The Evolution of Expander Graphs

A Thesis Presented

by

David Y. Xiao

to

Computer Science

in partial fulfillment of the honors requirement

for the degree of

Bachelor of Arts

Harvard College

Cambridge, Massachusetts

April 8th, 2003

Acknowledgements

First and foremost I'd like to thank my advisor Prof. Salil Vadhan for many fruitful discussions and for his support and guidance. Many thanks to Prof. Michael Rabin and Minh-Huyen Nguyen for their encouragement and help in the editing process. And my deepest gratitude to my parents for giving me the opportunity to reach this point. All that I've accomplished is because of you.

Contents

Chapter 1

Introduction and Preliminaries

1.1 Introduction

1.1.1 History

The study of expander graphs has been a rapidly developing subject in discrete mathematics and computer science in the past three decades. It has been motivated from many directions, including network design, algorithms, coding, cryptography, and pseudorandomness. Why has this kind of graph structure been so useful in so many diverse areas? Expander graphs have very special properties, indeed properties that at first seem to be self-conflicting.

Intuitively, a regular undirected graph is an *expander* if it is highly connected. That is, it is easy to get from any vertex to any other vertex in very few steps. In order for such graphs to be interesting, we also impose that they have low degree, because in all applications the graph's "cost" is related to its degree. Because these two requirements are in tension, it is remarkable that these graphs exist at all. However, a line of work initiated by Pinsker [23] and culminating in the recent work by Friedman [11] showed using the probabilistic method that randomly chosen graphs in fact enjoy these properties with high probability.

Many measurements have been posited to quantify this expansion property. Such measures include vertex expansion and edge expansion, properties that unfortunately are hard to compute [7]. However, as was shown in the classic work by Alon, Milman, and Tanner [3, 5, 27], the vertex expansion of a graph is intimately tied to its spectrum, and in particular the second-largest eigenvalue. This offers us an efficiently computable quantification of the expansion of a graph. More importantly, it allows us to apply the tools of linear algebra in analyzing expander graphs.

Using this approach, Alon and Boppana [3] showed the upper bound on expansion that we can hope for in an infinite family of graphs. Soon afterward, Lubotzky, Phillips, and Sarnak [20] and Margulis [21] independently constructed families of graphs reaching this bound. In the process, [20] coined the term Ramanujan graph, which has since been applied to all graphs reaching the bound given in [3].

There has been an explosion of interest recently in expander graphs in theoretical computer science, particularly with regard to their role in reducing the dependence of probabilistic algorithms on uniformly random bits. As such, explicit constructions of expanders have been very important. However, until recently all of the expander graph constructions have been algebraic, usually leveraging the theory of finite fields and the Cayley graphs of certain groups. Only recently has there been any success using combinatorial tools to create families of expander graphs. The zig-zag product of Reingold, Vadhan, and Widgerson [25] gives a recursive construction of expander graph families that yields to a pleasingly intuitive analysis.

1.1.2 Overview

This thesis has two goals. First, we present an introduction to the line of work that began with the study of expander graphs in the non-constructive setting, which then led to the algebraic construction of expanders, and finally has recently produced combinatorial constructions. Second, we extend the work on combinatorial constructions by new analyses and new constructions.

The thesis is arranged as follows. In Chapter 1 we define the terminology and prove several basic results about graph spectra. In Chapter 2 we apply spectral methods to the analysis of expander graphs and expander graph families, concluding with a discussion of the application of expander graphs to the study of derandomization. In Chapter 3 we give two expander graph family constructions, contrasting classical algebraic techniques and newer, more intuitive combinatorial techniques. In Chapter 4 we present work extending the zig-zag product to a group-theoretic interpretation. In Chapter 5 we explore the problem of finding lower bounds for the expansion of zig-zag products. In Chapter 6 we analyze a new combinatorial operation on graphs inspired by the zig-zag product that we call the semi-square, and we proceed to prove some results on its relation to expanders.

1.1.3 Contributions

Specific constributions of this thesis follow. In Chapter 2 we give a novel proof of Theorem 2.3.2 which shows that infinite families of graphs have bounded expansion. In Chapter 3 we give a new non-bipartite treatment of the Gabber-Galil construction in Theorem 3.1.1, and we give an improvement on the original result of Theorem 3.2.6 bounding the expansion of the zig-zag product. In Chapter 4 we coin the term wide zig-zag product and present the relationship between the wide zig-zag product and the semi-direct product from group theory (Theorem 4.3.3) differently, and hopefully more clearly, than in the original work.

Chapters 5 and 6 present entirely new material. In Chapter 5 we explore further questions concerning the second largest eigenvalue of the zig-zag product. We prove some results in special cases about a lower bound on the second largest eigenvalue. In the process, we invent the semi-square graph operation (Definition 5.3.3), a new combinatorial operation on graphs. We also exhibit experimental data that point to universal lower bounds for the zig-zag product. Finally, in Chapter 6 we analyze the semi-square. We show in Proposition 6.1.6 that, for certain graphs, semi-squaring always gives better expansion than plain squaring, and in Theorem 6.2.3 we specify completely the class of graphs where semi-squaring offers absolutely no benefit compared to squaring.

1.2 Definitions

There are many different ways to quantify expansion, all of them related. Here we present two of them: vertex expansion and spectral expansion. For most of this thesis we will be concerned with spectral expansion, but the fact that there is a link between the combinatorial property of vertex expansion and the algebraic property of spectral expansion is inherently intriguing and thus worth discussing. We will prove this loose equivalence in Chapter 2.

1.2.1 Graph Theory Terminology and Notation

Definition 1.2.1. We will employ the following graph theory terminology, all of which are standard.

- An *undirected multi-graph* G consists of a vertex set V and an edge multi-set $E \subset V \times V$. We will use $N = |V|$ to denote the *size* of the graph. The edge multi-set E is allowed to contain repeated edges between the same pair of edges, as well as any number of self-loops.
- For any subset of vertices $S \subset V$, we define its *neighbor set* $N(S) = \{j \in V \mid \exists i \in S \text{ s.t. } \{i,j\} \in E\}$.
- The *degree* of a vertex is the number of incident edges, where a self-loop counts as a single edge. A graph is said to be *regular* with degree D if all its vertices have degree D.
- A graph is *connected* if for any pair of vertices i and j, there is a path of edges from i to j.
- A graph is *bipartite* if there is a partition of V into subsets X and Y so that there are no edges between any two vertices of X or any two vertices of Y.
- A *subgraph* H of G is a set $H = (V', E')$ were $V' \subset V$ and $E' \subset (V' \times V') \cap E$. That is, H's vertices are vertices of G, and its edges are also edges in G.
- An *induced subgraph* H of G is a subgraph $H = (V', E')$ where $E' = (V' \times V') \cap E$. H is uniquely determined by its vertex set V' since its edge set includes all edges in G that connect vertices in V'.
- Let $p_{i,j}$ be a path from vertex i to vertex j, and let $|p_{i,j}|$ be its length. Then the *distance* between i and j is the length of the shortest path between them: $d(i,j) = \min_{p_{i,j}} |p_{i,j}|$.
- The *diameter* K of the graph is longest of the distances between any pair of vertices, $K = \max_{i,j} d(i,j)$.

Remark 1.2.2. We will work almost exclusively with connected regular undirected non-bipartite multi-graphs, which we will simply call *graphs* unless otherwise noted.[1]Also, unless otherwise noted, we will adopt the convention of labelling the vertices of a graph on N vertices using the integers $[N] = \{1, 2, \ldots, N\}$.

[1] The study of expander graphs has many interesting results on bipartite graphs, and indeed much of the classical work was done in a bipartite setting. However, because most recent work has focused on the non-bipartite setting, we will cast the classical results in this light.

1.2.2 Vertex Expansion

Definition 1.2.3. A graph $G = (V, E)$ on N vertices is called a γ-vertex expander if

$$S \subset V,\ |S| \leq N/2 \implies |N(S)| \geq \gamma |S|$$

We would like γ to be as large as possible, and in particular we would like $\gamma > 1$, since then the neighbor set of a subset of vertices will be larger than the starting subset. It is evident that this definition matches our intuition of expansion: any set of vertices will "expand" into a larger set of vertices when we follow the edges of some subset of vertices.

1.2.3 Algebra Terminology and Notation

Definition 1.2.4. We will employ the following algebra terminology. Most of the terminology is standard, and we preface non-standard definitions with $**$.

- We will work mainly over $\mathbb{R}$, though sometimes we will pass into $\mathbb{C}$. Where applicable below, our definitions over $\mathbb{R}$ extend to $\mathbb{C}$ in the obvious way.
- Let $\mathbb{R}^{n\times m}$ denote the space of $n \times m$ real matrices. Let $M_n(\mathbb{R}) = \mathbb{R}^{n\times n}$. We will usually treat column vectors $v \in \mathbb{R}^n$ as matrices in $\mathbb{R}^{n\times 1}$.
- Let I_n denote the $n \times n$ identity matrix. We write I when there is no confusion about the size of the matrix.
- Let A^t denote the matrix *transpose* of A. A is *symmetric* if $A = A^t$. A is *orthogonal* if $AA^t = I$.
- The *hermitian adjoint* of a matrix $A \in M_n(\mathbb{C})$ is denoted by $A^* = \overline{A^t}$, i.e. the conjugate transpose of A. A is *hermitian* if $A = A^*$. A is *unitary* if $AA^* = I$.
- A real matrix $A \in M_n(\mathbb{R})$ with entries $a_{ij} \geq 0$ is called *doubly stochastic* if $\sum_j a_{kj} = \sum_i a_{ik} = 1$ for all k.
- The *dot product* of two vectors is $\langle x, x\rangle = x^* x$. The L_1-*norm* of a vector $x \in \mathbb{R}^n$ is denoted by $|x|_1$ and equals $\sum_{i=1}^n |x_i|$. The L_2-*norm* of x is denoted by $\|x\|$ and equals $\sqrt{\langle x, x\rangle}$.
- Two vectors x, y are said to be *orthogonal* if $\langle x, y\rangle = 0$. If x, y are orthogonal, we write $x \perp y$. A set of vectors $v_1, \ldots, v_k$ are *orthonormal* if they are pair-wise orthogonal and $\|v_i\| = 1$ for all i.
- $\lambda \in \mathbb{R}$ is an *eigenvalue* of $A \in M_n(\mathbb{R})$ if there exists a corresponding non-zero *eigenvector* $x \in \mathbb{R}^n$ such that $Ax = \lambda x$. The set of all eigenvalues of a matrix A is called the *spectrum* of A.
- The *multiplicity* of an eigenvalue is the number of times it is repeated in the spectrum. Since we will work solely with diagonalizable matrices, we do not distinguish between the algebraic and geometric multiplicities of eigenvalues.
- Let $\mathrm{Sp}(v_1, \ldots, v_n)$ denote the span of vectors $v_1, \ldots, v_n$.
- The span of some eigenvectors of a matrix is called an *eigenspace*.

- For any subspace $W \subset V$, define the *orthogonal complement* $W^{\perp} = \{x \in V \mid x \perp w \; \forall w \in W\}$ to be the subspace of all vectors orthogonal to all vectors in W.
- $**$ Let $u \in \mathbb{R}^n$ denote the *uniform distribution* or *uniform vector*, $u = [\frac{1}{n}, \ldots, \frac{1}{n}]^t$.
- $**$ A vector $v \in \mathbb{R}^n$ is *anti-uniform* if $v \perp u$.
- $**$ Likewise, a subspace $V \subset \mathbb{R}^n$ is *anti-uniform* if it is contained in $\mathrm{Sp}(u)^{\perp}$.

1.2.4 Spectral Expansion

When we refer to any algebraic properties (e.g. spectrum, eigenvectors, etc.) of a graph G, we will mean the properties of its normalized adjacency matrix A, which is defined as follows. We will always work with this normalized version of the adjacency matrix.

Definition 1.2.5. With our convention of labelling the vertices of G by $[N]$, we define the *normalized adjacency matrix* of a graph G to be $A = [a_{ij}]$, where

$$a_{ij} = \frac{d_{ij}}{D}$$

where d_{ij} is the number of edges between vertices i and j and D is the degree of the graph.

We may view A as the doubly stochastic transition matrix for the random walk on G. It is doubly stochastic because G is regular, and therefore $\sum_j a_{kj} = \sum_i a_{ik} = 1$ for any fixed k. This matrix is real, non-negative and symmetric since the graph is undirected, and so from the spectral theorem of real symmetric matrices from linear algebra we know that all its eigenvalues are real. We state the spectral theorem here without proof as it is a standard result; the interested reader may refer to [6] for a proof.

Theorem 1.2.6. *Let A be a real symmetric $N \times N$ matrix. Then there exist orthogonal (equiv. orthonormal) eigenvectors $\{v_1, \ldots, v_N\}$ corresponding to N (not necessarily distinct) real eigenvalues $\lambda_1, \ldots, \lambda_N$.*

We will show in the next section that 1 is an eigenvalue with the uniform distribution u as the corresponding column eigenvector. We will also show 1 is the largest eigenvalue in absolute value. All other eigenvalues are at most 1 in absolute value, and we will be interested in the second-largest eigenvalue in absolute value. From here on, "second largest eigenvalue" will always refer to the second largest eigenvalue in absolute value unless otherwise noted.

Definition 1.2.7. Let $\lambda_1, \lambda_2, \ldots \lambda_N$ be the spectrum of the graph G, with the ordering $1 = \lambda_1 \geq |\lambda_2| \geq \ldots \geq |\lambda_N|$. We say that G is a λ-spectral expander if $|\lambda_2| \leq \lambda$. Furthermore, define the function $\lambda_2(G) = |\lambda_2|$.

In this definition, we want $|\lambda_2| \leq 1$ to be bounded as far away from 1 as possible.

We will often augment this definition with the parameters N and D, so that a graph G is an (N, D, λ)-spectral expander if it has N vertices, it is regular with degree D, and has $\lambda_2(G) \leq \lambda$.

We appeal to a different intuition to see the motivation behind this definition. One may say a graph is a good expander if, starting from any initial probability distribution on its vertices, taking a random walk on the graph will converge to the uniform distribution on the vertices very quickly. This matches up with our above intuition because the graph must be "highly connected" in order for this to happen.

Translating this to the spectral domain, consider any probability distribution $x \in \mathbb{R}^N$ over the vertices of the graph, where $\forall i \; x_i \geq 0$ and $\sum_{i=1}^N x_i = 1$. Since A is real and symmetric, we know from the spectral theorem that we can write x as a linear combination of the eigenvectors $v_1, \ldots, v_N$, corresponding to $\lambda_1, \ldots, \lambda_N$, where $v_1 = u$ the uniform distribution and the other v_i are anti-uniform. It is clear the coefficient of u in this decomposition is 1 since $\sum_{i=1}^N x_i = 1$. Therefore we have

$$x = u + c_2 v_2 + \ldots + c_N v_N$$

where $c_i \in \mathbb{R}$. If we repeatedly apply the normalized adjacency matrix A, we see that

$$A^k x = u + \lambda_2^k c_2 v_2 + \ldots + \lambda_N^k c_N v_N$$

Now if we consider the distance of $A^k x$ from u, say using the L_2 norm, it is clear that the smaller λ_2, the quicker the higher-order terms vanish and the faster $A^k x$ converges to u. In terms of being an expander, this means the smaller λ_2 is, the fewer steps it takes to reach all vertices in the graph with (almost) equal probability.

Remark 1.2.8. Another useful way of looking at the action of A on a distribution x is to envision each vertex i sending an equal amount of its old weight along each edge to its neighbors, and in turn receiving its new weight from its neighbors in the same manner. This is borne out in the formula for matrix product: $(Ax)_i = \frac{1}{D} \sum_{\{i,j\} \in E} x_j$.

1.3 Basic Graph Spectra Properties

To gain an intuition of how a graph's spectrum is related to its structure and to build the tools that we will use to analyze graphs, we begin by proving some simple results about the relationship between certain elementary graph properties and the eigenvalues of the graph.

Lemma 1.3.1. *The eigenvalues of any graph G have absolute value at most 1. G has an eigenvalue of 1 with the uniform vector u as an eigenvector.*

Proof. Consider any eigenvalue λ of G and its corresponding eigenvector v. There is some i such that $|v_i| = \max_j |v_j|$, and where $|v_i| > 0$ since $v \neq 0$. Then, using the triangle inequality and the fact that $|v_i|$ is maximal, we have that $|\lambda| \leq 1$ because

$$|\lambda||v_i| = \left| \sum_j a_{ij} v_j \right| \leq |v_i| \cdot \sum_j |a_{ij}| = |v_i|$$

$\sum_j |a_{ij}| = 1$ because A is non-negative and doubly stochastic. The formula above shows that equality occurs for the uniform distribution u. □

Remark 1.3.2. Equality occurs only if we have either

1. $\forall j, \; a_{ij} \neq 0 \implies v_j = v_i$
2. $\forall j, \; a_{ij} \neq 0 \implies v_j = -v_i$

For example, this occurs with u.

Because we wish to bound the second largest eigenvalue of a graph in absolute value, it is often convenient to have a guarantee that the eigenvalues of the graph are non-negative. To assure this is the case, we will often work with the square of a graph G^2:

Definition 1.3.3. The *square* of a graph $G = (V, E)$ is the graph $G^2 = (V, E^2)$ where (with appropriate multiplicity) $\{i, j\} \in E^2 \Leftrightarrow \exists k \in V \ s.t. \ \{i, k\} \in E$ and $\{k, j\} \in E$. If the normalized adjacency matrix of G is A, then the normalized adjacency matrix of G^2 is the matrix square A^2.

It is clear that if λ is an eigenvalue of G, then λ^2 is an eigenvalue of G^2 with the same eigenvector.

Using these simple tools, we may relate the spectrum of a graph G to the properties of connectedness and bipartiteness.

Lemma 1.3.4. *A graph G is connected if and only if the eigenvalue 1 occurs with multiplicity 1.*

Proof. Suppose G is disconnected, then we show that the eigenvalue 1 has multiplicity at least 2. G has a connected component $X \subsetneq V$, and let $Y = V - X$. For any set $S \subset V$, let χ_S be its characteristic vector, i.e. it is 1 for all elements of S and 0 elsewhere. We claim that $x = \chi_X$ and $y = \chi_Y$ are both eigenvectors with eigenvalue 1.

$$(Ax)_i = \sum_j a_{ij} x_j = \sum_{j \in X} a_{ij} x_j = \left\{ \begin{array}{l} 1, \, i \in X \\ 0, \, i \notin X \end{array} \right.$$

which immediately implies $Ax = x$. An analogous argument holds for y.

Now suppose 1 occurs with multiplicity > 1. Then there is some anti-uniform eigenvector $x \perp u$ with eigenvalue 1. Let $X = \{i \mid x_i = \max_j x_j\}$. Since $\lambda = 1$, from Remark 1.3.2, this means for any $i \in X$, and for any $a_{ij} \neq 0$, it must be that $x_i = x_j$. So $a_{ij} = 0$ for all $i \in X, \ j \in V - X$, and so X is not connected to $V - X$. □

Remark 1.3.5. This lemma is easily generalized in one direction to non-regular graphs. The statement is that if G is disconnected, then the eigenvalues of each of its connected components occur with the same multiplicity in the spectrum of G. Let v be an eigenvector of a connected component H with eigenvalue λ, then the vector $\hat{v}$ is also an eigenvector of G with eigenvalue λ, where $\hat{v}$ is equal to v on H and is 0 on all vertices outside H.

Lemma 1.3.6. *A connected graph G is bipartite if and only if it has an eigenvalue of -1.*

Proof. If a graph is bipartite, then its square is disconnected. By Lemma 1.3.4, G^2 has eigenvalue 1 with multiplicity 2. Each eigenvalue of G^2 is the square of an eigenvalue of G. Since G is connected, it has eigenvalue 1 with multiplicity 1, so this corresponds with one of G^2's eigenvalues of 1. Since the other eigenvalue 1 of G^2 must also be the square of an eigenvalue of G, G must have an eigenvalue of -1.

If the graph has an eigenvalue of -1, then there is an eigenvector $x \perp u$ such that $Ax = -x$. Let $X = \{i \mid x_i = \max_j x_j\}$ and $Y = V - X$. Since $\lambda = -1$, from Remark 1.3.2 we have for any $i \in X$, and for any $a_{ij} \neq 0$, it must be that $x_i = -x_j$. So $a_{ij} = 0$ for any $i, j \in X$. Furthermore, since G is connected this means $\forall i \in Y$ we have $x_i = -\max_j x_j$, so therefore there are no edges between vertices in Y. So G is bipartite with left side X and right side Y. □

We may interpret both these results in terms of the random walk intuition. If the graph is disconnected, then no random walk can take us from one component to another, and so the distributions on each component is independent. Hence, the eigenvalue 1 occurs with multiplicity equal to the number of components, since they don't interact.

Similarly, when the graph is bipartite then we can take two "almost-independent" random walks simultaneously. One walk starts on the left-hand side and has negative weight, while the other starts on the right-hand side and has positive weight. The steps alternate between the two sides, and one may think of the sign of the weights as denoting to which of the walks the weight belongs to. Hence, the (-1)-eigenvector is uniform on each side but with different sign.

