \text{mult}_S(f(i^*, 0)) + \text{mult}_S(f(0, j^*))$ in the definition of Q, we see that

$$\begin{aligned} Q &= \sum_{iN_1 \subset \mathbb{F}_p^*} \sum_{jN_2 \subset \mathbb{F}_q^*} \Big(D_{i^*j^*} + \text{mult}_S(f(i^*, 0)) + \text{mult}_S(f(0, j^*)) \Big) \\ &= \left(\sum_{cH \subset \Gamma^\times} k \cdot D_{cH} \right) + \left(\sum_{iN_1 \subset \mathbb{F}_p^*} |\mathbb{F}_q^*/N_2| \cdot \text{mult}_S(f(i^*, 0)) \right) + \left(\sum_{jN_2 \subset \mathbb{F}_q^*} |\mathbb{F}_p^*/N_1| \cdot \text{mult}_S(f(0, j^*)) \right), \end{aligned}$$

where the first sum is taken over distinct cosets of H in $\Gamma^\times$. The first term arises from the fact that every coset of H is the union of exactly k cosets of $(N_1 \times N_2)$, and specifying a coset of $N_1 \times N_2 \subset \Gamma^\times$ is equivalent to specifying unique cosets $iN_1 \subset \mathbb{F}_p^*$ and $jN_2 \subset \mathbb{F}_q^*$. Now, $|\mathbb{F}_q^*/N_2| = |\mathbb{F}_q^*/\Phi_2| \cdot |\Phi_2/N_2|$ is divisible by k, as is $|\mathbb{F}_p^*/N_1|$, so each of these summations is divisible by k. Therefore, Q is divisible by k.

Moving on to statement (2), if $Q = 0$, then since every $\text{mult}_S(f(i^*, j^*))$ is non-negative, we conclude that for every pair of cosets $iN_1 \subset \mathbb{F}_p^*$ and $jN_2 \subset \mathbb{F}_q^*$, we have $\text{mult}_S(f(i^*, j^*)) = 0$. This implies that for every $1 \leq i \leq p-1$ and $1 \leq j \leq q-1$, we have

$$\begin{aligned} \text{mult}_S(f(i,j)) &= D_{ij} + \text{mult}_S(f(i,0)) + \text{mult}_S(f(0,j)) \\ &= D_{i^*j^*} + \text{mult}_S(f(i,0)) + \text{mult}_S(f(0,j)) \\ &= -\text{mult}_S(f(i^*,0)) - \text{mult}_S(f(0,j^*)) + \text{mult}_S(f(i,0)) + \text{mult}_S(f(0,j)). \end{aligned}$$

It is now clear that S is organized, since $\text{mult}_S(f(i,0)) - \text{mult}_S(f(i^*,0))$ depends only on i and $\text{mult}_S(f(0,j)) - \text{mult}_S(f(0,j^*))$ depends only on j.

Finally, statement (3) follows from the definitions of i^* and j^*. Specifically, we have

$$\begin{aligned}\text{mult}_S(f(i,j)) &= D_{ij} + \text{mult}_S(f(i,0)) + \text{mult}_S(f(0,j)) \\ &= D_{i^*j^*} + \text{mult}_S(f(i,0)) + \text{mult}_S(f(0,j)) \\ &\geq D_{i^*j^*} + \text{mult}_S(f(i^*,0)) + \text{mult}_S(f(0,j^*)) = \text{mult}_S(f(i^*,j^*)).\end{aligned}$$

Therefore,

$$|S| \geq \sum_{i=1}^{p-1}\sum_{j=1}^{q-1} \text{mult}_S(f(i,j)) \geq \sum_{i=1}^{p-1}\sum_{j=1}^{q-1} \text{mult}_S(f(i^*,j^*)) = |N_1| \cdot |N_2| \cdot Q,$$

since each pair (i^*, j^*) occurs in the sum exactly $|N_1| \cdot |N_2|$ times. □

In this section, we have assumed that S is not organized, so by statement (2) of the lemma, it must be that $Q > 0$. But by statement (1) of the lemma, this means that $Q \geq k$, and combined with statement (3), we get $|S| \geq k \cdot |N_1| \cdot |N_2| = |H|$, as desired. This completes the proof of Proposition 6.9. □