These elementary results hint at the power of the connection between a graph's spectrum and its combinatorial properties. [9] and [14] gives a detailed introduction to this study, relating other graph properties to the spectrum, including girth, diameter, and chromatic number. In the next chapter, we will elaborate on spectral characteristics peculiar to expander graphs.

Chapter 2

Expander Graph Spectra

In this chapter, we highlight some of the fundamental results of expanders independent of their construction. We develop the main lemmas that will be used to analyze expander graph family constructions. We establish the correspondence between vertex expansion and spectral expansion, and go on to show that there is a limit to how good the expansion of an infinite family of graphs may be. We conclude by showing how spectral analysis is used in the application of expander graphs to randomness reduction.

2.1 Eigenvalue Formulæ

We begin by reviewing the classical Rayleigh quotient from linear algebra, which will provide the necessary tool to calculate bounds on the second largest eigenvalue.

Lemma 2.1.1. *Let $A \in M_n(\mathbb{R})$ be symmetric. Let $\lambda_1 \geq \ldots \geq \lambda_n$ be its spectrum with the corresponding orthonormal eigenvectors $\{v_1, \ldots, v_n\}$. For any $k \in [n]$, define $W_k \subsetneq \mathbb{R}^n$ to be the proper subspace $\mathrm{Sp}(v_1, \ldots, v_k)$. Then we have for all k that*

$$x \in W_k \setminus \{0\} \quad \Longrightarrow \quad \frac{\langle Ax, x\rangle}{\langle x, x\rangle} \geq \lambda_k \tag{2.1}$$

$$y \in W_k^{\perp} \setminus \{0\} \quad \Longrightarrow \quad \frac{\langle Ay, y\rangle}{\langle y, y\rangle} \leq \lambda_{k+1} \tag{2.2}$$

Equality holds exactly when x, y are the eigenvectors corresponding to λ_k, λ_{k+1} respectively.

Proof. Since $v_1, \ldots v_k$ form a basis of W_k, we may write $x = \sum_{i=1}^{k} c_i v_i$ for some constants $c_i \in \mathbb{R}$. By orthonormality, we get

$$\begin{aligned}\frac{\langle Ax, x\rangle}{\langle x, x\rangle} &= \frac{\sum_{i=1}^{k} \lambda_i c_i^2 \|v_i\|^2}{\langle x, x\rangle} = \frac{\sum_{i=1}^{k} \lambda_i c_i^2}{\langle x, x\rangle} \\ &\geq \frac{\sum_{i=1}^{k} \lambda_k c_i^2}{\langle x, x\rangle} = \frac{\lambda_k \sum_{i=1}^{k} c_i^2}{\langle x, x\rangle} \\ &= \lambda_k\end{aligned}$$

Similarly, since $v_1, \dots v_n$ span $\mathbb{R}^n$, this means $W_k^{\perp} = \mathrm{Sp}(v_{k+1}, \dots v_n)$. Thus, using an analogous argument gives the upper bound on $\frac{\langle Ay, y\rangle}{\langle y, y\rangle}$.

The equality condition is obvious given the above derivation. □

The following lemma uses the Rayleigh quotient and provides us with our primary method of upper-bounding the second largest eigenvalue of an arbitrary graph.

Lemma 2.1.2. *For any graph G with normalized adjacency matrix A,*

$$\lambda_2(G) = \max_{x \perp u} \frac{|\langle Ax, x\rangle|}{\langle x, x\rangle} \tag{2.3}$$

$$= \max_{x \perp u} \frac{\|Ax\|}{\|x\|} \tag{2.4}$$

$$= \max_{x \perp u} \left| 1 - \frac{1}{D\|x\|^2} \sum_{\{i,j\} \in E} (x_i - x_j)^2 \right| \tag{2.5}$$

Proof. Label the eigenvalues and eigenvectors as in Lemma 2.1.1. From Lemma 1.3.1 we know the spectrum of G has $\lambda_1 = 1$ and $v_1 = u$.

First we show Equation (2.3). If $\lambda_2(G) > 0$, then $\lambda_2(G) = \lambda_2$, so the Equation (2.2) from Lemma 2.1.1 becomes an equality by taking $k = 1$ and letting $y = v_2 \in W_1^{\perp}$. If $\lambda_2(G) < 0$, then $\lambda_2(G) = \lambda_n$, so take $k = n$ and $x = v_n \in W_n$, and thus Equation (2.1) becomes an equality with $\frac{\langle Ax, x\rangle}{\langle x, x\rangle} = \lambda_n$ and hence $\frac{|\langle Ax, x\rangle|}{\langle x, x\rangle} = |\lambda_n|$.

To show Equation (2.4), we see that $\|Ax\| = \sqrt{|\langle Ax, Ax\rangle|} = \sqrt{|\langle A^2x, x\rangle|}$. By Equation (2.3), the maximum of this is $\lambda_2(G) \cdot \|x\|$.

To show Equation (2.5), we have by the definition of dot product that

$$\begin{aligned} |\langle Ax, x\rangle| &= \left| \sum_{i,j} a_{ij} x_i x_j \right| = \left| \frac{2}{D} \sum_{\{i,j\} \in E} x_i x_j \right| = \left| \|x\|^2 - \frac{1}{D} \left(\|x\|^2 D - \sum_{\{i,j\} \in E} 2x_i x_j \right) \right| \\ &= \left| \|x\|^2 - \frac{1}{D} \sum_{\{i,j\} \in E} (x_i - x_j)^2 \right| \end{aligned}$$

where last equality holds because $\|x\|^2 \cdot D = \sum_{\{i,j\} \in E} (x_i^2 + x_j^2)$. □

Remark 2.1.3. We can compute *any* eigenvalue λ_n of G using Equations (2.3) or (2.5) by removing the absolute value and restricting our choice of x to the appropriate subspace as stated in Lemma 2.1.1 instead of simply restricting $x \perp u$. If we are just concerned with the eigenvalues of matrices, we may use Equations (2.3) and (2.4) to compute the eigenvalues of A a complex hermitian matrix, since it is well-known that the eigenvalues of a hermitian matrix are real [6].

In order to more precisely quantify the expansion of certain graphs, we will also be interested in lower-bounding the second-largest eigenvalue. One useful tool for this is the following eigenvalue interlacing theorem originally stated by Courant and Hilbert.

Lemma 2.1.4 ([10]). *Let A be a real symmetric $n \times n$ matrix, and let R be a $n \times m$ matrix such that $R^tR = I$. Then $B = R^tAR$ is a $m \times m$ matrix whose eigenvalues interlace those of*

A. That is, if we let $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$ be the eigenvalues of A and $\mu_1 \geq \mu_2 \geq \ldots \geq \mu_m$ be the eigenvalues of B, then for all $k \in [m]$ we have

$$\lambda_k \geq \mu_k \geq \lambda_{n-m+k}$$

Proof. Let $v_1, \ldots, v_n$ be the orthogonal eigenvectors of A corresponding to $\lambda_1, \ldots, \lambda_n$ and let $w_1, \ldots, w_n$ be the orthogonal eigenvectors of B corresponding to $\mu_1, \ldots, \mu_m$. Let $V_k = \mathrm{Sp}(v_1, \ldots, v_k)$ and let $W_k = \mathrm{Sp}(w_1, \ldots, w_k)$. For any k, the dimension of W_k is k, while the dimension of $R^t V_{k-1}$ is at most $k-1$, so its orthogonal complement has dimension at least $n-k+1$. Thus W_k and $(R^t V_{k-1})^\perp$ have a non-zero intersection, so there is some non-zero vector $y \in W_k \cap (R^t V_{k-1})^\perp$. By definition, $\langle y, R^t v_i \rangle = 0$ for all $i \in [k-1]$, and so by symmetry we have $\langle Ry, v_i \rangle = 0$ so we have $Ry \in V_{k-1}^\perp$. We may now apply $R^t R = I$, $R^t A R = B$, and Lemma 2.1.1 to get

$$\lambda_k \geq \frac{\langle ARy, Ry \rangle}{\langle Ry, Ry \rangle} = \frac{\langle By, y \rangle}{\langle y, y \rangle} \geq \mu_k$$

Applying the same analysis with $-A$ and $-B$ gives the lower bound. □

This theorem is often stated in the special case where R contains an identity submatrix and is otherwise zero, in which case B is a principal submatrix of A. However, we will need the more general form above in order to use it to analyze graphs.

We can immediately apply Lemma 2.1.4 to the analysis of spectral expansion. The corresponding statement is that the expansion of a graph G is bounded by the expansion of any of its induced subgraphs. However, we need to clarify the normalized adjacency matrix we use to describe the subgraph, which is different than usual because we normalize the degrees with respect to the host graph.

Definition 2.1.5. The *normalized adjacency matrix of H as a subgraph of G* is given by the matrix $B = [b_{ij}]$

$$b_{ij} = \frac{d_{ij}}{D}$$

where d_{ij} is the number of edges between vertices i and j in H and D is the degree of the host graph G.

Let M be the size of H. The *spectrum of H as a subgraph of G* is of the form $1 \geq \lambda_1^G \geq |\lambda_2^G| \geq \ldots \geq |\lambda_M^G|$, where λ_1^G is not necessarily 1 since B is not necessarily doubly stochastic. Let $\lambda_i^G(H) = |\lambda_i^G|$.

Corollary 2.1.6 ([14]). *Let G be a graph with second largest eigenvalue $\lambda_2(G)$. Let H be any induced subgraph of G. Then $\lambda_2^G(H) \leq \lambda_2(G)$.*

Proof. Let N be the size of G and M be the size of H, and let the vertices in the induced subgraph be $\{n_1, \ldots, n_M\} \subset [N]$. Let $R = [r_{ij}]$ be the $N \times M$ matrix defined as

$$r_{ij} = \begin{cases} 1, \, i = n_j \\ 0, \text{ else} \end{cases}$$

A simple computation shows $R^t R = I$. We claim $B = R^t A R$ is the normalized adjacency matrix of H. This is because if we let $B = [b_{ij}]$, we get

$$b_{ij} = \sum_{k,\ell} r_{ki} a_{k\ell} r_{\ell j} = a_{n_i n_j}$$

So by Lemma 2.1.4 the eigenvalues of H interlace the eigenvalues of G, and therefore $\lambda_2^G(H) \leq \lambda_2(G)$. □

We can compute $\lambda_i^G(H)$ for any i using Equations (2.3) or (2.4) from Lemma 2.1.2 and restricting to the appropriate subspace. However, we cannot use Equation (2.5) directly because part of its proof assumed the graph is regular. The following is the analogue for the case of $\lambda_i^G(H)$.

Lemma 2.1.7. *Let H be an induced subgraph of G with adjacency matrix $B = R^tAR$. Let M be the size of H. Then, with the maximum restricted to the appropriate subspace as specified in Lemma 2.1.1,*

$$\lambda_i^G(H) = \max_x \left| 1 - \frac{1}{D} \sum_{\{i,j\} \in E} ((Rx)_i - (Rx)_j) \right| \tag{2.6}$$

where E is the edge set of G (not H)

Proof. We know

$$\lambda_i^G(H) = \max_x |\langle R^tARx, x\rangle| = \max_x |\langle ARx, Rx\rangle|$$

Thus we may apply the analysis for Equation (2.5) using Rx instead of x. □

These are the main tools that we will use in our analysis of expander graphs. More case-specific techniques will be developed later, of course, but it is worthwhile to remark that these two results, and especially Lemma 2.1.2, will be referred to throughout this thesis.

2.2 Vertex Expansion vs. Spectral Expansion

We are now ready to prove the relationship between the two definitions of expansion given in Section 1.2. This remarkable relationship was first discovered in a series of works by Alon, Milman, and Tanner [3, 5, 27]. They provide the conceptual bridge between the combinatorial notion of vertex expansion and the algebraic notion of spectral expansion.

This connection is useful in several contexts. For example, it has been shown that computing the vertex expansion of a graph is **coNP**-complete [7], whereas computing the spectrum of a graph can be done in polynomial time. However, computing the spectrum is still unwieldy because it takes time polynomial in the size of the graph (the number of vertices). Unfortunately, the representation of a graph is often logarithmic in its size; for example it may simply be a function that computes the c'th neighbor of a vertex i.

A more important result of this relationship is that we are now able to apply tools from linear algebra and analysis in analyzing the expansion of families of graph. Thus, constructing a family of good spectral expanders immediately implies constructing a family of good vertex expanders. In many cases, we can prove interesting results about applications of expander graphs directly using spectral expansion.

However, in some contexts it is more convenient to use the vertex expansion formulation. For example, Sipser and Spielman [26] use a form of vertex expansion to construct expander codes.

The relationship is given in the following two theorems, which show that there is a loose equivalence between the two forms of expansion.

Theorem 2.2.1 ([5, 27]). *If G is a λ-spectral expander, then it is also a $\frac{2}{\lambda^2+1}$-vertex expander.*

Proof. Our intuition for the theorem is as follows: for any $S \subset V$, consider the probability distribution $u_S = \frac{1}{|S|}\chi_S$ that is uniform on S and zero elsewhere. Because $\lambda < 1$, we know that applying the normalized adjacency matrix A "flattens" the distribution considerably according to the L_2 norm. Since the L_1 norm remains unchanged, this means that the weight previously resting solely on S must have been spread to a larger set $N(S)$.

Formally, we can write $u_S^\perp = u_S - u$, where we can check $u_S^\perp \perp u$ since $\langle u_S^\perp, u\rangle = \langle u_S, u\rangle - \langle u, u\rangle = 0$. Simple linear algebra shows that $\|u_S^\perp\|^2 = \|u_S\|^2 - \|u\|^2 = \frac{1}{|S|} - \frac{1}{N}$.

Claim. For any distribution x, $\|x\|^2 \geq \frac{1}{|x_+|}$, where $x_+ = \{i \mid x_i > 0\}$ is the support of x.

This follows from the Cauchy-Schwarz inequality:

$$1 = \left(\sum_{i \in x_+} x_i\right)^2 = \langle \chi_S, x\rangle^2 \leq \|\chi_S\|^2 \|x\|^2 = |x_+| \cdot \sum_{i \in x_+} x_i^2$$

which immediately implies the claim.

Using these facts we see that

$$\frac{1}{|N(S)|} = \frac{1}{|(Au_S)_+|} \leq \|Au_S\|^2$$

A simple computation reveals $Au_S^\perp \perp u$, so we may write

$$\frac{1}{|N(S)|} - \frac{1}{N} \leq \|Au_S\|^2 - \|u\|^2 = \|A(u_S - u)\|^2 = \|Au_S^\perp\|^2$$

By Lemma 2.1.2, we know

$$\|Au_S^\perp\|^2 \leq \lambda^2 \|u_S^\perp\|^2 = \lambda^2 \left(\frac{1}{|S|} - \frac{1}{N}\right)$$

and so we can conclude

$$\frac{1}{|N(S)|} - \frac{1}{N} \leq \lambda^2 \left(\frac{1}{|S|} - \frac{1}{N}\right)$$

Rearranging gives us that

$$|N(S)| \geq \frac{1}{\lambda^2(1 - \frac{|S|}{N}) + \frac{|S|}{N}} |S|$$

and the theorem follows since $|S| \leq \frac{N}{2}$. □

Note that this relationship behaves as we expect: the smaller the λ, the better the vertex expansion. Furthermore, λ strictly less than 1 implies vertex expansion strictly greater than 1. Kahale [18] presents a tighter relationship between spectral and vertex expansion, which we omit here because it does not present further intuition into the relationship.

The next theorem gives the converse relationship: the better the vertex expansion, the better the spectral expansion.

Theorem 2.2.2 ([3]). *If G is a $(1+\alpha)$-vertex expander, then it is also a λ-spectral expander for $\lambda = \sqrt{1 - \frac{\alpha^2}{D^2 \cdot (8+4\alpha^2)}}$.*

Proof. Let A^2 be the normalized adjacency matrix of G^2. For the remainder of the proof, we work with $G^2 = (V, E)$ in order to guarantee that all eigenvalues are positive. The vertex expansion of G^2 must be at least $1+\alpha$: for any subset $S \subset V$ of size at most $N/2$, we know $|N(S)| \geq (1+\alpha)|S|$. Now if $|N(S)| \leq N/2$, we may apply vertex expansion again. If $|N(S)| > N/2$, we may consider any subset $S' \subset N(S)$ of size $N/2$, and we get

$$|N(N(S))| \geq |N(S')| \geq (1+\alpha)N/2 \geq (1+\alpha)|S|$$

Let x be the eigenvector of A^2 with eigenvalue $\lambda_2(G^2)$, the second largest eigenvalue, which in this case must be positive. Since $x \perp u$, x must have both positive and negative entries. Define

$$V_+ = \{i \mid x_i > 0\} \qquad\qquad V_- = \{i \mid x_i \leq 0\}$$

Both V_+ and V_- are non-empty, and we take *w.l.o.g.* $|V_+| \leq N/2$. Finally, define the vector $\bar{x}$ where

$$\bar{x}_i = \begin{cases} x_i,\ x_i > 0 \\ 0, \text{ else} \end{cases}$$

A simple computation verifies that

$$\lambda_2(G^2) = \frac{\langle A^2 x, \bar{x}\rangle}{\langle \bar{x}, \bar{x}\rangle}$$

We wish to change this expression into a form similar to Equation (2.5) from Lemma 2.1.2. We perform the following manipulations to do this:

$$\begin{aligned}
\langle A^2 x, \bar{x}\rangle &= \sum_{i,j} a_{ij} x_i \bar{x}_j \\
&= \|\bar{x}\|^2 - \frac{1}{D^2}\left(\|\bar{x}\|^2 D^2 - \sum_{\substack{i \in V_+ \\ \{i,j\} \in E}} 2\bar{x}_i x_j\right) \\
&= \|\bar{x}\|^2 - \frac{1}{D^2}\left(\|\bar{x}\|^2 D^2 - \sum_{\substack{i,j \in V_+ \\ \{i,j\} \in E}} 2\bar{x}_i \bar{x}_j - \sum_{\substack{i \in V_+,\ j \notin V_+ \\ \{i,j\} \in E}} 2\bar{x}_i x_j\right)
\end{aligned}$$

where D^2 is the degree of G^2. Since $x_j \leq 0$ for $j \notin V_+$, the above becomes

$$\begin{aligned}
|\langle A^2 x, \bar{x}\rangle| &\leq \|\bar{x}\|^2 - \frac{1}{D^2}\left(\|\bar{x}\|^2 D^2 - \sum_{\{i,j\} \in E} 2\bar{x}_i \bar{x}_j\right) \\
&\leq \|\bar{x}\|^2 - \frac{1}{D^2} \sum_{\{i,j\} \in E} (\bar{x}_i - \bar{x}_j)^2
\end{aligned}$$

and therefore

$$\lambda_2(G^2) \leq 1 - \frac{1}{D^2} \frac{\sum_{\{i,j\} \in E} (\bar{x}_i - \bar{x}_j)^2}{\sum_{i \in V} \bar{x}_i^2} \tag{2.7}$$

Note the direction of this inequality is opposite from the one obtained from Equation (2.5).

Now we appeal to the max-flow min-cut theorem. Let us construct the directed weighted graph $\bar{G}$ with vertex set $\{s,t\} \cup V_+ \cup V$ where s is the source and t is the sink, with the understanding that here V_+ and V are copies, so as to be disjoint. We define the edges and weights in the graph as follows:

1. For each $i \in V_+$, the edge (s,i) has capacity $1+\alpha$.
2. For each $i \in V_+,\ j \in V$, the edge (i,j) has capacity 1 if $\{i,j\} \in E$ in the original graph, and zero otherwise.
3. For each $j \in V$, the edge (j,t) has capacity 1.

The min-cut of this network is $(1+\alpha) \cdot |V_+|$. It is clear that the cut consisting of all arcs (s,i) for all $i \in V_+$ has this capacity. To see that it is the minimum, consider any other cut C, and let $W = \{i \in V_+ \mid (s,i) \notin C\}$. The intuition behind the reason this directed graph is linked with G^2's expansion is because these sets W must expand by at least $(1+\alpha)$. Therefore, the flow through them is large enough so as to be greater than the cut constructed above.

Formally, since $|W| \leq \frac{N}{2}$, we must have $|N(W)| \geq (1+\alpha)|W|$, and for each $j \in N(W)$, C must contain an edge incident to j. So the capacity of the cut must be at least $(1+\alpha)|V_+ - W| + |N(W)| \geq (1+\alpha)|V_+|$.

Because the min-cut is large, by the max-flow min-cut theorem, there then exists a function $F : V \times V \to \mathbb{R}$ with the following properties. Here, $\tilde{E}$ is the set of ordered edges: $\{i,j\} \in E \implies (i,j),(j,i) \in \tilde{E}$.

1. $\forall (i,j) \notin \tilde{E}$, $F(i,j) = 0$, and $\forall (i,j) \in \tilde{E},\ 0 \leq F(i,j) \leq 1$
2. For any fixed $i \in V_+$, $\sum_{j|(i,j)\in\tilde{E}} F(i,j) = 1+\alpha$ and $F(i,j) = 0$ for $i \notin V_+$
3. For any fixed $j \in V$, $\sum_{i|(i,j)\in\tilde{E}} F(i,j) \leq 1$.

This function F is the crucial link we need to show the vertex expansion of G^2. We will be able to use the following two bounds to bound expression (2.7).

In the summations below, we take care to write $(i,j) \in \tilde{E}$ to specify summing over ordered pairs (so that each $\{i,j\} \in E$ is counted twice). Applying the above and the fact that $2(a^2+b^2) \geq (a+b)^2$ for any real a,b gives

$$\begin{aligned}
\sum_{(i,j)\in\tilde{E}} F^2(i,j)(\bar{x}_i + \bar{x}_j)^2 &\leq 2\sum_{(i,j)\in\tilde{E}} F^2(i,j)(\bar{x}_i^2 + \bar{x}_j^2) \\
&= 2\sum_{i\in V} \bar{x}_i^2 \left(\sum_{(i,j)\in\tilde{E}} F^2(i,j) + \sum_{(i,j)\in\tilde{E}} F^2(j,i) \right) \\
&\leq (4+2\alpha^2)\sum_{i\in V} \bar{x}_i^2 \\
\sum_{(i,j)\in\tilde{E}} F(i,j)(\bar{x}_i^2 - \bar{x}_j^2) &= \sum_{i\in V} \bar{x}_i^2 \left(\sum_{(i,j)\in\tilde{E}} F(i,j) - \sum_{(i,j)\in\tilde{E}} F(j,i) \right) \\
&\geq \alpha \sum_{i\in V} \bar{x}_i^2
\end{aligned}$$

Multiplying by expression (2.7) by $1 = \frac{\sum_{(i,j)\in\tilde{E}} F^2(i,j)(\bar{x}_i+\bar{x}_j)^2}{\sum_{(i,j)\in\tilde{E}} F^2(i,j)(\bar{x}_i+\bar{x}_j)^2}$ and then applying Cauchy-Schwarz, we get that

$$\begin{aligned}
\lambda_2(G^2) &\leq 1 - \frac{1}{D^2}\frac{\sum_{\{i,j\}\in E}(\bar{x}_i - \bar{x}_j)^2}{\sum_{i\in V}\bar{x}_i^2} \\
&= 1 - \frac{1}{D^2}\frac{\sum_{\{i,j\}\in E}(\bar{x}_i - \bar{x}_j)^2 \cdot \sum_{(i,j)\in\tilde{E}} F^2(i,j)(\bar{x}_i+\bar{x}_j)^2}{\sum_{i\in V}\bar{x}_i^2 \cdot \sum_{(i,j)\in\tilde{E}} F^2(i,j)(\bar{x}_i+\bar{x}_j)^2} \\
&\leq 1 - \frac{1}{D^2}\frac{(\sum_{(i,j)\in\tilde{E}} F(i,j)(\bar{x}_i^2 - \bar{x}_j^2))^2}{2(4+2\alpha^2)(\sum_{i\in V}\bar{x}_i^2)^2} \\
&\leq 1 - \frac{\alpha^2}{D^2(8+4\alpha^2)}
\end{aligned}$$

From this bound on $\lambda_2(G^2)$, it immediately follows that $\lambda_2(G) \leq \sqrt{1 - \frac{\alpha^2}{D^2(8+4\alpha^2)}}$. □

Remark 2.2.3. The above bound may be substantially improved if we know that the second largest eigenvalue of G is positive. In that case, we can analyze G directly instead of G^2, and the bound becomes $1 - \frac{\alpha^2}{D(8+4\alpha^2)}$.

2.3 Expander Graph Families

For many applications we do not know *a priori* the size of the expander we wish to employ. So instead of focusing on specific graphs with good expansion, we will be interested in infinite families of graphs where each graph exhibits good expansion. Beginning with a simple cased proven in Pinsker [23], there has been a large body of work that uses the probabilistic method to show that nearly all large enough graphs are expanders. Friedman [11] recently proved Alon's original conjecture of that nearly all large enough graphs are expanders in an extremely general setting. Because the result is very recent and the proof is extremely long and technical, we simply state the result here without proof.

Theorem 2.3.1 ([11]). *Fix any $\varepsilon > 0$ and an even positive integer D. Then $\exists c > 0$ such that for a random degree D graph on N vertices, we have with probability $1 - \frac{c}{n^{\Omega(\sqrt{D})}}$ that*

$$|\lambda_2(G)| \leq \frac{2\sqrt{D-1}}{D} + \varepsilon$$

For a description of the distribution from which these graphs are drawn, please see Definition 5.4.1.

The above bound is indicative of a fundamental limit on the spectral expansion of increasingly large graphs. [22] proved a matching negative result, which has become something of a "gold standard" in evaluating the quality of infinite expander families. Currently, the only known families that achieve the lower bound are algebraic constructions, such as the Ramanujan graphs of [20].

Theorem 2.3.2 ([22]). *For any infinite family $\{G_i\}$ of degree D graphs, $\lim_{i\to\infty}\lambda_2(G_i) \geq \frac{\sqrt{2D-1}}{D}$.*

Here we give a novel proof of the result. To analyze a bound on an infinite family of graphs, we need a property of the graph that goes to infinity as $N \to \infty$. In this case, the appropriate property is the diameter K of a graph.

Lemma 2.3.3. *For any degree D graph on N vertices, $K = \Omega(\log N)$.*

Proof. We prove this by showing a bound on N given D and K. Let i be one endpoint of the path whose length is K. We examine the worst case, which is when the graph is a tree rooted at i with maximal branching. In such a graph, at depth 0 from the root we have the single vertex i, at depth 1 we have at most D distinct children, and so on so that at depth n we have at most $D(D-1)^{n-1}$ vertices. The depth is at most K, so we have

$$N \leq 1 + D + D(D-1) + \ldots + D(D-1)^{K-1} = \frac{D(D-1)^K - 2}{D-2}$$

Treating D as constant and solving for K gives us the lemma. □

Proof of Theorem 2.3.2. We will show that, for any graph G, $\lambda_2(G) \geq \frac{2\sqrt{D-1}}{D} - O(\frac{1}{K})$. Here we assume *w.l.o.g.* that $\lambda_2(G) \geq 0$ since for any family $\{G_i\}$ we can bound its eigenvalues in the limit by analyzing the family $\{G_i^2\}$. Since any infinite family of graphs must have unbounded N, by Lemma 2.3.3 we have $N \to \infty \implies K \to \infty$ and so the theorem follows.

Let $e_1 = \{i_1, i_2\}$, $e_2 = \{j_1, j_2\}$ be two edges of G where the distance from either endpoint of e_1 to either endpoint of e_2 is at least $2k+2$. We use the distance between two edges instead of two vertices purely to simplify the notation of later expressions; it is easy to check from this definition that $k = \Theta(K)$ where K is the diameter of the graph.

The intuition behind the proof is that we can exhibit a disconnected induced subgraph H of G. Let H_I and H_J be the connected components of H, which by construction are centered about e_1 and e_2 respectively. Since H is disconnected, Remark 1.3.5 says that $\lambda_2^G(H)$ is at least the minimum of $\lambda_1^G(H_I), \lambda_1^G(H_J)$ the *largest* eigenvalues of H_I and H_J. Then, using the induced subgraph eigenvalue bound given in Corollary 2.1.6, we have that $\lambda_2(G) \geq \lambda_2^G(H)$.

For any set $S \subset V$, we define $C_r(S) \subset V$ to be the circle of radius r about S. For any vertex $i \in C_r(S)$, the distance of i from some vertex in S is exactly r, and from all other vertices in S the distance is at least r. Formally, we have:

$$C_r(S) = \{i \in V \mid \exists j_0 \in S \ s.t. \ d(i, j_0) = r \text{ and } \forall j \in S, \ d(i,j) \geq r\}$$

Let us define the set $I_r = C_r(\{i_1, i_2\})$. Define $I = \bigsqcup_{r=0}^{k} I_r$ the disjoint union of the I_r. $|I_{r+1}| \leq (D-1)|I_r|$ for all $0 \leq r \leq k-1$ because each vertex in I_r can have at most $D-1$ neighbors in the direction away from $\{i_1, i_2\}$. Let H_I be the induced subgraph with vertices in I.

We can lower-bound the largest eigenvalue $\lambda_1^G(H_I)$ of H_I by applying Lemma 2.1.7 without any restrictions on x (see Remark 2.1.3). To compute the lower bound, we adapt a construction invented in [22]. Note that $\lambda_1^G(H_I)$ is positive because the normalized adjacency matrix of H_I is non-negative. Also, note that H_I is necessarily non-regular since the vertices on the outermost circle I_k have degree $< D$, so its largest eigenvalue is strictly less than 1.

For convenience, we restrict $\|x\| = 1$. Let us construct x where $x_i = a(D-1)^{-r/2}$ iff $i \in I_r$, and where we pick $a \in \mathbb{R}$ such that $\|x\| = 1$. To summarize, x is constructed such that

$$1 = \|x\|^2 = a^2 \sum_{r=0}^{k} \frac{|I_r|}{(D-1)^r} \tag{2.8}$$

We know the normalized adjacency matrix of H_I has the form $B = R^t AR$, where A is the normalized adjacency matrix of G and R has the form shown in Lemma 2.1.6. Lemma 2.1.7 tells us

$$\lambda_1^G(H_I) \geq 1 - \frac{1}{D} \sum_{\{i,j\} \in E} ((Rx)_i - (Rx)_j)^2$$

There are at most $D - 1$ edges from any vertex $i \in I_r$ to vertices in I_{r+1}. Substituting our choice of x, and noticing that vertices in H_I are connected only to vertices in H_I except for the outermost circle I_k, we get

$$\lambda_1^G(H_I) \geq 1 - \frac{1}{D} \left(a^2 \left(\sum_{r=0}^{k-1} |I_r|(D-1) \left(\frac{1}{(D-1)^{r/2}} - \frac{1}{(D-1)^{(r+1)/2}} \right)^2 + |I_k| \frac{D-1}{(D-1)^k} \right) \right)$$

Note that $\frac{|I_r|}{(D-1)^r}$ never increases with increasing r, which means $\frac{1}{k+1} \frac{|I_k|}{(D-1)^k} \leq \sum_{r=0}^{k} \frac{|I_r|}{(D-1)^r}$. Using this fact and a little manipulation, we get

$$\lambda_1^G(H_I) \geq 1 - \frac{1}{D} \left((D - 2\sqrt{D-1}) \cdot \left(a^2 \sum_{r=0}^{k} \frac{|I_r|}{(D-1)^r} \right) + \frac{2\sqrt{D-1} - 1}{k+1} \cdot \left(a^2 \sum_{r=0}^{k} \frac{|I_r|}{(D-1)^r} \right) \right)$$

Applying Equation (2.8) gives

$$\lambda_1^G(H_I) \geq 1 - \frac{1}{D} \left((D - 2\sqrt{D-1}) \cdot 1 + \frac{2\sqrt{D-1} - 1}{k+1} \cdot 1 \right) = \frac{2\sqrt{D-1}}{D} + \frac{2\sqrt{D-1} - 1}{D(k+1)} \quad (2.9)$$

We may construct a graph H_J similarly starting with $\{j_1, j_2\}$, and an identical analysis follows to get the same bound on $\lambda_1^G(H_J)$. The graphs H_I and H_J have no edges between them because the distance between edges e_1, e_2 is at least $2k+2$. So the induced subgraph H on vertices $I \cup J$ is disconnected with connected components H_I and H_J. By Remark 1.3.5 we know since H is disconnected that $\lambda_2^G(H) \geq \min\{\lambda_1^G(H_I), \lambda_1^G(H_J)\}$. Corollary 2.1.6 tells us that, because of interlacing, we have $\lambda_2(G) \geq \lambda_2^G(H)$. Applying Equation (2.9) above and using $k = \Theta(K)$, we have $\lambda_2(G) \geq \frac{2\sqrt{D-1}}{D}$ as $K \to \infty$. □

2.4 Application in Randomness Reduction

As we noted in the introduction, expander graphs are useful in many areas of computer science. They have been applied in constructing efficient computer networks [8], in the study of complexity [28], and in generating error-correcting codes [26]. More recently, they have found a prominent place in the study of derandomization and pseudorandomness. In particular, they are useful in reducing the amount of randomness required for probabilistic algorithms.

Probabilistic algorithms have proven extremely useful in solving diverse classes of problems; many protocols such as zero-knowledge proofs, any secure cryptosystem, and many others could not exist without randomness. These algorithms also offer massive performance gains in some cases when deterministic algorithms are known, such as primality testing [24, 1].

However, the analysis of these algorithms usually depends on a uniform, unbiased source of random bits, which is difficult to implement. True physical randomness is hard to capture both because natural sources tend to produce biased bits and because of the deterministic nature of the available hardware. Programmers usually resort to using pseudorandom generators (PRG's) as a substitute for truly random bits. The most popular family of PRG's are linear congruential generators, and in general it is unknown how well randomized algorithms perform when using bits from such a PRG (though Impagliazzo and Zuckerman [16] do present some positive results in this area). In some cases it has been shown that these generators are completely useless, for example in most cryptographic applications. The study of pseudorandomness is most concerned with finding ways to "amplify" random bits, and, ideally, to obviate the need for them at all.

One of the earliest examples of expanders being used in the context of derandomization was in the parallel sorting algorithm given by Ajtai, Komlós, and Szemerédi [2]. Indeed, their algorithm is remarkable in that they used expander graphs to transform a randomized algorithm into a wholy deterministic one. In general, such results are not known to hold, and expanders are used most often to reduce the number of random bits a randomized algorithm requires to achieve exponentially small error.

We present the results of using expanders to reduce the number of random bits to get exponentially small error for algorithms in the three most basic randomized complexity classes **RP**, **coRP**, **BPP**. For the following definitions, we take all languages to be a subset of $\{0,1\}^*$. If $x \in \{0,1\}^*$, we let $|x|$ denote the binary length of x. By $x \overset{R}{\leftarrow} X$, we mean take x to be an element of X chosen uniformly (here X must be finite).

2.4.1 Expander walks for RP and coRP

Definition 2.4.1. A language $L \subset \{0,1\}^*$ is said to be in **RP** if there exists a polynomial-time algorithm A and a polynomial r with the following properties: for every $x \in \{0,1\}^*$ and for $s \overset{R}{\leftarrow} S$, where $S = \{0,1\}^r$ and $r = r(|x|)$, we have that

1. If $x \in L$, then A accepts x with probability at least 1/2, i.e.

$$\Pr_{s \overset{R}{\leftarrow} S} [A(x,s) = 1] \geq 1/2$$

2. If $x \notin L$, then A always rejects x, i.e. $A(x,s) = 0$ for all $s \in \{0,1\}^r$.

As usual, the class **coRP** is defined as all languages L such that $\overline{L} \in$ **RP**. The error for algorithms in **RP** vanishes exponentially with repeated independent trials, since if $x \in L$, then

$$\Pr_{s_1 \ldots s_n \overset{R}{\leftarrow} S} [A(x,s_1) = 0 \wedge \ldots \wedge A(x,s_n) = 0] \leq 2^{-n}$$

An analogous result holds for **coRP**. However, note that it takes nr random bits in order to reduce the error exponentially in this fashion. Since randomness is a valuable resource, we would like to reduce this quantity. To do so, we use the following technique of taking a random walk on an expander to show that we can get a similar exponential reduction in error using only $r + O(n)$ random bits.

We state the technique for **RP**; the case for **coRP** is entirely similar. Let L be any **RP** language, and define A and r as above in Definition 2.4.1. Let us also suppose we have a graph $G = (V,E)$ that is a $(N = 2^r, D, \lambda)$-spectral expander for some fixed constant $\lambda < 1$ and where $D \ll N$. We perform a random walk of length n starting at a randomly chosen inital vertex, and we use each vertex that we hit in the walk as the r-bit random input to S. That is, we pick $s_1 \overset{R}{\leftarrow} V$

using some r random bits. Then at each step $i < n$ in the walk, we pick a random neighbor $s_{i+1} \overset{R}{\leftarrow} N(\{s_i\})$. Note that picking a random neighbor only takes $\log_2 D$ bits. Finally, we evaluate $A(x, \cdot)$ at each of the points $s_1 \dots s_n$, and if any of them accept, we accept x.

Theorem 2.4.2 ([15]). *Using the above technique, a random string of length $r + \log_2 D \cdot n$ gives us a probability of error $2^{-\Omega(n)}$.*

Proof. Let B be the set of bad vertices for a particular $x \in L$.

$$B = \{s \in V \mid A(x, s) = 0\}$$

Define $\alpha = \frac{|B|}{|V|} \leq 1/2$. We want to show that the probability of a random walk staying entirely in B goes down exponentially.

To do this, we define the restriction matrix R that restricts a vector to B. That is, if we let $R = [r_{ij}]$ where

$$r_{ij} = \begin{cases} 1, \, i = j \text{ and } i \in B \\ 0, \text{ else} \end{cases}$$

then $(Rv)_i > 0 \implies i \in B$. The probability that a walk of length n stays entirely in B is then $|R(AR)^n u|_1$, where u is our initial uniform vector. This is because R effectively removes the part of the distribution that falls outside of B, so applying it at each step keeps exactly the components of $(AR)^k u$ that come only from walks that have stayed entirely within B.

Noticing that R is idempotent (i.e. $R^2 = R$) and using Cauchy-Schwarz, we may write

$$|R(AR)^n u|_1 = |(RAR)^n u|_1 \leq \sqrt{N} \|(RAR)^n u\| \tag{2.10}$$

We may use the expansion of G to bound $\|(RAR)^n x\|$ for all x. Let us define $y = Rx$. Recall that we can separate $y = y^{\|} + y^{\perp}$ where, since y is non-negative, $y^{\|} = \frac{|y|_1}{N} u$ and $y^{\perp} \perp u$. We compute

$$\begin{aligned} |\langle RARx, x\rangle| &= |\langle ARx, Rx\rangle| \\ &= |\langle Ay, y\rangle| \\ &= |\langle A(y^{\|} + y^{\perp}), y^{\|} + y^{\perp}\rangle| \\ &= |\langle Ay^{\|}, y^{\|}\rangle + 2\langle Ay^{\|}, y^{\perp}\rangle + \langle y^{\perp}, y^{\perp}\rangle| \end{aligned}$$

We may apply Lemma 2.1.2 and the fact that $Ay^{\|} = y^{\|}$ to get that

$$\begin{aligned} |\langle RARx, x\rangle| &= |\langle Ay^{\|}, y^{\|}\rangle + \langle Ay^{\perp}, y^{\perp}\rangle| \\ &\leq \|y^{\|}\|^2 + \lambda \|y^{\perp}\|^2 \end{aligned}$$

Since $y^{\|} = \frac{|y|_1}{N} \cdot u$, we have $\|y^{\|}\| = \frac{|y|_1}{N} \cdot \|u\| = \frac{|y|_1}{\sqrt{N}}$. We can use Cauchy-Schwarz and the fact that y is non-negative to check

$$|y|_1 = \sum_{i=1}^{N} |y_i| \leq \|\chi_B\| \cdot \|y\| = \sqrt{\alpha N} \|y\|$$

This gives us $\|y^{\|}\| \leq \sqrt{\alpha} \cdot \|y\|$.

Applying this to the above, we have that

$$\begin{aligned}|\langle RARx, x\rangle| &\leq \|y^{\|}\|^2 + \lambda\|y^{\perp}\|^2 \\ &= \|y^{\|}\|^2 + \lambda(\|y\|^2 - \|y^{\|}\|^2) \\ &= \|y^{\|}\|^2(1-\lambda) + \lambda\|y\|^2 \\ &\leq \alpha\|y\|^2(1-\lambda) + \lambda\|y\|^2 \\ &= (\alpha + \lambda(1-\alpha))\|x\|^2\end{aligned}$$

where the last line applies $\|y\|^2 \leq \|x\|^2$.

Since the above bound holds for all x, we may apply Lemma 2.1.2 without the restriction that $x \perp u$, to get that $\max_x \frac{|\langle RARx,x\rangle|}{\langle x,x\rangle} = \max_x \frac{\|RARx\|}{\|x\|}$. This implies

$$\|RARx\| \leq (\alpha + \lambda(1-\alpha))\|x\|$$

for all x. Applying this to Equation (2.10), we obtain

$$\begin{aligned}\|(RAR)^n u\| &\leq (\alpha + \lambda(1-\alpha))^n\|u\| \\ |(RAR)^n u|_1 &\leq (\alpha + \lambda(1-\alpha))^n\end{aligned}$$

Since $\lambda < 1$ and $\alpha \leq \frac{1}{2}$, it follows that

$$\|(RAR)^n u\| \leq 2^{-\Omega(n)}$$

□

2.4.2 Expander Walks for BPP

The classes **RP** and **coRP** are rather restrictive since they do not allow for decision algorithms that may err in both directions. We define the class **BPP** that allows for decision algorithms that may err in both directions:

Definition 2.4.3. A language $L \subset \{0,1\}^*$ is said to be in **BPP** if there exists a polynomial-time algorithm A and a polynomial r with the following properties: for every $x \in \{0,1\}^*$ and for $s \overset{R}{\leftarrow} S$, where $S = \{0,1\}^r$ and $r = r(|x|)$, we have that

1. If $x \in L$, then A accepts x with probability at least 2/3, i.e.

$$\Pr_{s \overset{R}{\leftarrow} S}[A(x,s) = 1] \geq 2/3$$

2. If $x \notin L$, then A accepts x with probability at most 1/3, i.e.

$$\Pr_{s \overset{R}{\leftarrow} S}[A(x,s) = 1] \leq 1/3$$

We use a Chernoff bound to show that repeated independent samples and taking a majority vote produces exponentially small error. The Chernoff bound gives us an exponentially small of deviating away from the mean when taking the average of an increasingly large number of independent samples.