Evidently, Theorem 1.3 comes out as a direct consequence of Lemma 6.7, Proposition 6.8, and Proposition 6.9. Furthermore, it's worth noting that the condition of being organized is extremely restrictive: the multiplicities of $f(i,1)$ and $f(1,j)$ for every $i, j \neq 0$ completely determine the multiplicities of all $f(i',j')$ with $i', j' \neq 0$. In some sense, then, "almost all" Cayley graphs on $\mathbb{Z}/(pq)\mathbb{Z}$ satisfy the conditions of Theorem 1.3.

6.3 The Organized Case

In Section 5, obtaining a bound on the multiplicity of the second-largest eigenvalue in the organized case required putting together a patchwork of clever little statements to catch every edge case. Here, some headway can be made, but it involves some rather torturous casework and doesn't have a satisfying resolution (in my opinion). Instead, here is the example that breaks the barrier of d:

Example 6.11. Let $p < q$ be odd primes, and let $n = pq$ and $\Gamma = \mathbb{Z}/n\mathbb{Z}$. Let $S = \{m \neq 0 : \gcd(m,n) \neq p\}$. Then, the Cayley graph $G = Cay(\Gamma, S)$ has degree $|S| = (p-1)q$, but the nontrivial eigenvalues of G can be shown (using our usual techniques) to be 0 with multiplicity $p(q-1)$ and $-q$ with multiplicity $p-1$. Therefore, the second-largest adjacency matrix eigenvalue has multiplicity $p(q-1) = pq - p > pq - q = (p-1)q = d$.

Note that the degree of this example is very large (almost n), and we still have a bound on the multiplicity of $\frac{p}{p-1} \cdot d$. I am not aware of any smaller examples that have second-largest eigenvalue multiplicity exceeding d, which is why I believe that Conjecture 1.4 is true.

A Cyclotomic Fields and Goursat's Lemma

Here, we state and prove some facts about cyclotomic fields and sums of roots of unity.

Fact A.1. *[Mar18, Theorem 3] Given the cyclotomic field $\mathbb{Q}[\zeta_n]$, we have $[\mathbb{Q}[\zeta_n] : \mathbb{Q}] = \varphi(n)$, where φ is Euler's totient function. Furthermore, the minimal polynomial of ζ_n over $\mathbb{Q}$ has degree $\varphi(n)$ and has roots*

$$\{\zeta_n^m : 1 \le m \le n, \gcd(m,n) = 1\}.$$

Lemma A.2. *If n is odd and squarefree, then $\{\zeta_n^m : 1 \le m \le n, \gcd(m,n) = 1\}$ is a basis for $\mathbb{Q}[\zeta_n]$ as a $\mathbb{Q}$-vector space. In particular, these roots of unity are $\mathbb{Q}$-linearly independent.*

Proof of Lemma A.2. Let $n = p_1 p_2 \cdots p_r$ be a squarefree odd integer, where the p_i's are distinct odd primes, and let $B = \{\zeta_n^m : 1 \le m \le n, \gcd(m,n) = 1\}$. By Fact A.1, the size of any $\mathbb{Q}$-basis for $\mathbb{Q}[\zeta_n]$ is $\varphi(n) = |B|$, so it suffices to show that B spans all of $\mathbb{Q}[\zeta_n]$. In fact, $\mathbb{Q}[\zeta_n]$ is obviously spanned by $\{1, \zeta_n, \zeta_n^2, \ldots, \zeta_n^{n-1}\}$, so we only need to show that each of these powers of ζ_n can be represented as a $\mathbb{Q}$-linear combination of the elements of B.