Proposition 2.4.4. *For any language $L \in$ **BPP** with decision algorithm A and for any string x, let $V(x, s_1 \ldots s_n)$ be the majority vote of $A(x, s_1) \ldots A(x, s_n)$. Then*

$$\Pr_{s_1 \ldots s_n \overset{R}{\leftarrow} S}[V(x, s_1 \ldots s_n) \text{ incorrect}] \leq 2^{-\Omega(n)}$$

We begin by proving the following Chernoff bound.

Lemma 2.4.5. *Let $X_1, \ldots, X_n$ be independent $\{0,1\}$ random variables, and where* $\Pr[X_i = 1] = p$ *for all i. Let $\overline{X} = \frac{1}{n}\sum_{i=1}^{n} X_i$ be their mean, then*

$$\Pr[\overline{X} > p + \varepsilon] \leq e^{-2\varepsilon^2 n}$$

Proof. We raise e to the power of the above quantities multiplied by an optimization factor γ to be set later, and then apply Markov's inequality and use independence to get the following:

$$\begin{aligned}\Pr[e^{\gamma \overline{X}} > e^{\gamma(p+\varepsilon)}] &\leq e^{-\gamma(p+\varepsilon)} E[e^{\gamma \overline{X}}] \\ &= e^{-\gamma(p+\varepsilon)}(1 - p + pe^{\gamma/n})^n \\ &= e^{-\gamma\varepsilon}(e^{-\frac{\gamma p}{n}}(1-p) + pe^{\frac{\gamma(1-p)}{n}})^n\end{aligned}$$

Applying the fact from analysis that $e^{-\frac{\gamma p}{n}}(1-p) + pe^{\frac{\gamma(1-p)}{n}} \leq e^{(\gamma/n)^2/8}$ for $0 \leq p \leq 1$, and solving for the γ that minimizes the exponent, we get $\gamma = 4\varepsilon n$, and substituting in gives us

$$\begin{aligned}\Pr[\overline{X} > p + \varepsilon] &\leq e^{-4\varepsilon^2 n + 2\varepsilon^2 n} \\ &= e^{-2\varepsilon^2 n}\end{aligned}$$

□

Proof of Proposition 2.4.4. We examine the case for input $x \notin L$; the case for $x \in L$ follows analogously from a symmetric Chernoff bound. Let $p = \Pr_{s \overset{R}{\leftarrow} S}[A(x,s) = 1] \leq 1/3$, then from Lemma 2.4.5 we have that

$$\begin{aligned}\Pr[V(x, s_1 \ldots s_n) \text{ incorrect}] &= \Pr\left[\frac{1}{n}\sum_{i=1}^{n} A(x, s_i) \geq 1/2\right] \\ &\leq e^{-n/18} \\ &= 2^{-\Omega(n)}\end{aligned}$$

□

Again, we use nr random bits to achieve this error reduction. As with the **RP** and **coRP** case, using a random walk on an expander gives us an improved efficiency of $r + O(n)$. Gillman [13] showed that this randomness-efficient procedure gives a similar exponentional reduction in error similar to that for **RP**.

Theorem 2.4.6 ([13]). *For any language $L \in$ **BPP** with decision algorithm A and any string x, we may use a random walk on an expander to produce strings $s_1 \ldots s_n$, and output the majority vote $V(x, s_1 \ldots s_n)$. Then the probability of deciding x in error is $2^{-\Omega(n)}$.*

Proof Sketch. We sketch a simplified version of Gillman's proof for the case $x \notin L$. The proof has a similar flavor to the proof of the Chernoff bound in Lemma 2.4.5, where we begin by raising the terms to an exponent and introduce an optimization factor γ, employ Markov's inequality to get a bound, and then use independence and some analysis result to optimize the resulting bound.

Define $t(s_1 \dots s_n) = \sum_{i=1}^{n} A(x, s_i)$ to be the time spent inside the set B of "bad vertices" (where $A(x, s) = 1$). We wish to bound

$$\begin{aligned} \Pr[V(x, s_1 \dots s_n) \text{ incorrect}] &= \Pr[t(s_1 \dots s_n) \geq n/2] && (2.11) \\ &= \Pr[e^{\gamma t} \geq e^{\gamma n/2}] && (2.12) \\ &\leq e^{-\gamma n/2} E[e^{\gamma t}] && (2.13) \end{aligned}$$

where γ is the optimization factor and the last line follows from Markov's inequality. The probabilities are taken over $s_1 \dots s_n$, which are taken from a random walk. We notice that $E[e^{\gamma t}]$ is the average of $e^{\gamma t(s_1 \cdots s_n)}$ over all walks $s_1 \dots s_n$. To express it more concretely, we define a perturbation matrix $A(\gamma) = [\alpha_{ij}]$ of the normalized adjacency matrix A with entries

$$\alpha_{ij} = \begin{cases} e^{\gamma} a_{ij},\ i = j, i \in B \\ \quad a_{ij}, \text{ else} \end{cases}$$

From this definition, we can show that $E[e^{\gamma t}] = |A(\gamma)^n u|_1$. Consider $|A^n u|_1$, which is the sum of the probabilities all length n walks. The probability of a walk is the product of the chain of a_{ij} defining the walk. But taking $|A(\gamma)^n u|_1$ gives us an extra factor of e^{γ} each time the walk hits an element of B, and so summing over all walks we get $E[e^{\gamma t}]$.

Let $\mu(\gamma)$ be the largest eigenvalue of $A(\gamma)$. Cauchy-Schwarz gives us $|A(\gamma)^n u|_1 \leq \sqrt{N} \|A(\gamma)^n u\|$, and we know $\|A(\gamma)^n u\| \leq \frac{\mu(\gamma)^n}{\sqrt{N}}$. Applying this to Equation (2.13) gives us the following:

$$\begin{aligned} \Pr[t(s_1 \dots s_n) \geq n/2] &\leq e^{-\gamma n/2} |A(\gamma)^n u|_1 \\ &\leq e^{-\gamma n/2} \mu(\gamma)^n \\ &= e^{-n(\gamma/2 - \log \mu(\gamma))} \end{aligned}$$

Here, as in the proof of the generic Chernoff bound, we require a result from analysis to bound the exponent. In this case, a much more difficult result is needed, and [13] uses some deep analysis and linear algebra results to show that $\log \mu(\gamma) \leq \gamma/3 + 5\gamma^2/(1 - \lambda_2(G))$. Using this fact, we get that

$$\Pr[t(s_1 \dots s_n) \geq n/2] \leq e^{-n(\gamma/6 - 5\gamma^2/(1-\lambda_2(G)))}$$

We may choose γ to maximize $\gamma/6 - 5\gamma^2/(1 - \lambda_2(G))$, which gives us

$$\begin{aligned} \Pr[t(s_1 \dots s_n) \geq n/2] &\leq e^{-n \frac{(1-\lambda_2(G))}{720}} \\ &= 2^{-\Omega(n)} \end{aligned}$$

This shows that $\forall x \notin L$, the probability of error is $2^{-\Omega(n)}$. A similar proof shows the analogous result $\forall x \in L$. □

Chapter 3

Constructing Expanders

In this chapter we examine two constructions of expander graph families. The classical constructions of expanders are drawn from algebra (some examples may be found in [9]), notably the study of linear and cyclic groups. Their analyses usually draw on advanced algebra results, and sometimes obscure the relationship between the graph's structure and its expansion. To give a taste of the somewhat esoteric techniques used in the proof of expansion for algebraic expander constructions, we first present the Gabber-Galil Construction. This will contrast with the Zig-Zag Product Construction, presented later in this chapter, whose analysis meshes more closely with intuition.

3.1 The Gabber-Galil Construction

The original construction [12] gave a family of bipartite expander graphs. Here we modify the construction to give a family of non-bipartite expanders with bounded second eigenvalue. We base our proof on the technique introduced by Jimbo and Maruoka [17].

Theorem 3.1.1 ([12, 17]). *Let $G_n = \{V_n, E_n\}$. Let $\mathbb{Z}_n$ be the ring of integers modulo n, and set $V = \mathbb{Z}_n \times \mathbb{Z}_n$. We define any $(x, y) \in V$ to be connected to the vertices $\rho_i(x, y)$ with ρ_i as defined below:*

$$\begin{array}{llll} \rho_1(x,y) = (x, y+2x) & \rho_2(x,y) = (x, y+2x+1) & \rho_3(x,y) = (x, y-2x) & \rho_4(x,y) = (x, y-2x-1) \\ \rho_5(x,y) = (x+2y, y) & \rho_6(x,y) = (x+2y+1, y) & \rho_7(x,y) = (x-2y, y) & \rho_8(x,y) = (x-2y-1, y) \end{array}$$

Then $\lambda_2(G_n) \leq \frac{5\sqrt{2}}{8}$ for all n, and $\{G_n\}$ is a family of $(n^2, 8, \frac{5\sqrt{2}}{8})$-spectral expanders.

It is difficult to see intuitively why these graphs are good expanders. One interpretation might be that the permutations ρ_i operate on the two coordinates "independently", and so it is easy to get from one coordinate pair to another. In addition, the +1 introduced in ρ_2, ρ_6 make sure we avoid problems with parity, since otherwise even coordinates would only go to other even coordinates. However, even this very weak intuition is not used in the proof of expansion, which directly attacks the problem of bounding the eigenvalue using analytic and algebraic techniques.

Because the proof does not follow a clear intuition, we outline the structure here. First, we use the Fourier transform to find a matrix A' that is similar to A and study it instead. This facilitates the analysis because the Fourier transform has nice properties when applied to the matrices of the above permutations ρ_i. In particular, A' has a distinct submatrix for which we need only analyze the *largest* eigenvalue instead of the second largest, and it is this submatrix which admits to the $\frac{5\sqrt{2}}{8}$ bound.

For this proof we will use much of the algebra terminology defined in Subsection 1.2.3.

For convenience and for simplicity of notation, for the duration of this section we will adopt the convention of indexing $n^2 \times n^2$ matrices by coordinates $i, j \in \mathbb{Z}_n^2$. As usual, i indexes the row and j indexes the column. For example, if $n = 10$ and A is a $n^2 \times n^2$ matrix, then the entry $((0,5),(9,2))$ is in row $(0,5)$ and in column $(9,2)$.

We will extend the notation of $\lambda_i(B)$ to also refer to the i'th largest eigenvalue in absolute value of some matrix B.

Finally, let $\omega = \exp(2\pi\sqrt{-1}/n)$, and let $\mathbf{0} = (0,0)$.

We begin by defining the Fourier transform we use.

Definition 3.1.2. Let $\mathfrak{F}$ be the normalized matrix of the discrete two-dimensional Fourier transform from $\mathbb{C}^{n\times n} \to \mathbb{C}^{n\times n}$. To represent it in two dimensions, we consider the matrix in $\mathfrak{F} \in M_{n^2}(\mathbb{C})$ where $\mathfrak{F} = [\mathfrak{f}_{ij}]$ is given by

$$\mathfrak{f}_{ij} = \frac{1}{n}\omega^{\langle i,j\rangle}$$

where $\langle i, j\rangle$ is the standard dot product of $i, j \in \mathbb{Z}_n^2$.

Lemma 3.1.3. *For any graph G with normalized adjacency matrix A, we have that*

$$\mathfrak{F}A\mathfrak{F}^* = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & & & \\ \vdots & & H & \\ 0 & & & \end{bmatrix}$$

where H is a $(n^2-1)\times(n^2-1)$ hermitian matrix. Furthermore, $\lambda_2(G)$ is equal to $\lambda_1(H)$ the largest eigenvalue of H in absolute value.

Proof. Since A is symmetric and real, it is hermitian, and therefore $\mathfrak{F}A\mathfrak{F}^*$ is hermitian, so H must be hermitian.

Let $\mathfrak{F}A\mathfrak{F}^* = [b_{ij}]$, then we have

$$b_{ij} = \sum_{k,\ell} \mathfrak{f}_{ik} a_{k\ell} \bar{\mathfrak{f}}_{j\ell}$$

From the definition of $\mathfrak{F}$, we have $\mathfrak{f}_{\mathbf{0},k} = \frac{1}{n}$ for any k. Recall that $\sum_i a_{ij} = \sum_i a_{ji} = 1$ for any j, so we get $b_{\mathbf{0},\mathbf{0}} = \frac{1}{n^2} \cdot \sum_{k,\ell} a_{k,\ell} = 1$. Using the same facts we can also compute for $j \neq \mathbf{0}$ that

$$b_{\mathbf{0},j} = \sum_{k,\ell} \mathfrak{f}_{\mathbf{0},k} a_{k,\ell} \bar{\mathfrak{f}}_{j,\ell} = \sum_{\ell} \mathfrak{f}_{\mathbf{0},\ell} \bar{\mathfrak{f}}_{j,\ell} = [\mathfrak{F}\mathfrak{F}^*]_{\mathbf{0},j} = 0$$

where the last equality follows because of the well-known fact that $\mathfrak{F}$ is unitary. $b_{j,\mathbf{0}}$ is computed similarly, and so $\mathfrak{F}A\mathfrak{F}^*$ has the form stated. Also, $\mathfrak{F}A\mathfrak{F}^{-1} = \mathfrak{F}A\mathfrak{F}^*$ since $\mathfrak{F}$ is unitary, so $\mathfrak{F}A\mathfrak{F}^*$ is similar to A. Therefore they have the same eigenvalues. Since we already have an eigenvalue of 1 on the diagonal (with the first standard basis e_1 as eigenvector), the largest eigenvalue of H in absolute value must equal the second largest eigenvalue of A. □

We will use the following two lemmas to analyze the largest eigenvalue of the submatrix H. In both lemmas we implicitly use the fact that if B is a real symmetric matrix with non-negative entries, then its largest eigenvalue non-negative.

Lemma 3.1.4. *Let $H \in M_n(\mathbb{C})$ be hermitian, and let $B \in M_n(\mathbb{R})$ be non-negative and symmetric. If $|h_{ij}| \leq b_{ij}$ for all i, j, then for $\lambda_1(H)$ the largest eigenvalue of H in absolute value, we have $\lambda_1(H) \leq \lambda_1(B)$ the largest eigenvalue of B.*

Proof. By Remark 2.1.3, we can extend Lemma 2.1.2 to compute the largest eigenvalue of complex hermitian matrices, which gives

$$\lambda_1(H) = \max_{\|x\|=1} |\langle Hx, x\rangle|$$

We apply this to get

$$\lambda_1(H) = \max_{\|x\|=1} |\langle Hx, x\rangle| = \max_{\|x\|=1} \left| \sum_{i,j} h_{ij}\bar{x}_i x_j \right| \leq \max_{\|x\|=1} \sum_{i,j} |h_{ij}||\bar{x}_i||x_j| \leq \max_{\|x\|=1} \sum_{i,j} b_{ij} x_i x_j = \lambda_1(B)$$

where the last inequality holds since $\max \sum b_{ij} x_i x_j$ occurs on a vector x with non-negative entries. This is because we have removed the restriction that $x \perp u$. □

Lemma 3.1.5. *Let $B \in M_n(\mathbb{R})$ be symmetric, and where $b_{ij} \geq 0$ if $i \neq j$. Let $\lambda_1(B)$ be the largest eigenvalue of B. Take d_{ij} for $1 \leq i, j \leq n$ to be arbitrary real numbers that for all i, j have $d_{ij} > 0$ and $d_{ij} = \frac{1}{d_{ji}}$. Then $\lambda_1(B) \leq \max_i \sum_j b_{ij} d_{ij}$.*

Proof. Let x be the eigenvector of B with maximal eigenvalue $\lambda_1(B)$, normalized so that $\|x\| = 1$. Simple computation reveals that

$$\lambda_1(B) = \langle Bx, x\rangle = \sum_{i,j} b_{ij} x_i x_j = \frac{1}{2} \sum_{i,j} 2 b_{ij} x_i x_j$$

Using the simple fact that, for any $\alpha, \beta \in \mathbb{R}$ and any $d > 0$, we have $(\sqrt{d}\alpha - \frac{\beta}{\sqrt{d}})^2 = d\alpha^2 + \frac{\beta^2}{d} - 2\alpha\beta \geq 0$, which implies $2\alpha\beta \leq d\alpha^2 + \frac{\beta^2}{d}$, we get

$$\begin{aligned}
\langle Bx, x\rangle &\leq \frac{1}{2} \sum_{i,j} b_{ij}(d_{ij} x_i^2 + d_{ji} x_j^2) \\
&\leq \frac{1}{2} \sum_{i,j} (b_{ij} d_{ij} x_i^2 + b_{ji} d_{ji} x_j^2) \\
&= \sum_{i,j} b_{ij} d_{ij} x_i^2 \\
&\leq \left(\max_i \sum_j b_{ij} d_{ij} \right) \left(\sum_i x_i^2 \right) \\
&= \max_i \sum_j b_{ij} d_{ij}
\end{aligned}$$

□

Now we consider how $\mathfrak{F}$ acts on certain permutations of $\mathbb{Z}_n^2$, the vertex set of the graph. For any matrix in $B \in M_2(\mathbb{Z}_n)$, we denote the (linear) permutation $\mathbb{Z}_n^2 \to \mathbb{Z}_n^2$ defined by B by where $\sigma_B(i) = Bi$. Likewise, for any $t \in \mathbb{Z}_n^2$, we denote the (translation) permutation $\mathbb{Z}_n^2 \to \mathbb{Z}_n^2$ defined by t by $\sigma_t(i) = i + t$. We will use $\mathcal{M}$ to denote the group of affine permutations generated by permutations of the form σ_B and σ_t.

Let $P(\sigma) \in M_{n^2}(\{0,1\})$ denote the permutation matrix of any permutation on $\mathbb{Z}_n^2$. More formally, taking $\delta(i,j)$ be the Kronecker delta, we may define $P(\sigma)$ to have entries $\delta(i,\sigma(j))$. It is clear that $P(\sigma\sigma') = P(\sigma)P(\sigma')$ and that $P(\sigma^{-1}) = P(\sigma)^t$.