Let m be any integer such that $0 \le m \le n-1$, and let $g = \gcd(m,n)$. If $g = 1$, then ζ_n^m is in B, so it is definitely in the span of b. Otherwise, we can assume WLOG that there exists a positive integer $k \le r$ such that $g = p_1 \cdots p_k$. We claim that

$$\zeta_n^m = (-1)^k \sum_{a_1=1}^{p_1-1} \sum_{a_2=1}^{p_2-1} \cdots \sum_{a_k=1}^{p_k-1} \zeta_n^{f(a_1,a_2,\ldots,a_k,m,m,\ldots,m)},$$

where $f(x_1, \ldots, x_r)$ is the unique integer modulo n that is equivalent to $x_i \bmod p_i$ for every i. By assumption, m is not divisible by p_{k+1} through p_r, so this summation represents a linear combination of elements of B, since every exponent is relatively prime to n. In other words, we have reduced the problem to proving this claim.

To prove the claim, we induct on k. The base case of $k = 0$ is immediate: the equation just reads $\zeta_n^m = \zeta_n^{f(m,m,\ldots,m)}$, which is true because $m = f(m, m, \ldots, m)$. For the inductive step,

observe that we can rearrange the order of summation to get

$$\begin{aligned}(-1)^k \sum_{a_1=1}^{p_1-1}\sum_{a_2=1}^{p_2-1}\cdots\sum_{a_k=1}^{p_k-1}\zeta_n^{f(a_1,a_2,\ldots,a_k,m,m,\ldots,m)} &= -\sum_{a_k=1}^{p_k-1}\left[(-1)^{k-1}\sum_{a_1=1}^{p_1-1}\sum_{a_2=1}^{p_2-1}\cdots\sum_{a_{k-1}=1}^{p_{k-1}-1}\zeta_n^{f(a_1,a_2,\ldots,a_{k-1},a_k,m,\ldots,m)}\right] \\ &= -\sum_{a_k=1}^{p_k-1}\zeta_n^{f(0,\ldots,0,a_k,m,\ldots,m)} \\ &= \zeta_n^m\cdot\left(-\sum_{a_k=1}^{p_k-1}\zeta_n^{f(0,\ldots,0,a_k,0,\ldots,0)}\right),\end{aligned}$$

where we applied the inductive hypothesis for the quantity in brackets for each fixed a_k. Finally, notice that for each $j \in \{1,\ldots,p_k-1\}$, we have $j\cdot(n/p_k) = f(0,\ldots,0,j\cdot n/p_k,0,\ldots,0)$ where the nonzero entry is in the kth position, and $p_k \nmid j\cdot(n/p_k)$ by definition. Therefore,

$$-\sum_{a_k=1}^{p_k-1}\zeta_n^{f(0,\ldots,0,a_k,0,\ldots,0)} = -\sum_{j=1}^{p_k-1}\zeta_n^{j\cdot(n/p_k)} = -\sum_{j=1}^{p_k-1}\zeta_{p_k}^{j} = 1,$$

which completes the proof of the claim. □

Corollary A.3. *If p is prime and $G = Cay(\mathbb{Z}/p\mathbb{Z}, S)$ is any Cayley graph, then a nontrivial eigenvalue of G is an integer if and only if* $\mathrm{mult}_S(1) = \mathrm{mult}_S(2) = \cdots = \mathrm{mult}_S(p-1)$.

Proof of Corollary A.3. By Corollary 2.24, the nontrivial eigenvalues of the Cayley graph $G = Cay(\mathbb{Z}/p\mathbb{Z}, S)$ are given by $\alpha_r = \sum_{s\in S}\zeta_p^{rs} = \mathrm{mult}_S(0) + \sum_{j=1}^{p-1}\mathrm{mult}_S(j)\cdot\zeta_p^{rj}$ for $r = 1,2,\ldots,p-1$. Since $\mathrm{mult}_S(0)$ is an integer, the only way for α_r to be an integer is if $\sum_{j=1}^{p-1}\mathrm{mult}_S(j)\cdot\zeta_p^{rj}$ is an integer. Because p is prime and $r \in \{1,2,\ldots,p-1\}$, the exponents in the summation are just a rearrangement of the numbers from 1 to $p-1$.