Lemma 3.1.6. *Let $B \in M_2(\mathbb{Z}_n)$ be an invertible matrix. Then we have*

$$\mathfrak{F}P(\sigma_B)\mathfrak{F}^* = P(\sigma_{(B^{-1})^t})$$

Proof. Using the definition of $\mathfrak{F}$ and $P(\sigma_B)$, we get that

$$\begin{aligned}
[\mathfrak{F}P(\sigma_B)\mathfrak{F}^*]_{ij} &= \frac{1}{n^2}\sum_{k,\ell}\omega^{\langle i,k\rangle}\delta(k,B\ell))\omega^{-\langle j,\ell\rangle} \\
&= \frac{1}{n^2}\sum_{\ell}\omega^{\langle i,B\ell\rangle-\langle j,\ell\rangle} \\
&= \frac{1}{n^2}\sum_{\ell}\omega^{\langle B^t i-j,\ell\rangle}
\end{aligned}$$

It is well-known that the Fourier basis functions are orthogonal in the one-dimensional case, and it is a simple exercise to show the same in the two-dimensional case we employ here. Therefore, the above summation term is n^2 when $B^t i = j$ and 0 otherwise, so this gives us

$$[\mathfrak{F}P(\sigma_B)\mathfrak{F}^*]_{ij} = \delta(i,\sigma_{(B^{-1})^t}(j))$$

which immediately implies the lemma. □

Lemma 3.1.7. *Let $t \in \mathbb{Z}_n^2$. Then $\mathfrak{F}P(\sigma_t)\mathfrak{F}^*$ is a diagonal matrix, where the i'th entry is $\omega^{\langle t,i\rangle}$.*

Proof. Using the definition of $\mathfrak{F}$ and $P(\sigma_t)$, we get that

$$\begin{aligned}
[\mathfrak{F}P(\sigma_t)\mathfrak{F}^*]_{ij} &= \frac{1}{n^2}\sum_{k,\ell}\omega^{\langle i,k\rangle}\delta(k,\ell+t)\omega^{-\langle j,\ell\rangle} \\
&= \frac{1}{n^2}\sum_{\ell}\omega^{\langle i,\ell+t\rangle-\langle j,\ell\rangle} \\
&= \frac{\omega^{\langle t,i\rangle}}{n^2}\sum_{\ell}\omega^{\langle i-j,\ell\rangle} \\
&= \omega^{\langle t,i\rangle}\cdot\delta(i,j)
\end{aligned}$$

□

Proof of Theorem 3.1.1. Let A be the normalized adjacency matrix of G_n as in the statement of the theorem. A is real and symmetric, and hence hermitian. Note that we can decompose the normalized adjacency matrix of our graph into $A = \frac{1}{8}\sum_{i=1}^{8} P(\rho_i)$. It is clear $\rho_i \in \mathcal{M}$ for all i. Furthermore, if we define matrices B_1, B_2 and vectors t_1, t_2 as follows:

$$B_1 = \begin{bmatrix}1 & 2\\ 0 & 1\end{bmatrix}, B_2 = \begin{bmatrix}1 & 0\\ 2 & 1\end{bmatrix}, t_1 = \begin{bmatrix}1\\ 0\end{bmatrix}, t_2 = \begin{bmatrix}0\\ 1\end{bmatrix}$$

then we may rewrite the ρ_i as products of the permutations σ_{B_i} and σ_{t_i}, giving us

$$A = \frac{1}{8}\sum_{i=1}^{2}((I + P(\sigma_{t_i}))P(\sigma_{B_i}) + (I + P(\sigma_{t_i}^{-1}))P(\sigma_{B_i}^{-1})) \tag{3.1}$$

We now take $\mathfrak{F}A\mathfrak{F}^*$, which has the form as guaranteed in Lemma 3.1.3 with a hermitian submatrix H. To apply Lemma 3.1.4, we define the real symmetric $(n^2-1) \times (n^2-1)$ matrix $C = [c_{ij}]$ in relation to A. Here, for the sake of simplifying notation without sacrificing correctness, we take the unusual step of indexing C and H with indices in $\mathbb{Z}_n^2$ without the first index, which we write as $\mathbb{Z}_n^2 \setminus \{\mathbf{0}\}$. This is because we wish to define the entries of C in terms of entries of A, but we wish to ignore the first row and column of A because C has smaller dimensions. With this clarification, we define C as follows.

$$c_{ij} = \frac{1}{8}\left(|1+\omega^{i_1}| \cdot \left(\delta(i, \sigma_{B_2}(j)) + \delta(i, \sigma_{B_2}^{-1}(j))\right) + |1+\omega^{i_2}| \cdot \left(\delta(i, \sigma_{B_1}(j)) + \delta(i, \sigma_{B_1}^{-1}(j))\right)\right)$$

where i_1, i_2 are the components of i.

Using the facts that $B_1 = B_2^t$ and $\langle t_j, i\rangle = i_j$, we can check by applying applying Lemmas 3.1.6 and 3.1.7 to the decomposition of A in Equation (3.1) that in fact $c_{ij} \geq |h_{ij}|$ for all $i, j \in \mathbb{Z}_n^2 \setminus \{\mathbf{0}\}$.

So by Lemma 3.1.4, it suffices to find the largest eigenvalue of C. We do so by applying Lemma 3.1.5. We define a sequence d_{ij} for $1 \leq i, j \leq n-1$ where $d_{ij} = \frac{1}{d_{ji}}$.

To do so, let us first specify the function $\phi : \mathbb{Z}_n^2 \setminus \{\mathbf{0}\} \to \mathbb{R}$ given by

$$\phi(i) = \cos((2\pi/n)i_1) + \cos((2\pi/n)i_2)$$

where i_1, i_2 are the components of i. Now we take the sequence d_{ij} to be

$$d_{ij} = \begin{cases} \frac{1}{\sqrt{2}},\ \phi(i) > \phi(j) \\ 1,\ \phi(i) = \phi(j) \\ \sqrt{2},\ \phi(i) < \phi(j) \end{cases}$$

This sequence is useful because of the following claim.

Claim. Let $\mathcal{S} = \{\sigma_{B_1}, \sigma_{B_2}, \sigma_{B_1}^{-1}, \sigma_{B_2}^{-1}\}$. For any $i \in \mathbb{Z}_n^2 \setminus \{\mathbf{0}\}$, define sets

$$\begin{aligned} U &= \{\sigma \mid \sigma \in \mathcal{S}, \phi(i) > \phi(\sigma(i))\} \\ V &= \{\sigma \mid \sigma \in \mathcal{S}, \phi(i) < \phi(\sigma(i))\} \end{aligned}$$

We have the following:

1. If $\phi(i) > 0$, then $|V| \leq 1$ and $|U| - |V| \geq 2$.
2. Otherwise, one of $|U| \geq 1$ or $|V| \leq 2$ is true.

We omit the details of this claim, which follow from straightforward but tedious and uninsightful trigonometric manipulations. The interested reader may refer to the appendix of the original work [17] for the specifics.

Returning to the application of Lemma 3.1.5, we recall

$$\lambda_1(C) \leq \max_i \sum_j d_{ij}c_{ij}$$

Since the c_{ij} are defined as a sum of δ's, it becomes evident why the above claim is interesting: it says that $\phi(i) > \phi(\sigma(i))$ occurs frequently enough so we can bound the above right-hand side to be strictly less than 1. Formally, if we pick any i, then if $\phi(i) > 0$, using $|1 + \omega^{i_k}| \leq 2$ and part 1 of the above claim we have

$$\sum_j d_{ij} c_{ij} \leq \frac{1}{4}\left(\sqrt{2} + \frac{3}{\sqrt{2}}\right) = \frac{5\sqrt{2}}{8}$$

For the case $\phi(i) \leq 0$, suppose *w.l.o.g.* that $|1 + \omega^{i_1}| \geq |1 + \omega^{i_2}|$. We get

$$\sum_j d_{ij} c_{ij} \leq \frac{1}{8}\left(\sqrt{2}(2|1 + \omega^{i_1}| + |1 + \omega^{i_2}|) + \frac{|1 + \omega^{i_2}|}{\sqrt{2}}\right) = \frac{\sqrt{2}}{8} \cdot \frac{4|1 + \omega^{i_1}| + 3|1 + \omega^{i_2}|}{2}$$

A simple calculation shows that for any real a, b, we have

$$(4a + 3b)^2 = 16a^2 + 9b^2 + 24ab \leq 16a^2 + 9b^2 + 16b^2 + 9a^2 = 25(a^2 + b^2)$$

Substituting in $a = \frac{|1+\omega^{i_1}|}{2}$ and $b = \frac{|1+\omega^{i_2}|}{2}$, we now claim that $a^2 + b^2 \leq 1$. This is because

$$\begin{aligned} a^2 + b^2 &= \frac{1}{4}(|1 + \omega^{i_1}|^2 + |1 + \omega^{i_2}|^2) \\ &= 1 + \frac{1}{2}\Re(\omega^{i_1} + \omega^{i_2}) \leq 1 \end{aligned}$$

where $\Re$ denotes the real part of a complex number. The last line above follows because

$$\Re(\omega^{i_1} + \omega^{i_2}) = \cos((2\pi/n)i_1) + \cos((2\pi/n)i_2) = \phi(i) \leq 0$$

Therefore $4a + 3b \leq 5$, and we get

$$\sum_j d_{ij} c_{ij} \leq \frac{5\sqrt{2}}{8}$$

□

The construction here misses the optimal bound for infinite families given in Theorem 2.3.2 by a factor of approximately $\sqrt{2}$. Lubotzky, Phillips, and Sarnak [20] and Margulis [21] independently gave a Cayley graph (see Definition 4.1.1) construction based on $PSL(2, q)$ that can be shown to reach the optimal eigenvalue bound using the Ramanujan conjecture and Kazhdan's Property T. [20] coined the term "Ramanujan graphs" for all families of graphs that reach the optimal eigenvalue bound. However, we will not elaborate on the constructions of [20, 21] because of the depth of the algebraic analysis required. As our goal is to show the progression of expander constructions from algebraic to combinatorial, we now turn to the zig-zag product.

3.2 The Zig-Zag Product Construction

3.2.1 Definition and Intuition

The zig-zag product was introduced by Reingold, Vadhan, and Widgerson [25], and offers an elegant yet powerful recursive and combinatorial construction of expander families given a good expander graph of fixed size. It depends on a fixed (but arbitrary) labelling of the edges, which we define presently. In the next chapter, we present new results that indicate that this labelling has little effect on the eigenvalues of the resulting graph.

Definition 3.2.1. Let $G = (V, E)$ be a graph with degree D. For each $i \in V$, we fix an ordering of its incident edges $\{i, j_1\}, \{i, j_2\}, \ldots, \{i, j_D\}$. We define a *rotation map* of G to be a map $\mathrm{Rot}_G : V \times [D] \to V \times [D]$ such that $\mathrm{Rot}_G(i, c) = (j, d)$ if and only if i is the d'th neighbor of j and j is the c'th neighbor of i.

Remark 3.2.2. Here we present a useful alternative formulation of rotation maps. Say that each edge $e \in E$ is colored with two colors in $[D]$, where the colors may be the same. Each color of e corresponds to an endpoint of e. We restrict the coloring such that no vertex has two incident edges colored the same at that endpoint. Then we may say that $\mathrm{Rot}_G(i, c) = (j, d)$ iff there is an edge $\{i, j\} \in E$ is colored c at endpoint i and is colored d at endpoint j.

In the ideal and simplest case, this would be an edge-coloring of the graph using D colors, such that $\mathrm{Rot}_G(i, c) = (j, c)$ for each edge $\{i, j\} \in E$. However, because it is **NP**-complete to compute whether a coloring with D colors even exists [19], we allow for arbitrary rotation maps.

With the concept of a rotation map, we are equipped to define the zig-zag product.

Definition 3.2.3. Let $G = (V_1, E_1)$ be a degree D_1 graph on N vertices, and let $H = (V_2, E_2)$ be a degree D_2 graph on D_1 vertices. Fix rotation maps Rot_G of G and Rot_H of H. Then the *zig-zag product* $\Gamma = G \textcircled{z} H$ is a degree D_2^2 graph on ND_1 vertices with vertex set $V = V_1 \times V_2$. The edges of Γ are defined by the rotation map

$$\mathrm{Rot}_\Gamma : (V_1 \times V_2) \times ([D_2] \times [D_2]) \to (V_1 \times V_2) \times ([D_2] \times [D_2])$$

given by $\mathrm{Rot}_\Gamma((i, c), (k, \ell)) = ((j, d), (m, n))$ where there exist some $p, q \in V_2$ such that

$$1.\quad \mathrm{Rot}_H(c, k) = (p, n) \qquad 2.\quad \mathrm{Rot}_G(i, p) = (j, q) \qquad 3.\quad \mathrm{Rot}_H(q, \ell) = (d, m)$$

We will commonly refer to G and H in the above definition as the "large graph" and the "small graph" respectively because of the following intuition. To grasp the reasoning behind the zig-zag product, it is useful to imagine a graph identical to G, but with each vertex replaced by a cloud of vertices, where each cloud is a copy of H. At the i'th cloud, the c'th vertex in H corresponds to the c'th edge of i, as given by Rot_G. Taking a random step in $G \textcircled{z} H$ corresponds to three steps in the smaller graphs:

1. First we take a random step from c to p in the H-cloud.
2. Then, according to the vertex labelling of H and the edge labelling of G, we follow the edge of G that corresponds to p, and arrive at the neighboring H-cloud. At the new H-cloud, the edge is labelled as q.
3. At this new H-cloud, we take another random step q to d.

Example 3.2.4. We refer to Figures 3.1 and 3.2. G is the complete graph on 5 vertices and H is the cyclic graph on 4 vertices. The zig-zag product graph $\Gamma = G \textcircled{z} H$ in Figure 3.2 illustrates the cloud analogy, where each of the clouds 1-5 are a copy of H. Starting at the c'th vertex of cloud 2, we take a step inside the cloud to d. d corresponds to the edge going between clouds to cloud 5, where it is labelled c. Finally, we take another step in cloud 5 to the d'th vertex. The solid line from $(2, c)$ to $(5, d)$ represents this edge. Other edges in Γ are similar.

To see why $G \textcircled{z} H$ should be a good expander if G and H are good expanders, let us consider any distribution x on the vertices. We may break x into the sum of two orthogonal parts $x^\perp$ and $x^\parallel$: $x^\perp$ is anti-uniform on each cloud, and $x^\parallel$ is uniform on each cloud. By this, we mean

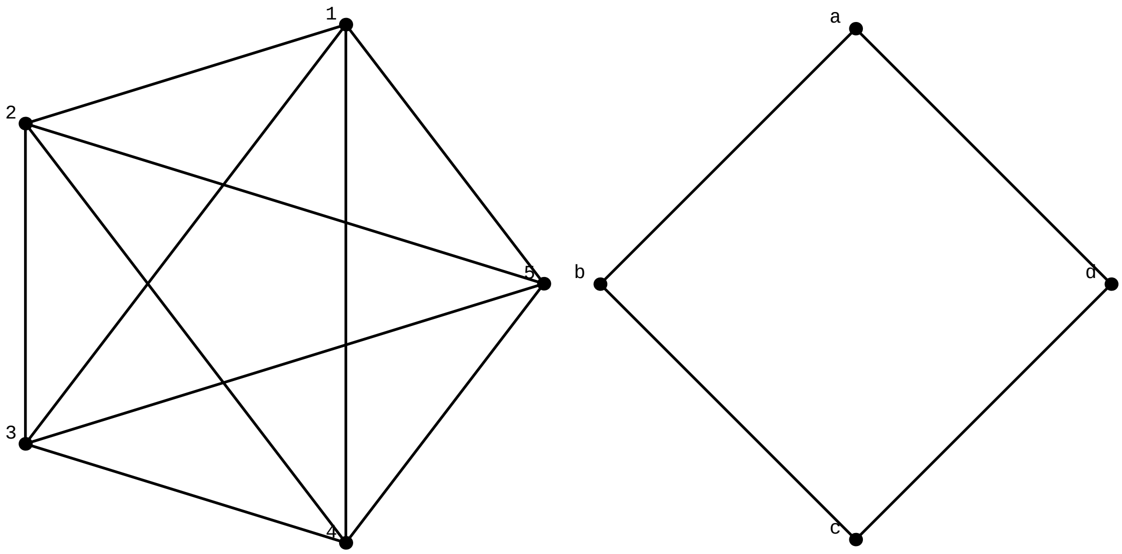

Figure 3.1: G and H

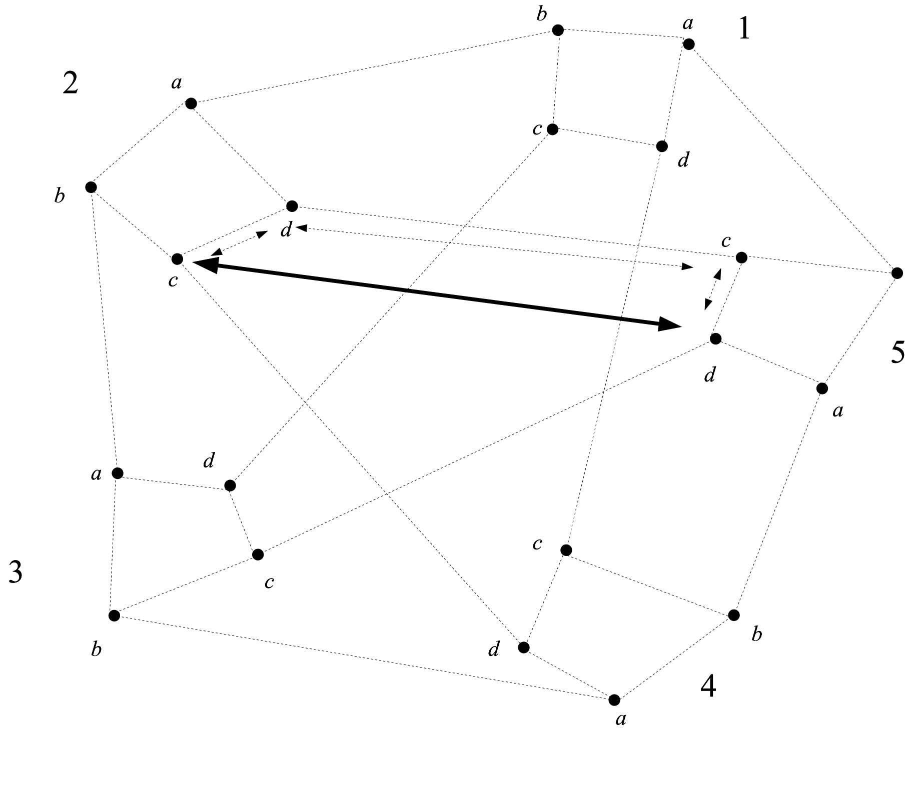

Figure 3.2: Γ

that if we look only at the components of $x^{\|}$ that correspond to a single cloud, they are uniform. Likewise, the components of $x^{\perp}$ belonging to a single cloud are anti-uniform.

Step 1 makes $x^{\perp}$ closer to uniform because it is analogous to taking a random step in H, and H is an expander. This is enough, since steps 2 and 3 do not make $x^{\perp}$ any less uniform. Step 1 has no effect on $x^{\|}$, but because step 2 is analogous to a step on G and because G is an expander, step 2 disperses the uniformity from each cloud and mixes them. When we get to step 3, we have that each cloud is far from uniform, so when we take a random step 3 (analogous to a random step on H), we move closer to uniform because H is an expander. Finally, because both components of x become more uniform, x itself becomes more uniform. We formalize this in the next section.

3.2.2 Constructing Expander Families

To facilitate the statement of the proof of expansion, we introduce tensor terminology.

Definition 3.2.5. Let A and B be real matrices. Then their matrix tensor product is $C = A \otimes B$, where $c_{(i,c),(j,d)} = a_{ij} \cdot b_{cd}$.

Furthermore, by treating vectors as $n \times 1$ matrices, we have an analogous definition of the tensor product of two vectors. Simple computation shows that

$$(A \otimes B)(x \otimes y) = (Ax) \otimes (By)$$

The key to constructing expander families using the zig-zag product is the following bound on the expansion of the product graph. The original work [25] presents two analyses of the zig-zag product, which we refer to as the "simple bound" and the "advanced bound". Our analysis here is better than the simple bound but worse than the advanced bound. However, it follows the simple bound in form, and does not involve any geometric arguments such as those employed in [25]'s proof of the advanced bound. Furthermore, this analysis carries over directly to the wide zig-zag product case (Definition 4.3.1), whereas it is unclear whether the advanced bound extends to the wide zig-zag product.

Theorem 3.2.6. *If G is a (N, D_1, λ)-spectral expander and H is a (D_1, D_2, μ)-spectral expander, then $G \text{ⓩ} H$ is a $(ND_1, D_2^2, f(\lambda, \mu))$-spectral expander where*

$$f(\lambda,\mu) \leq \frac{1}{2}(\lambda + \mu^2) + \frac{1}{2}|\lambda - \mu^2|\sqrt{1 + \left(\frac{2\mu}{\lambda - \mu^2}\right)^2}$$

Furthermore, $f(\lambda, \mu) < 1$ if $\lambda < 1$ and $\mu < 1$.

Despite the messiness of the above bound, by splitting up the radical using the triangle inequality, we can easily see that

$$f(\lambda,\mu) \leq \max\{\lambda + \mu,\ \mu^2 + \mu\}$$

Proof. We will apply the intuition developed in the previous subsection. First we define some notation. Let 1_n be the 1's column vector of length n. Let us index vectors $x \in \mathbb{R}^{ND_1}$ with two coordinates, $i \in [N]$ and $c \in [D_1]$. i indexes the which cloud we are in, and c indexes a vertex in the i'th cloud. For example, a tensor such as $x' = 1_N \otimes x/N$ means that $x'_{(i,c)} = \frac{1}{N}x_c$ for any i; this means that all clouds have equal weight, and the distribution within a cloud is identical on all clouds.

By x_i we shall mean the vector in $\mathbb{R}^{D_1}$ that consists of the entries of $x_{(i,c)}$ for all $c \in [D_1]$. That is, x_i is restriction of x to the i'th cloud.

Let us first define the *collection map* $C : \mathbb{R}^{ND_1} \to \mathbb{R}^N$, where the i'th entry of $C(v)$ is the sum of all the weights on the i'th cloud of v. Formally, for any vector $v \in \mathbb{R}^{ND_1}$ with entries $v_{(i,c)}$ we define

$$C(v)_i = \sum_{c=1}^{D_1} v_{(i,c)}$$

It is clear that C is linear.

We want to decompose x into two orthogonal parts, one that is uniform on each cloud and one that is anti-uniform on each cloud. We may use the collection map to define it as such:

$$\begin{aligned} x^{\|} &= C(x) \otimes 1_{D_1}/D_1 \\ x^{\perp} &= x - x^{\|} \end{aligned}$$

Using the above notation, we take $x_i^{\|}$ to be the restriction of $x^{\|}$ to the i'th cloud. It is clear that for all i, $x_i^{\|}$ is a multiple of the uniform distribution, and it immediately follows that $x_i^{\perp} \perp u$ for each i. These are exactly the vectors we discussed in the intuition for the zig-zag product.

Let A and B be the normalized adjacency matrices of G and H respectively. Let M be the normalized adjacency matrix of $G \textcircled{z} H$. Because each step in $G \textcircled{z} H$ is the composition of three steps in the subgraphs, we may decompose $M = B'A'B'$, where $B' = I_N \otimes B$ and A' is a permutation matrix where

$$a_{(i,c),(j,d)} = \begin{cases} 1, \, \mathrm{Rot}_G(i,c) = (j,d) \\ 0, \text{ else} \end{cases}$$

Note that both A' and B' are symmetric.

Using Lemma 2.1.2, we try to bound the maximum of $|\langle Mx, x\rangle| = |\langle B'A'B'x, x\rangle|$ where $x \in \mathbb{R}^{ND_1}$ and $x \perp u$. Applying our decomposition, we have

$$\lambda_2(G \textcircled{z} H) = \max_{x \perp u} \frac{|\langle B'A'B'(x^{\|} + x^{\perp}), x^{\|} + x^{\perp}\rangle|}{\langle x, x\rangle}$$

Notice that

$$B'x^{\|} = (I_N \otimes B)(C(x) \otimes 1_{D_1}/D_1) = C(x) \otimes 1_{D_1}/D_1 = x^{\|}$$

Using this substitution and the fact that B' is symmetric, we have

$$|\langle B'A'B'(x^{\|} + x^{\perp}), x^{\|} + x^{\perp}\rangle| = |\langle A'x^{\|} + A'B'x^{\perp}, x^{\|} + B'x^{\perp}\rangle|$$

Since A' is a permutation, it is length-preserving, and we may apply Cauchy-Schwarz to get

$$|\langle B'A'B'(x^{\|} + x^{\perp}), x^{\|} + x^{\perp}\rangle| \le |\langle A'x^{\|}, x^{\|}\rangle| + 2\|x^{\|}\| \cdot \|B'x^{\perp}\| + \|B'x^{\perp}\|^2$$

We now bound each of these terms individually.

Firstly, we have

$$\|B'x^{\perp}\|^2 = \sum_i \|Bx_i^{\perp}\|^2 \le \mu \sum_i \|x_i^{\perp}\|^2 = \mu \|x^{\perp}\|^2$$

Secondly, recall the definition of A'. Applying A' to $x^{\|}$ distributes the weight at $x_i^{\|}$ to the neighbors of i when considered as a vertex of G. This means that $C(A'e_i \otimes 1_{D_1}/D_1) = Ae_i$, and since $\{e_i\}$ forms a basis of $\mathbb{R}^N$, and $x^{\|} = C(x) \otimes 1_{D_1}/D_1$, this means we may write $C(A'x^{\|}) = AC(x)$. Hence we can check that

$$\begin{aligned}
\langle A'x^{\|}, x^{\|}\rangle &= \sum_{i,j} (A'x^{\|})_{(i,j)} \frac{C(x)_i}{D_1} \\
&= \frac{1}{D_1} \sum_{i=1}^{N} C(x)_i \sum_{j=1}^{D_1} (A'x^{\|})_{(i,j)} \\
&= \frac{1}{D_1} \sum_{i=1}^{N} C(x)_i C(A'x^{\|})_i \\
&= \frac{1}{D_1} \sum_{i=1}^{N} C(x)_i (AC(x))_i \\
&= \frac{\langle AC(x), C(x)\rangle}{D_1}
\end{aligned}$$

In the following derivation, we apply the collection map, the easily verifiable fact that $\forall v, w \in \mathbb{R}^N,\ \langle v \otimes 1_{D_1}, w \otimes 1_{D_1}\rangle = D_1\langle v, w\rangle$, and the expansion of G.

$$\begin{aligned}
|\langle A'x^{\|}, x^{\|}\rangle| &= \frac{|\langle AC(x), C(x)\rangle|}{D_1} \\
&\leq \lambda \frac{\langle C(x), C(x)\rangle}{D_1} \\
&= \lambda \langle x^{\|}, x^{\|}\rangle
\end{aligned}$$

Combining these facts with the above expression, we have

$$|\langle B'A'B'(x^{\|} + x^{\perp}), x^{\|} + x^{\perp}\rangle| \leq \lambda \|x^{\|}\|^2 + 2\mu \|x^{\|}\| \cdot \|x^{\perp}\| + \mu^2 \|x^{\perp}\|^2$$

Let θ be the angle between $x^{\|}$ and x, then we may substitute $\cos\theta = \frac{\|x^{\|}\|}{\|x\|}$ and $\sin\theta = \frac{\|x^{\perp}\|}{\|x\|}$ to get

$$\lambda_2(G \text{ⓩ} H) = \max_{x \perp u} \frac{|\langle Mx, x\rangle|}{\langle x, x\rangle} \leq \max_{\theta \in [0, \frac{\pi}{2}]} \lambda \cos^2\theta + 2\mu \cos\theta \sin\theta + \mu^2 \sin^2\theta$$

Let us define

$$g(\theta) = \lambda \cos^2\theta + 2\mu \cos\theta \sin\theta + \mu^2 \sin^2\theta$$

We wish to maximize g over all $\theta \in [0, \pi/2]$, so we take $g'(\theta)$ and find that θ achieves a critical point when

$$\tan 2\theta = \frac{2\mu}{\lambda - \mu^2} \overset{\text{def}}{=} \eta$$

There are two θ that achieve this (in opposite quadrants of the Cartesian plane). We claim that there exists a unique $\theta_0 \in [0, \pi/2]$ such that $\tan 2\theta_0 = \eta$, and where θ_0 satisfies $g''(\theta_0) < 0$, and therefore g achieves a maximum at θ_0. We can compute

$$g''(\theta) = -2(\lambda - \mu^2)\cos 2\theta - 4\mu \sin 2\theta$$

Now if $\lambda - \mu^2 > 0$, then $\eta > 0$ and we pick the unique $\theta_0 \in [0, \pi/4] \subset [0, \pi/2]$ satisfying $\tan 2\theta_0 = \eta$, which gives us $\cos 2\theta_0 > 0$ and $\sin 2\theta_0 > 0$, which implies $g''(\theta_0) < 0$. If $\lambda - \mu^2 < 0$, then $\eta < 0$ and we pick the unique $\theta_0 \in [\pi/4, \pi/2] \subset [0, \pi/2]$ satisfying $\tan 2\theta_0 = \eta$, which implies $\cos 2\theta_0 < 0$ and $\sin 2\theta_0 > 0$, again giving us $g''(\theta_0) < 0$.

It is a simple exercise to see that the other critical points are never maxima in either of the above cases, and that the endpoints of the range are never maxima. Therefore θ_0 is a global maximum in $[0, \pi/2]$. Using a little trigonometry, we get that

$$g(\theta) = \frac{1}{2}(\lambda + \mu^2) + \frac{1}{2}(\lambda - \mu^2)\cos 2\theta + \mu \sin 2\theta$$

In the case where $\eta > 0$, we substitute $\theta = \theta_0$ to get

$$\begin{aligned} g(\theta_0) &= \frac{1}{2}(\lambda + \mu^2) + \frac{1}{2}(\lambda - \mu^2)\frac{1}{\sqrt{1+\eta^2}} + \mu\frac{\eta}{\sqrt{1+\eta^2}} \\ &= \frac{1}{2}(\lambda + \mu^2) + \frac{1}{2}(\lambda - \mu^2)\frac{1 + \frac{4\mu^2}{(\lambda-\mu)^2}}{\sqrt{1+\eta^2}} \\ &= \frac{1}{2}(\lambda + \mu^2) + \frac{1}{2}(\lambda - \mu^2)\sqrt{1+\eta^2} \end{aligned}$$

If $\eta < 0$, a similar derivation gives us

$$g(\theta_0) = \frac{1}{2}(\lambda + \mu^2) + \frac{1}{2}(-\lambda + \mu^2)\sqrt{1+\eta^2}$$

This shows the first part of the theorem.

To show the last part of the theorem, we separate the two cases: when $\|x^\perp\| \leq \frac{1-\lambda}{3\mu}\|x\|$, we have

$$|\langle Mx, x\rangle| \leq \lambda\|x\|^2 + 2\mu\left(\frac{1-\lambda}{3\mu}\right)\|x\|^2 + \mu^2\left(\frac{1-\lambda}{3\mu}\right)^2\|x\| < \left(1 - \frac{1-\lambda}{9}\right)\|x\|^2$$

If $\|x^\perp\| > \frac{1-\lambda}{3\mu}\|x\|$, then we have

$$|\langle Mx, x\rangle| \leq \|x^\| + B'x^\perp\| = \|x^\|\|^2 + \|B'x^\perp\|^2 \leq \left(1 - (1-\lambda^2)\left(\frac{1-\lambda}{3\mu}\right)^2\right)\|x\|^2$$

Since it holds for both cases, we have $f(\lambda, \mu) < 1$ for any $\lambda, \mu < 1$. □

Theorem 3.2.6 is enough for us to construct a family of degree D^2 expanders. Let us fix some graph H which is a $(D^4, D, 1/5)$-spectral expander. Then we may define $G_1 = H^2$ and $G_i = G_{i-1}^2 \textcircled{z} H$ for all $i \geq 2$. It is easy to check that all graphs $G_i = (V_i, E_i)$ in this family have eigenvalue at most $\frac{2}{5}$. Unfortunately, this construction is inefficient because computing neighbors in the graph takes time polynomial in the size $N_i = |V_i| = D^{4i}$. We would like a family where computing neighbors takes time polynomial in $\log N_i$ because in most applications, such as the derandomization application discussed in Section 2.4, though we only look at one vertex at a time, the size of the graphs required is actually exponential.

To achieve such a family, we use the graph analogue of the matrix tensor product.

Definition 3.2.7. Let $G = (V_1, E_1)$ and $H = (V_2, E_2)$ be degree D_1, D_2 graphs, respectively. Fix some rotation maps Rot_G and Rot_H of G and H. Then their *tensor product* $\Gamma = G \otimes H$ is the graph with vertex set $V_1 \times V_2$ and with rotation map

$$\mathrm{Rot}_\Gamma : (V_1, V_2) \times ([D_1] \times [D_2]) \to (V_1, V_2) \times ([D_1] \times [D_2])$$

is given by $\mathrm{Rot}_\Gamma((i,c),(k,\ell)) = ((j,d),(m,n))$ iff $\mathrm{Rot}_G(i,k) = (j,m)$ and $\mathrm{Rot}_H(c,\ell) = (d,n)$.

A vertex in Γ can be thought of as being simultaneously two independent vertices in G and H. Taking a step in Γ can be thought of as independently taking steps in G and H. It is easy to check that if A and B are the normalized adjacency matrices of G and H, then $A \otimes B$ is the normalized adjacency matrix $G \otimes H$.

Lemma 3.2.8. *The eigenvalues of $G \otimes H$ are exactly all $\lambda\mu$ where λ is an eigenvalue of G and μ is an eigenvalue of H.*

Proof. Consider the eigenvectors $\{v_i\}$ of G with eigenvalues $\{\lambda_i\}$ and the eigenvectors $\{w_i\}$ of H with eigenvalues $\{\mu_i\}$. From the definition of tensor product, we see that for any i, j:

$$(A \otimes B)(v_i \otimes w_j) = (Av_i) \otimes (Bw_j) = \lambda_i \mu_j (v_i \otimes w_j)$$

□

Remark 3.2.9. If G and H are connected, then $G \otimes H$ is disconnected if and only if both G and H are bipartite. If both are bipartite, then both have -1 as an eigenvalue so $G \otimes H$ has 1 with multiplicity two, which by Lemma 1.3.4 means $G \otimes H$ is disconnected. If at least one of G or H is non-bipartite, then it is impossible for $G \otimes H$ to have a second eigenvalue of 1, so it is connected.

Corollary 3.2.10.

$$\lambda_2(G \otimes H) = \max(\lambda_2(G), \lambda_2(H))$$

Proof. Since 1 is an eigenvalue of both G and H, it follows that both $\lambda_2(G)$ and $\lambda_2(H)$ are eigenvalues of $G \otimes H$, and so the greater of the two is the second largest eigenvalue of $G \otimes H$. □

Now we can construct an efficiently computable family of degree D^2 expanders.

Let H be a fixed (D^8, D, λ)-spectral expander for some small λ. Let $G_1 = H^2$ and $G_2 = H \otimes H$. Then for $i > 2$ define

$$G_i = \left(G_{\lceil \frac{i-1}{2} \rceil} \otimes G_{\lfloor \frac{i-1}{2} \rfloor}\right)^2 \textcircled{z} H$$

It is easy to check that G_i is a degree D^2 graph on $N_i = D^{8i}$ vertices. Computing the rotation map is efficient because the depth of the recursion is $\log_2 i$, and at each level there are 4 recursive calls, for a total of $4^{\log_2 i} = i^2$ for a graph on D^{8i} vertices. The eigenvalue of G_i can be analyzed by solving the recurrence relation given by applying Theorem 3.2.6 to the recursive definition of G_i. [25] shows that this gives us $\lambda_2(G_i) = \lambda + O(\lambda^2)$, which is bounded < 1 if λ is small enough.

However, note that this does not give an arbitrarily good family of expanders. Requiring λ to be extremely small forces the degree D of the graph H to grow, and indeed the best asymptotic behavior of the second largest eigenvalue that this construction allows is $O(1/\sqrt[4]{D})$. The derandomized zig-zag product included in [25] gives a better asymptotic bound of $O(1/\sqrt[3]{D})$, but it is currently unknown how to construct an infinite family of graphs with second largest eigenvalue $O(1/\sqrt{D})$ using the zig-zag product or a related graph product.

One may also ask where we obtain the base graphs H. Note that here we would like families that do not have constant degree; instead we would like an optimal degree/eigenvalue relationship but where the degree grows (and so the second largest eigenvalue shrinks) so that we can get sufficiently good spectral expanders to use as base graphs in the above recursive constructions. [25] gives a list of suggested base graphs, which are algebraic constructions based on finite fields.

Chapter 4

The Zig-Zag Product and the Algebra Connection

The zig-zag product is remarkable in being the first tool to construct infinite families of expander graphs relying on combinatorial rather than algebraic arguments. It is thus somewhat ironic that there is a natural relationship between the wide zig-zag product and one of the fundamental operations of abstract algebra, the semi-direct product from group theory. In this chapter, we will introduce and explain the term *wide zig-zag product* and present a novel and expanded exposition of the fascinating connection established by Alon, Lubotzky, and Widgerson [4] between the wide zig-zag product and the semi-direct product.

4.1 Group Theory Terminology

Let us first recall some definitions and simple facts from group theory. A group $(A, \times)$ is a set of elements A and a rule of composition $\times : A \times A \to A$ with the three following properties.

Identity $\exists 1 \in A$ such that $\forall a \in A$, $1 \times a = a \times 1 = a$.

Inverse $\forall a \in A$, $\exists b \in A$ such that $a \times b = b \times a = 1$. We call $b = a^{-1}$.

Associativity $\forall a, b, c \in A$, we have $(a \times b) \times c = a \times (b \times c)$.

We will usually omit the explicit symbol $\times$ and write the product of two elements implicitly ($a \times b = ab$). Likewise we will often refer to the group by its set of elements A without mentioning the rule of composition.

$S \subset A$ is said to *generate* A if we can write any $a \in A$ as a finite product of elements of S. Formally, $\forall a \in A$, $\exists s_1 \dots s_n \in S$ where $n < \infty$ such that $a = \prod s_i$. S is symmetric if $\forall s \in S$, $s^{-1} \in S$. We will also consider symmetric generating *multi*-sets S, where elements may be repeated (their inverses being repeated with the same multiplicity).

The *direct product* $A \times B$ of two groups A and B is the group with elements in the Cartesian product $A \times B$ and with the rule of composition $(a, b)(c, d) = (ac, bd)$.

The link between algebra and graphs is provided by the notion of the Cayley graph.

Definition 4.1.1. The *Cayley graph* of a group A over the symmetric generating multi-set S is the $|S|$-regular graph $C(A, S) = (V, E)$ where $V = A$ and $\{a, b\} \in E$ if $as = b$ for some $s \in S$ (with appropriate multiplicity).

Note that there is a natural rotation map associated with any Cayley graph, namely $\mathrm{Rot}_G(a, s) = (b, s^{-1})$ iff $as = b$. As a warm-up, let us consider the relationship between the direct product of groups and the tensor product of their Cayley graphs.

Proposition 4.1.2. *Let A, B be groups with respective generating multi-sets α, β. Then the Cayley graph of $A \times B$ over $\alpha \times \beta$ is isomorphic to the tensor product graph $C(A, \alpha) \otimes C(B, \beta)$.*

Proof. It is clear that $G = C(A \times B, \alpha \times \beta)$ and $H = C(A, \alpha) \otimes C(B, \beta)$ have the same vertex sets and are regular with the same degree. It remains to be shown that they have the same edges.

Any edge in G has the form $\{(a, b), (as, bt)\}$ for some $(a, b) \in A \times B$ and some $(s, t) \in \alpha \times \beta$. But this means that $\{a, as\}$ is an edge in $C(A, \alpha)$ and $\{b, bt\}$ is an edge in $C(B, \beta)$, so from the definition of the graph tensor product it must be that $\{(a, b), (as, bt)\}$ is an edge in H. The reverse direction is similar. □

4.2 The Semi-Direct Product

We now strengthen this relationship to include a richer set of group products. We need a little more vocabulary before proceeding. A group *homomorphism* between A and B is a map $\phi : A \to B$ that preserves the rule of composition:

$$\forall a, a' \in A,\ \phi(aa') = \phi(a)\phi(a')$$

A homomorphism is an *isomorphism* if it is one-to-one and onto, and an isomorphism is an *automorphism* if it maps to and from the same group. We will use $\mathrm{Aut}(G)$ to denote the set of automorphisms of a group G; it is easy to check that $\mathrm{Aut}(G)$ is a group where the rule of composition is the composition of maps.

Let A and B be groups. We say that B *acts on* A if there is a homomorphism $\Phi : B \to \mathrm{Aut}(A)$. The action of $b \in B$ on an element $a \in A$ is denoted $b \cdot a = (\Phi(b))(a)$. We will always explicitly write the operator $\cdot$ to denote the action of B on A. The *orbit* of $a \in A$ under some action of B is the set

$$B \cdot a = \{a' \in A \mid \exists b \in B \ s.t.\ b \cdot a = a'\}$$

Definition 4.2.1. The *semi-direct product* $A \ltimes B$ of two groups A and B has elements in $A \times B$ with the rule of composition

$$(a, b)(c, d) = (a(b \cdot c), bd)$$

using the action of B on A defined by any homomorphism Φ as above.

Note that this means the semi-direct product is ambiguous unless we explicitly state which action we are using to define $b \cdot c$. Also note that in the trivial case where $\forall a, b \in B,\ b \cdot a = a$ (i.e. Φ in the above definition is trivial), then the semi-direct product degenerates to the direct product. The semi-direct product is interesting because it offers a much richer vocabulary with which to express the structure of many common groups. Conceptually, the homomorphism defines the behavior of $a \in A$ under conjugation by an element $b \in B$.

Example 4.2.2. As an example, consider the dihedral group D_n, which is generated by two elements ρ and τ with the relations

$$\rho^n = \tau^2 = 1 \qquad\qquad \tau^{-1}\rho\tau = \rho^{-1}$$

Recall that the dihedral group D_n may be thought of as the symmetries of the regular n-gon, consisting of all the rotations and reflections preserving the n-gon.

We may represent D_n as the semi-direct product $C_n \ltimes C_2$ where C_n and C_2 are the cyclic groups of orders n and 2. Let ρ be a generator of C_n and τ be a generator of C_2. Then the homomorphism $\Phi : C_2 \to \mathrm{Aut}(C_n)$ of this semi-direct product is $\Phi(1)(x) = x\ \forall x \in C_n$ and $\Phi(\tau)(x) = x^{-1}\forall x \in C_n$. To verify that this is D_n, we check

$$(\rho, 1)^n = (1, \tau)^2 = 1 \qquad\qquad (1, \tau^{-1})(\rho, 1)(1, \tau) = (\rho^{-1}, \tau^{-1})(1, \tau) = (\rho^{-1}, 1)$$

since $\tau^{-1} = \tau$. Notice here that, for any $a \in C_2$, the homomorphism $\varphi = \Phi(a)$, $\varphi : C_n \to C_n$ gives exactly the result of conjugating an element of C_n by a.

We will discuss generating multi-sets α of groups A, where there is an action of some other group B on A. We will be particularly interested when α is composed of relatively few B-orbits. If this is the case, we can show that using $\alpha \subset A$ and $\beta \subset B$ generating multi-sets of their respective groups, we can define a (small) generating multi-set $\gamma \subset A \ltimes B$ such that $C(A \ltimes B, \gamma)$ is equal to the wide zig-zag product of the Cayley graphs of A and B.

4.3 The Wide Zig-Zag Product

Here we introduce the wide zig-zag product, which is a natural generalization of the zig-zag product of the previous chapter.