By the preceding lemma, we know that $B := \{\zeta_p, \zeta_p^2, \ldots, \zeta_p^{p-1}\}$ are linearly independent. Furthermore, $-c\cdot\zeta_p - c\cdot\zeta_p^2 - \cdots - c\cdot\zeta_p^{p-1} = c$ for any number c, which when c is an integer must represent the unique $\mathbb{Q}$-linear combination of elements of B that produces the value c. Therefore, in order for α_r to be an integer, it must be the case that $\mathrm{mult}_S(1) = \mathrm{mult}_S(2) = \cdots = \mathrm{mult}_S(p-1)$. Finally, if $\mathrm{mult}_S(1) = \cdots = \mathrm{mult}_S(p-1)$, then $\alpha_r = \mathrm{mult}_S(0) - \mathrm{mult}_S(1)$ is an integer, which proves the reverse direction. □

Fact A.4 (Goursat's Lemma, [Ser16, Proposition 1.6])**.** *Let Γ_1 and Γ_2 be two groups, and let $\Gamma = \Gamma_1\times\Gamma_2$ be their direct product, which comes with projections $p_1 : \Gamma\to\Gamma_1$ and $p_2 : \Gamma\to\Gamma_2$. Let $H \subseteq \Gamma$ be a subgroup, and define $H_1 = p_1(H)$ and $H_2 = p_2(H)$. The following statements are true:*

1. *H_1 is a subgroup of Γ_1, H_2 is a subgroup of Γ_2, and H is a subgroup of $H_1\times H_2$.*

2. *If* $N_1 = \{x \in \Gamma_1 : (x, 1_{\Gamma_2}) \in H\}$ *and* $N_2 = \{y \in \Gamma_2 : (1_{\Gamma_1}, y) \in H\}$*, then* N_1 *is a normal subgroup of* H_1 *and* N_2 *is a normal subgroup of* H_2*.*

3. H *is a union of cosets of* $(N_1 \times N_2)$*.*

4. $H_1/N_1 \cong H_2/N_2$*, and the image of* H *in* $H_1/N_1 \times H_2/N_2$ *is the graph of an isomorphism between* H_1/N_1 *and* H_2/N_2*.*

References

[BB24] Igor Balla and Matija Bucić. "Equiangular lines via improved eigenvalue multiplicity". In: *arXiv preprint arXiv:2409.16219* (2024).

[Con] Keith Conrad. "Characters of Finite Abelian Groups". Online resource. URL: https://kconrad.math.uconn.edu/blurbs/grouptheory/charthy.pdf.

[DJRVZ24] Tommaso d'Orsi et al. *Sparsest cut and eigenvalue multiplicities on low degree Abelian Cayley graphs*. 2024. arXiv: 2412.17115 [cs.DS]. URL: https://arxiv.org/abs/2412.17115.

[HSZZ23] Milan Haiman et al. *Graphs with high second eigenvalue multiplicity*. 2023. arXiv: 2109.13131 [math.CO]. URL: https://arxiv.org/abs/2109.13131.

[JTYZZ21] Zilin Jiang et al. "Equiangular lines with a fixed angle". In: *Annals of Mathematics* 194.3 (2021), pp. 729–743.

[Mar18] Daniel A. Marcus. *Number fields*. eng. Second edition. Universitext. Cham: Springer, 2018 - 2018. ISBN: 9783319902326.

[ML08] Y Makarychev and J Lee. "Eigenvalue multiplicity and volume growth". In: *arXiv preprint arXiv:0806.1745* (2008).

[MRS23] Theo McKenzie, Peter M. R. Rasmussen, and Nikhil Srivastava. *Support of Closed Walks and Second Eigenvalue Multiplicity of the Normalized Adjacency Matrix*. 2023. arXiv: 2007.12819 [math.CO]. URL: https://arxiv.org/abs/2007.12819.

[Ser16] Jean-Pierre Serre. *Finite groups : an introduction*. eng. Surveys of modern mathematics ; volume 10. Somerville, Massachusetts: International Press, 2016. ISBN: 9781571463203.

[Spi25] Daniel Spielman. "Spectral and Algebraic Graph Theory". Incomplete Draft. 2025. URL: http://cs-www.cs.yale.edu/homes/spielman/sagt/sagt.pdf.

[Tre16] Luca Trevisan. "Lecture Notes on Graph Partitioning, Expanders and Spectral Methods". In: *URL: https://lucatrevisan.github.io/books/expanders-2016.pdf* (2016).