Definition 4.3.1. Let $G = (V_1, E_1)$ be a graph on N vertices with degree nD_1, and let $H = (V_2, E_2)$ be a graph on D_1 vertices with degree D_2. Fix an edge-labelling of G with repetitions, such that for all vertices i, there are exactly n incident edges labelled with each of $1, \ldots, D_1$.

Then the *wide zig-zag product* $G ⓦ H$ is defined to be the graph on ND_1 vertices with degree nD_2^2 where an edge $\{(i, c), (j, d)\}$ is in the graph with the same multiplicity as there are pairs x, x' such that:

1. $\{c, x\} \in E_2$
2. $\exists$ an edge $e = \{i, j\} \in E_1$ incident to i that is labelled x at vertex i and is labelled x' at vertex j
3. $\{x', d\} \in E_2$

A rotation map for $G ⓦ H$ may be defined analogous to the rotation map for the regular zig-zag product.

This definition is similar to the definition of the standard zig-zag product, except the intermediate step on G is "wider": instead of a unique intermediate step defined according to the edge-labelling, there are n such possible intermediate steps.

Corollary 4.3.2. *If G is a (N, nD_1, λ)-spectral expander and H is a (D_1, D_2, μ)-spectral expander, then $G \textcircled{w} H$ is a $(ND_1, nD_2^2, f(\lambda, \mu))$-spectral expander for the same f as in Theorem 3.2.6.*

Proof. This follows almost directly from the proof of Theorem 3.2.6. Recalling the terminology used there, the normalized adjacency matrix M of $G \textcircled{w} H$ is also a product of the form $B'A'B'$. Because the intermediate step is now non-deterministic, A' is no longer a permutation matrix. However, it is still doubly stochastic, so the property $\|A'x^{\|}\| \leq \|x^{\|}\|$, holds.

We also desire the property $\langle A'x^{\|}, x^{\|}\rangle = \langle AC(x), C(x)\rangle / D_1$. This still holds because, as in the original proof, $C(A'(e_i \otimes 1_{D_1}/D_1)) = Ae_i$.

Since these are the only two properties of A' that we relied on, and B' is the same in both cases, the proof directly carries over to the wide zig-zag product case. □

We now show that by appropriately defining the generating set and the edge labelling respectively, the Cayley graph of the semi-direct product of two groups and the wide zig-zag product of the Cayley graphs of the two groups are one and the same.

Theorem 4.3.3. *Let $\alpha \subset A$ and $\beta \subset B$ be generating multi-sets of their respective groups. Fix some action of B on A. Suppose that there exists a set of $a_i \in \alpha$ such that $\alpha = \bigcup_{i=1}^{n} \alpha_i$, where α_i is the symmetric multi-set*

$$\alpha_i = (B \cdot a_i) \cup (B \cdot a_i^{-1})$$

That is, α is a union of B-orbits of n elements a_i and their inverses. Let γ be the following generating multi-set of $A \ltimes B$:

$$\gamma = \{(1, b)(a_i^s, 1)(1, b') \mid b, b' \in \beta, s \in \{\pm 1\}, i \in [n]\}$$

of size $2n|\beta|^2$. Let $\Gamma_1 = C(A \ltimes B, \gamma)$.

Now define $G = C(A, \alpha)$ and $H = C(B, \beta)$. Then there exists a labelling of G such that $\Gamma_2 = G \textcircled{w} H = \Gamma_1$.

Proof. First we specify the labelling of G. For any vertex x, we label an outgoing edge $e \in \alpha$, where $e = (b \cdot a_i)$ or $e = (b \cdot a_i^{-1})$ for some $i \in [n]$, to be $b \in B$. Call the neighbor $x' = xe$. We require that e is also labelled b at x', which is consistent with the above since the inverse of $(b \cdot a_i)$ is $(b \cdot a_i^{-1})$ and vice versa. This labelling is easily constructible when constructing the graph or when computing neighbors. One can check that this satisfies the constraint on the labelling of G specified in Definition 4.3.1: G has degree $2n|B|$, and there are exactly $2n$ vertices labelled by each element of $|B|$ at each vertex.

Showing that $\Gamma_1 = \Gamma_2$ follows straightforwardly from the definitions of the wide zig-zag product and the semi-direct product. It is clear that Γ_1 and Γ_2 share the same vertex set, and they both have degree $2n|\beta|^2$. So we need only show that the edges of the two graphs are identical.

An edge in Γ_1 leaving the vertex (x, y) given by $(1, b)(a_i^s, 1)(1, b') \in \gamma$ goes to the neighbor

$$(x, y)(1, b)(a_i^s, 1)(1, b') = (x(yb \cdot a_i^s), ybb')$$

We claim that this edge is also in Γ_2. Starting from (x, y), taking a "zig" step of b takes us to (x, yb); in the labelling of G, yb is a label of the edge going to $x(yb \cdot a_i^s)$, so we follow this edge to go from cloud x to cloud $x(yb \cdot a_i^s)$, and in this new cloud the edge is labelled yb as well by our choice of labelling. So we arrive at $(x(yb \cdot a_i^s), yb)$, and taking a final "zag" step of b' takes us to $(x(yb \cdot a_i^s), ybb')$.

The reverse direction is shown similarly. □

Example 4.3.4. As an example, consider the case where the action of B on A is trivial. Then it must be the case that any symmetric generating multi-set $\alpha \subset A$ is the union of trivial orbits, i.e. $\alpha_i = \{a_i\} \cup \{a_i^{-1}\}$. Using this action, the semi-direct product $A \ltimes B = A \times B$ the direct product. However, because $\gamma \neq \alpha \times \beta$, this does not follow the case of Proposition 4.1.2. Instead, we see that

$$\gamma = \{(a_i^s, bb') \mid b, b' \in \beta, s \in \{\pm 1\}, i \in [n]\}$$

So $G ⓦ H = G \otimes H^2 = C(A \times B, \gamma)$.

As this example shows, when n is large this relationship is uninteresting because the degree of the resulting product graph is large. We are most interested in groups generated by a small number of orbits, that is, where n is small. This means the degree of the wide zig-zag product graph is small, and so we get a family of low-degree expanders.

In [4] this relationship is used to show that expansion is not a group property, so the expansion properties of the Cayley graphs of a particular group depend on which generator multi-set one picks. They show that it is is possible to create two infinite families of graphs $\{G_i\}$ and $\{H_i\}$ where each G_i and H_i are Cayley graphs of the same group but whose edges are defined using different generator multi-sets. Remarkably, the expansion of G_i stays good as $i \to \infty$ but the expansion of H_i deteriorates as $i \to \infty$.

Chapter 5

Eigenvalue Lower Bounds for the Zig-Zag Product

As the wide zig-zag product/semi-direct product relationship given in Theorem 4.3.3 shows, the choice of labelling is extremely important in determining the product graph's structure. Specifically, given G and H such that $G \textcircled{z} H$ exists (i.e. G and H have the appropriate size and degree relationships), using different rotation maps for G may result in different (non-isomorphic) product graphs.[1]Theorem 3.2.6 gives us an upper bound on the expansion of the product graph independent of the rotation maps, which is sufficient to guarantee that $G \textcircled{z} H$ is a good expander when both G and H are good expanders. However, it seems the rotation maps come more explicitly into the analysis for the problem we discuss in this chapter.

5.1 The Question

We pose the question of whether it is *necessary* for both the component graphs G and H to be good expanders in order for $G \textcircled{z} H$ to be a good expander. For the remainder of this chapter, unless otherwise noted we will let G be a degree D_1 graph on N vertices and let H be a degree D_2 graph on D_1 vertices. It is easy to establish a lower bound for how well $G \textcircled{z} H$ can expand based on the expansion of G.

Proposition 5.1.1. *For any graphs G, H such that $G \textcircled{z} H$ exists, it is true regardless of the choice of* $\mathrm{Rot}_G, \mathrm{Rot}_H$ *that* $\lambda_2(G \textcircled{z} H) \geq \lambda_2(G)$.

Proof. Let $M = B'A'B'$ be the normalized adjacency matrix of $G\textcircled{z}H$ (see the proof of Theorem 3.2.6 for the terminology), and let v be the eigenvector of G corresponding to $\lambda_2(G)$. Taking $v' = v \otimes 1_{D_1}/D_1$, we get

$$\begin{aligned}
\langle Mv', v' \rangle &= \langle A'B'v', B'v' \rangle && (5.1)\\
&= \langle A'(v \otimes 1_{D_1}), v \otimes 1_{D_1} \rangle / D_1^2 && (5.2)\\
&= \langle Av, v \rangle / D_1 && (5.3)\\
&= \lambda_2(G)\langle v', v' \rangle && (5.4)
\end{aligned}$$

[1] It is easy to check that if we fix Rot_G, then the different zig-zag products of G and H according to different rotation maps for H are still isomorphic. However, their rotation maps will label the edges differently.

We implicitly use the collection map C from the analysis of Theorem 3.2.6 to go from (5.2), which is a dot product of vectors in $\mathbb{R}^{ND_1}$, to (5.3), which is a dot product of vectors in $\mathbb{R}^N$. So $\lambda_2(G \textcircled{z} H) \geq \lambda_2(G)$. $\square$

This property makes sense because if G is a bad expander then it is "hard" to get from one vertex to another. Recalling the cloud intuition from Subsection 3.2.1, this means that it is hard to get from one cloud to another in $G \textcircled{z} H$, and so $G \textcircled{z} H$ cannot be a good expander.

Unfortunately, there is no clear lower bound relationship between $\lambda_2(G \textcircled{z} H)$ and $\lambda_2(H)$. In particular, this task is more difficult because it seems that any analysis of a lower bound must explicitly take into account the rotation map of G.

As a basis for comparison, we can establish some bounds based not on the expansion of H but on its degree. Applying the bound on infinite families of graphs from Theorem 2.3.2 says that, as $G \textcircled{z} H$ gets large, we have

$$\lambda_2(G \textcircled{z} H) \geq \frac{2\sqrt{D_2^2 - 1}}{D_2^2}$$

We can get a slightly stronger relationship between $\lambda_2(G \textcircled{z} H)$ and D_2 when G has diameter at least 3.

Proposition 5.1.2. *Take G, H graphs such that $G \textcircled{z} H$ exists. Suppose that G has vertices i, j such that the distance between them is at least 3. Then $\lambda_2(G \textcircled{z} H) \geq \frac{1}{\sqrt{D_2}}$.*

Proof. Since i, j have distance at least 3, they have no neighbors in common. Let us define the distribution $x = e_i \otimes 1_{D_1} + e_j \otimes (-1_{D_1})$ where the e_i are the standard basis vectors for $\mathbb{R}^N$. Clearly $x \perp u$, and $\|x\| = \sqrt{2D_1}$.

Recalling the cloud terminology from Subsection 3.2.1, let us see what happens when we apply $B'A'B'$ to x. The first step B' fixes x since x is uniform on each cloud. The second step A' sends the weight of each vertex c in the i to some neighboring cloud, and likewise for j. The last step B' takes a step within clouds on the resulting distribution.

Since we are concerned with lower-bounding $\lambda_2(G \textcircled{z} H)$, and since there is no "negative interference" between the positive and negative weights because i, j have no common neighbors, the worst case is when, in the middle step, the weight of each vertex in the i, j'th clouds is sent to different neighboring clouds. Formally, $\nexists k \in [N]$ such that $\exists c, d, e, f \in [D_2]$ where $\mathrm{Rot}_G(i, c) = (k, e)$ and $\mathrm{Rot}_G(i, d) = (k, f)$, and similarly for the j'th cloud. In this case, after the final step we have a D_2 number of weights of value $\frac{1}{D_2}$ in each of $2D_1$ clouds, and so $\|B'A'B'x\| = \sqrt{\frac{2D_1}{D_2}}$.

Therefore, by Lemma 2.1.2 we have $\lambda_2(G \textcircled{z} H) \geq \frac{\|B'A'B'x\|}{\|x\|} = \frac{1}{\sqrt{D_2}}$. $\square$

Let us return to the problem of relating the expansion of $G \textcircled{z} H$ to the expansion of H. Our intuition tells us that the expansion of $G \textcircled{z} H$ cannot be much better than that of H^2. It seems from the definition of the zig-zag product that $G \textcircled{z} H$ is related to H^2 because the "zig" and "zag" steps are actually two independent steps in H. Thus, if H is a bad expander then $G \textcircled{z} H$ cannot be very good because the two steps in H do not do much. However, the middle permutation induced by G obscures how one might use this intuition in a proof of a lower bound relationship. Nevertheless, we will present evidence that supports the following conjecture.

Conjecture 5.1.3. *For any graphs G, H such that $G \textcircled{z} H$ exists, it is true regardless of the choice of* $\mathrm{Rot}_G, \mathrm{Rot}_H$ *that* $\lambda_2(G \textcircled{z} H) = \Omega(\lambda_2(H)^2)$.

In Section 5.3 we examine some special cases where we can prove the conjecture in the strong sense $\lambda_2(G \textcircled{z} H) \geq \lambda_2(H)^2$, and in Section 5.4 we give experimental evidence that this bound is universal. Before proceeding, we need some vocabulary from combinatorics.

5.2 Combinatorics Terminology

We will employ the following combinatorics terminology in our discussion of rotation maps and in later sections. The terminology here is entirely standard.

- A *permutation* σ of a finite set is a bijection from that set to itself. When working with permutations of N elements, we will usually let $[N]$ be the set on which the permutations act.
- The set of all permutations of N elements forms the *symmetric group* S_N with the rule of composition being the composition of maps.
- A *cycle* $(i_1, i_2, \ldots i_k)$ denotes the permutation taking $i_j \mapsto i_{j+1}$ for all $1 \leq j < k$ and taking $i_k \mapsto i_1$. The cycle fixes all elements in $[N]$ not appearing in the cycle. For us, the product of two cycles st will denote the permutation given by first applying t, then applying s.
- Any permutation $\sigma \in S_N$ may be written in *cycle notation* as the product of disjoint cycles.
- A *transposition* is a permutation (i_1, i_2) for $i_1 \neq i_2$.
- A permutation σ is an *involution* if it is a product of disjoint transpositions. Equivalently, σ is an involution if $\sigma^2 = 1$.
- The *matrix of a permutation* σ acting on N elements is the matrix $P \in M_N(\{0, 1\})$ with entries $\delta(i, \sigma(i))$ (δ being the Kronecker delta).
- Let $G = (V, E)$ be a graph. A permutation σ of V is said to be an *automorphism* of G if for all $i, j \in V$, we have that $\{\sigma(i), \sigma(j)\} \in E$ with the exact same multiplicity as there are $\{i, j\} \in E$.
- The *automorphism group* of a graph G is the set of all automorphisms of G with the rule of composition being the composition of maps.

Remark 5.2.1. If P is the matrix of a permutation σ on N elements, then P's action by left multiplication on a matrix with N rows simply interchanges the rows of the matrix according to σ. P's left action on vectors follows similarly.

Remark 5.2.2. If we let P represent the matrix of the permutation σ on N elements and let A be the normalized adjacency matrix of a graph G on N vertices, then it is easy to check using the above definition that σ is an automorphism of G if and only if $AP = PA$.

5.3 Provable Lower Bounds in Special Cases

The simplest case where we can show a lower bound of $\lambda_2(G \textcircled{z} H)$ in relation to $\lambda_2(H)$ is when the rotation map of G is an edge-coloring of G.

Proposition 5.3.1. *Suppose there exists a rotation map* Rot_G *that is an edge-coloring using* D_1 *colors. That is, for any* $i \in [N], c \in [D_1]$, *we have* $\mathrm{Rot}(i, c) = (j, c)$ *for some* $j \in [N]$. *Then the zig-zag product* $G \textcircled{z} H$ *according to* Rot_G *has* $\lambda_2(G \textcircled{z} H) \geq \lambda_2(H)^2$.

Proof. Let v be an eigenvector of H corresponding to $\lambda_2(H)$. We use the notation of the proof of Theorem 3.2.6 and let $B'A'B'$ denote the normalized adjacency matrix of $G \textcircled{z} H$. Then setting $v' = 1_N \otimes v/N$, we can use Lemma 2.1.2 to see that

$$\max_{x \perp u} \frac{|\langle B'A'B'x, x\rangle|}{\langle x, x\rangle} \geq \frac{|\langle B'A'B'v', v'\rangle|}{\langle v', v'\rangle} = \frac{\lambda_2(H)^2 |\langle A'v', v'\rangle|}{\langle v', v'\rangle}$$

Recalling the definition of A' from the proof of Theorem 3.2.6, we see that it permutes v' by sending the $v'_{(i,c)}$ to $(A'v')_{\mathrm{Rot}_G(i,c)}$. Since $\mathrm{Rot}_G(i, c) = (j, c)$ for any $i \in [N]$, we have that v' is identical across clouds so that $v'_{(i,c)} = v'_{(j,c)}$ for all $i, j \in [N]$. It follows that by transposing the (i, c)'th and (j, c)'th entries, A' actually just fixes v'. Therefore $\frac{\lambda_2(H)^2|\langle A'v', v'\rangle|}{\langle v', v'\rangle} = \lambda_2(H)^2$. □

Unfortunately, as we noted in Section 3.2.1 an edge-coloring with D_1 color does not always exist, and it is hard to compute even if it does, so the above example does not have great generality.

An eminently more interesting example is when the graph has a rotation map with the following property of being a semi-coloring. We define it using the coloring interpretation of rotation maps discussed in Remark 3.2.2.

Definition 5.3.2. A rotation map Rot_G of a graph G is a *semi-coloring* if the edge colors come in fixed pairs: for any color c, there is a "partner color" d such that any edge that is colored c at one end is always colored d at the other end. (We allow the possibility that $c = d$.) Formally, this means $\forall c \in [D_1], \exists! d \in [D_1]$ such that $\forall i \in [N]$, $\exists! j \in [N]$ such that we have $\mathrm{Rot}_G(i, c) = (j, d)$.

The reason semi-colorings are more interesting than edge-colorings with D_1 colors is because the zig-zag product naturally preserves semi-colorings. Let G be the large graph and H be the small graph in the zig-zag product $\Gamma = G \textcircled{z} H$. It is true that, regardless of Rot_G, as long as Rot_H is a semi-coloring then the rotation map Rot_Γ is also a semi-coloring. This is because in the mapping $((i, c), (k, \ell)) \mapsto ((j, d), (m, n))$, G only influences the mapping $i \mapsto j$ and so the mapping $(k, \ell) \mapsto (m, n)$ depends solely on H. Thus if k, n and ℓ, m are fixed color pairs in Rot_H, then so is $(k, \ell), (m, n)$ in Rot_Γ.

It is worth noting that the zig-zag product does not preserve edge colorings with D_1 colors. If Rot_H has a coloring, then in Rot_Γ an edge colored (k, ℓ) on one end will be colored (ℓ, k) on the other.

It is easy to check that in all of the recursive expander constructions given in Subsection 3.2.2, all the graphs in the families have a semi-coloring as long as the base graph H has one. Furthermore, all of the suggested base graphs given in [25] (as well as all Cayley graphs, which include the Ramanujan graphs of [20, 21]) come with natural semi-colorings. Thus, it seems in practice a proof of a lower bound for this special case would be extremely interesting.

It turns out that finding a lower bound for $\lambda_2(G \textcircled{z} H)$ where Rot_G is a semi-coloring reduces to finding a lower bound on the expansion of the semi-square of H, which we define presently.

Definition 5.3.3. Let $G = (V, E)$ be a graph on N vertices that is regular with degree D, and let σ be any involution on N elements. Then the *semi-square* of G with respect to σ is the degree D^2 graph $G^{\lfloor 2 \rceil}$ on V where an edge $\{i, j\}$ is in $G^{\lfloor 2 \rceil}$ as many times as there exist $k \in V$ such that $\{i, k\} \in E$ and $\{\sigma(k), j\} \in E$. It follows that if the normalized adjacency matrix G is A and P is the matrix representation of σ, then the normalized adjacency matrix of $G^{\lfloor 2 \rceil}$ is APA.

Of course, the semi-square reduces to the regular square if σ is the identity involution. We can reduce the semi-coloring lower bound into a semi-squaring lower bound because the semi-coloring induces an involution on the vertices of H.

Proposition 5.3.4. *Take the zig-zag product $G \textcircled{z} H$ according to a rotation map Rot_G that is a semi-coloring of G. Rot_G defines an involution σ on $[D_1]$, and $\lambda_2(G \textcircled{z} H) \geq \lambda_2(H^{\lfloor 2 \rceil})$ where the semi-square $H^{\lfloor 2 \rceil}$ is taken with respect to σ.*

Proof. Let $B'A'B'$ be the normalized adjacency matrix of $G \textcircled{z} H$. For any vector $y \in \mathbb{R}^{D_1}$, let $y' = 1_N \otimes y/N$, and use Lemma 2.1.2 to get

$$\max_{x \perp u} \frac{|\langle B'A'B'x, x\rangle|}{\langle x, x\rangle} \geq \max_{y \perp u} \frac{|\langle B'A'B'y', y'\rangle|}{\langle y', y'\rangle}$$

Recalling the definition of A' from the proof of Theorem 3.2.6, we see that it permutes $B'y'$ by sending $(B'y')_{(i,c)}$ to $(A'B'y')_{(j,d)}$ where $\mathrm{Rot}_G(i, c) = (j, d)$. It must be that c, d are paired in the semi-coloring. For each such pairing, let us define $\sigma(c) = d$ and $\sigma(d) = c$. σ fixes all elements that are paired to themselves.

Recall that, by our definition of $B' = I_N \otimes B$ where B is the normalized adjacency matrix of H, we have that $B'y' = 1_N \otimes By/N$. Now consider the entries of $A'B'y'$. For any $c \in [D_1]$, the weight of the c'th vertex on any cloud of $A'B'y'$ depends only on the weight of the $\sigma(c)$'th vertex of some cloud in $B'y'$. But since $B'y'$ is identical across clouds, this means $A'B'y'$ is identical across all clouds. In particular, because A' permutes the vertices within a cloud according to σ, we have $A'B'y' = 1_N \otimes PBy/N$ where P is the $D_1 \times D_1$ permutation matrix of σ. Therefore we have

$$\max_{y \perp u} \frac{|\langle B'A'B'y', y'\rangle|}{\langle y', y'\rangle} = \max_{y \perp u} \frac{|\langle 1_N \otimes (BPBy)/N, 1_N \otimes y/N\rangle|}{\langle y', y'\rangle} = \max_{y \perp u} \frac{|\langle BPBy, y\rangle|}{\langle y, y\rangle}$$

The last step follows from a simple calculation expanding the dot product. Since BPB is exactly the normalized adjacency matrix of $H^{\lfloor 2 \rceil}$ according to σ, Lemma 2.1.2 gives us $\lambda_2(G \textcircled{z} H) \geq \lambda_2(H^{\lfloor 2 \rceil})$. □

So finding a lower bound on $\lambda_2(G \textcircled{z} H)$ reduces to finding a lower bound on $\lambda_2(H^{\lfloor 2 \rceil})$. This reduction is useful because as the following lemma shows, we can calculate exactly the expansion of the semi-square in a special case. Even more general results about the semi-square are established in Chapter 6.

Lemma 5.3.5. *Let G be a graph with A its normalized adjacency matrix. Take the semi-square $G^{\lfloor 2 \rceil}$ with respect to an involution σ. Then $\lambda_2(G^{\lfloor 2 \rceil}) \leq \lambda_2(G)^2$. Furthermore, if σ gives an automorphism of G, then $\lambda_2(G^{\lfloor 2 \rceil}) = \lambda_2(G)^2$.*

Proof. Let P be the matrix of the involution σ. To see that $\lambda_2(G^{\lfloor 2 \rceil}) \leq \lambda_2(G)^2$, we use Lemma 2.1.2, the fact that P is length-preserving, and the fact that P and A preserve anti-uniformity.

Since APA is the normalized adjacency matrix of $G^{\lfloor 2 \rceil}$, we have for all $x \perp u$ that

$$\frac{\|APA\|}{\|x\|} \le \frac{\lambda_2(G)\|PAx\|}{\|x\|} \le \frac{\lambda_2(G)\|Ax\|}{\|x\|} \le \lambda_2(G)^2$$

The second part of the lemma follows from the fact that, if P gives an automorphism of G, then $PA = AP$. Therefore, Lemma 2.1.2 reduces to:

$$\max_{x \perp u} \frac{\|APAx\|}{\|x\|} = \max_{x \perp u} \frac{\|PA^2x\|}{\|x\|} = \max_{x \perp u} \frac{\|A^2x\|}{\|x\|} = \lambda_2(G)^2$$

□

We quickly state a few simple consequences of the above.

Corollary 5.3.6. *If G is connected and non-bipartite, then any semi-square $G^{\lfloor 2 \rceil}$ is connected and non-bipartite.*

Proof. This follows immediately by combining Lemmas 1.3.4, 1.3.6, and 5.3.5. □

Corollary 5.3.7. *Let G, H, σ be as in Proposition 5.3.4. If σ is an automorphism of H, then $\lambda_2(G \textcircled{z} H) \ge \lambda_2(H)^2$.*

Proof. This follows immediately from Proposition 5.3.4 and Lemma 5.3.5. □

Unfortunately, it is hard to pin down exactly when the expansion of the semi-square of a graph hits the upper bound. This is especially true if we are considering all semi-squares of H with no restriction on the involution σ. In order to prove $\lambda_2(G \textcircled{z} H) \ge \lambda_2(H)^2$ for all semi-colorings of G, we would like our H to be such that all its possible semi-squares are no better expanders than H^2. We call this property semi-square invariance.

Definition 5.3.8. A graph G is called *semi-square invariant* if for any involution σ, the semi-square $G^{\lfloor 2 \rceil}$ with respect to σ has $\lambda_2(G^{\lfloor 2 \rceil}) = \lambda_2(G)^2$.

Corollary 5.3.9. *Let G, H be as in Proposition 5.3.4 and let Rot_G be a semi-coloring of G. If H is semi-square invariant then $\lambda_2(G \textcircled{z} H) \ge \lambda_2(H)^2$.*

Proof. This follows immediately from the Proposition 5.3.4 and the definition of semi-square invariance. □

We will explore the property of semi-square invariance in greater detail in Chapter 6. For example, in Proposition 6.1.4 we specify a class of semi-square invariant graphs. Thus if H is one of these graphs and if Rot_G is a semi-coloring, then we get that $\lambda_2(G \textcircled{z} H) \ge \lambda_2(H)^2$.

Proving Conjecture 5.1.3 for all zig-zag products using semi-colorings may be done by showing for any graph H that $\lambda_2(H^{\lfloor 2 \rfloor}) = \Omega(\lambda_2(H)^2)$ for all possible semi-squares $H^{\lfloor 2 \rfloor}$. At this time, it is unclear how one would go about this in full generality.

5.4 Experimental Lower Bounds

Here we present experimental evidence that the bound $\lambda_2(G \textcircled{z} H) = \Omega(\lambda_2(H)^2)$ is universal. We implemented the zig-zag product in Mathematica and computed the eigenvalues of the zig-zag products of certain graphs. The experiments in this section were run using (pseudo-) random rotation maps for G, which intuitively seem will give $G \textcircled{z} H$ the best expansion and hence the lowest $\lambda_2(G \textcircled{z} H)$.

We begin by defining the random regular graphs that we work with.

Definition 5.4.1. A *random regular graph* $R_{N,D}$ on N vertices with degree D is constructed as follows. Choose $\lfloor D/2 \rfloor$ random permutations from the symmetric group S_N and let $P_1, \ldots P_{\lfloor D/2 \rfloor}$ be their permutation matrices. If D is even, then the normalized adjacency matrix of $R_{N,D}$ is

$$\frac{1}{D} \sum_{i=1}^{\lfloor D/2 \rfloor} (P_i + P_i^t)$$

which guarantees it is doubly stochastic and symmetric. If D is odd, the normalized adjacency matrix of $R_{N,D}$ is

$$\frac{1}{D} \left(\sum_{i=1}^{\lfloor D/2 \rfloor} (P_i + P_i^t) + I \right)$$

i.e. we add one self-loop to each vertex.

Random regular graphs are used when we wish to find a large graph with good expansion, which happens with high probability according to Theorem 2.3.1. Our hardware allowed us to work with graphs on vertices of at most about 4000 vertices, depending slightly on the number of edges.

Our first series of experiments zig-zags together random graphs and cyclic graphs.

Definition 5.4.2. A *cyclic graph* C_N is a degree 2 graph on N vertices. Let us use $\mathbb{Z}_N$ to denote the vertex set, then each vertex $i \in \mathbb{Z}_N$ is connected to $i+1, i-1 \in \mathbb{Z}_N$.

Proposition 5.4.3. *Each $p \in \mathbb{Z}_N$ corresponds to an eigenvalue of $\cos \frac{2\pi p}{N}$ in the spectrum of C_N.*

Proof. Note that for the duration of this proof we work in $\mathbb{Z}_N$. It is easy to check that the matrix of C_N is $A = [a_{ij}]$, with $i, j \in \mathbb{Z}_N$, given by

$$a_{ij} = \begin{cases} \frac{1}{2},\ i - j = \pm 1 \\ 0, \text{ else} \end{cases}$$

We will show that it is similar to the diagonal matrix where the p'th diagonal is $\frac{2\pi p}{N}$. We use the one-dimensional discrete Fourier transform to do this, whose matrix $\mathfrak{F} = [\mathfrak{f}_{ij}]$ is given by $\mathfrak{f}_{ij} = \frac{1}{\sqrt{N}} \omega^{ij}$ for $i, j \in \mathbb{Z}_N$, where $\omega = \exp(2\pi\sqrt{-1}/N)$. (Contrast this with the *two-dimensional*

Fourier transform used in the proof of Theorem 3.1.1.) This gives us

$$\begin{aligned}
[\mathfrak{F}A\mathfrak{F}^*]_{ij} &= \sum_{k,\ell} \mathfrak{f}_{ik} a_{k\ell} \bar{\mathfrak{f}}_{j,\ell} \\
&= \frac{1}{2} \sum_{k} \mathfrak{f}_{ik} (\bar{\mathfrak{f}}_{j,k+1} + \bar{\mathfrak{f}}_{j,k-1}) \\
&= \frac{1}{2N} \sum_{k} \omega^{ik} (\omega^{-jk-j} + \omega^{-jk+j}) \\
&= \frac{(\omega^{-j} + \omega^{j})}{2N} \cdot \sum_{k} \omega^{(i-j)k}
\end{aligned}$$

By the orthogonality of the Fourier basis functions, we know that $\sum_k \omega^{(i-j)k} = 0$ if $i \neq j$ and $\sum_k \omega^{(i-j)k} = N$ if $i = j$. This means $a_{ij} = 0$ for $i \neq j$ and $a_{pp} = \frac{\omega^p + \omega^{-p}}{2} = \cos \frac{2\pi p}{N}$. □

We only work with odd-size cyclic graphs since even-size ones are bipartite. We choose cyclic graphs as our first test case because they intuitively seem like they might cause $\lambda_2(R_{N,D_1} \textcircled{z} C_{D_1}) \ll \lambda_2(C_{D_1})^2$. This is because though C_{D_1} is a bad expander, the middle step of the zig-zag product sends vertices in C_{D_1} to non-adjacent vertices, and since the main "cause" of C_{D_1}'s bad expansion is because some vertices are far away from each other, we might hope that the resulting product graph $R_{N,D_1} \textcircled{z} C_{D_1}$ has significantly better expansion than $C_{D_1}^2$.

We refer to Figures 5.1 and 5.2, which show the results of our series of experiments using random regular graphs and cyclic graphs. Here we define $\lambda = \lambda_2(R_{N,D_1}), \mu = \lambda_2(C_{D_1}), \eta = \lambda_2(R_{N,D_1} \textcircled{z} C_{D_1})$.

Samples 1-5 were done with random regular graphs $R_{50,11}, \ldots, R_{50,51}$ and with cyclic graphs $C_{11}, \ldots, C_{51}$ in increments of 10, and samples 6-10 were done with $R_{80,11}, \ldots, R_{80,41}$ and with cyclic graphs $C_{11}, \ldots, C_{41}$.

We see that although $\lambda_2(R_{N,D_1} \textcircled{z} C_{D_1}) < \lambda_2(C_{D_1})^2$ consistently, the two do not deviate far from each other. In particular, the behavior of $\lambda_2(R_{N,D_1} \textcircled{z} C_{D_1})$ and $\lambda_2(C_{D_1})^2$ are highly correlated with a correlation coefficient of 0.9433, whereas it is clear that $\lambda_2(R_{N,D_1} \textcircled{z} C_{D_1})$ and $\lambda_2(R_{N,D_1})$ are anti-correlated. This means, most likely, that the behavior of $\lambda_2(R_{N,D_1} \textcircled{z} C_{D_1})$ is tied to the behavior of $\lambda_2(C_{D_1})^2$. Furthermore, these numbers are well above the degree-based lower bounds of Theorem 2.3.2 and Proposition 5.1.2, which give us $\frac{\sqrt{3}}{2} \approx 0.866025$ and $\frac{1}{\sqrt{2}} \approx 0.707107$ respectively.

However, one reason why the above evidence may be misleading is because the eigenvalues of C_{D_1} are very evenly distributed, as according to Proposition 5.4.3.

Intuitively, if there are many eigenspaces with relatively large eigenvalue, it may be that the middle step of the zig-zag product is ineffectual. Consider v an eigenvector of C_{D_1} with eigenvalue $\lambda_2(C_{D_1})$. Then if there are many large eigenvalues, it may be that the middle step of the zig-zag product induced by R_{N,D_1} takes v to some v' in the eigenspace of $\lambda_k(C_{D_1})$, where the difference $|\lambda_2(C_{D_1}) - \lambda_k(C_{D_1})|$ may not be much. Thus, if we apply $B'A'B'$ to $v' = 1_N \otimes v/N$, in the first step B' shrinks the norm by $\lambda_2(C_{D_1})$, then A' takes us from the eigenspace of $\lambda_2(C_{D_1})$ to the eigenspace of $\lambda_k(C_{D_1})$, and in the final step B' shrinks the norm by about $\lambda_k(C_{D_1})$, and we get that $\|B'A'B'v\| \approx \lambda_2 \lambda_k \|v\| \approx \lambda_2^2 \|v\|$, and so the graph $R_{N,D_1} \textcircled{z} C_{D_1}$ exhibits poor expansion. It is conceivable then that for different graphs with more extreme distributions of eigenvalues, it may be that all the other eigenvalues $\lambda_k \ll \lambda_2$, in which case $R_{N,D_1} \textcircled{z} C_{D_1}$ might be a much better expander than C_{D_1}.

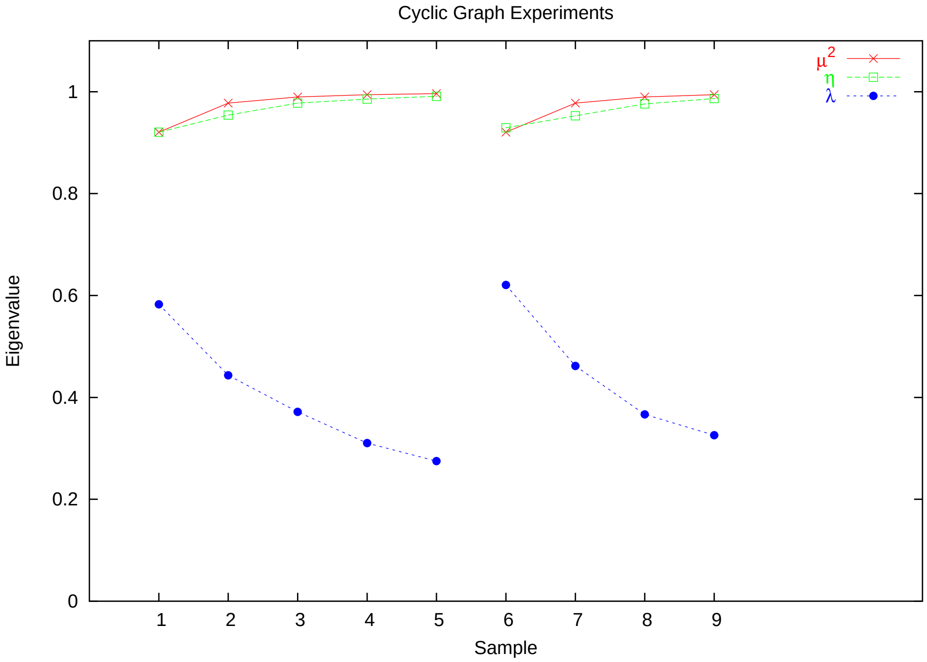

Figure 5.1: Zig-zag product of random regular graphs and cyclic graphs

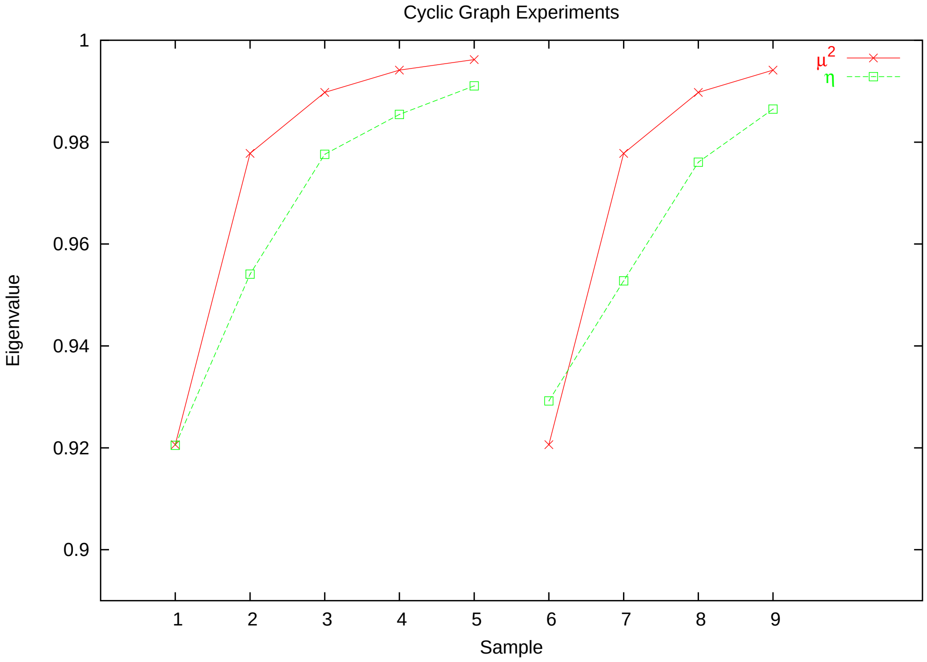

Figure 5.2: Zoomed view of Figure 5.1

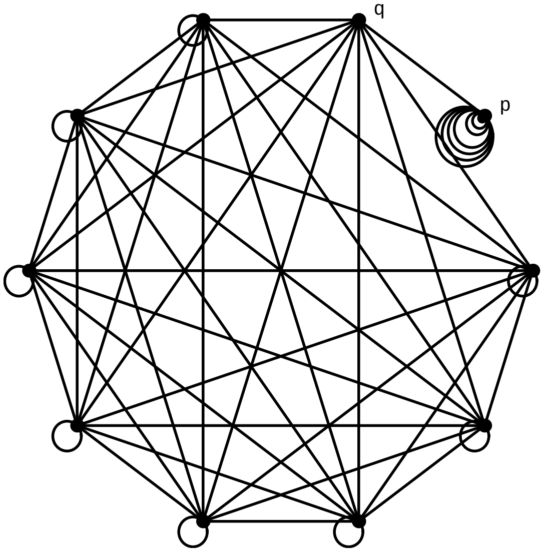

Figure 5.3: The Awful Graph Ω_{10}

Because of this possible concern, we introduce the following construction that has an extremely polarized spectrum. This means that all the eigenvalues of the graph are much less than the second largest. Experiments with these graphs dispel the above concern and provide further support for Conjecture 5.1.3.

Definition 5.4.4. An *awful graph* Ω_N on N vertices has degree $D = N-1$. We construct it by starting with the standard complete graph K_{N-1}, which has $N-1$ vertices and where each pair of vertices share a single edge and there are no self-loops. We add a new vertex p and connect it by a single edge to an arbitrary $q \in [N-1]$. We add $N-2$ self-loops to p and we add a single self-loop to all $i \in [N-1]$ where $i \neq q$.

As an example, Figure 5.3 shows the awful graph Ω_{10}. The following proposition shows that awful graphs have the kind of spectrum we want.

Proposition 5.4.5. *Let Ω_N be an awful graph. Then, letting $D = N-1$ be the degree, it has spectrum of the form $\lambda_1 = 1, \lambda_2 = 1 - \frac{1}{D}, \lambda_3 = -\frac{1}{D}, \lambda_k = 0$ for all $k \geq 4$.*

Proof. We show this by constructing all the eigenvectors. We will use the interpretation of a step in the graph as vertices sending their weights to their neighbors, as discussed in Remark 1.2.8. Let p, q be as in Definition 5.4.4. λ_1 corresponds to uniform.

$\lambda_2 = 1 - \frac{1}{D}$ corresponds to x where

$$x_i = \begin{cases} 1 - \frac{1}{D}, \, i = p \\ 0, \, i == q \\ -\frac{1}{D}, \, \text{else} \end{cases}$$

Since p has $D-1$ self-loops, it receives $(1-\frac{1}{D})^2$ from itself and nothing from q. q receives $\frac{D-1}{D^2}$ to q. q also receives $-\frac{1}{D^2}$ from each of its $D-1$ neighbors, so its final weight is 0. All other $i \neq q$ receive $-\frac{1}{D^2}$ to themselves because of the self-loop, and they receive $-\frac{D-2}{D^2}$ from all the other vertices that are neither p nor q. So i's final weight is $-\frac{1}{D}(1-\frac{1}{D})$.

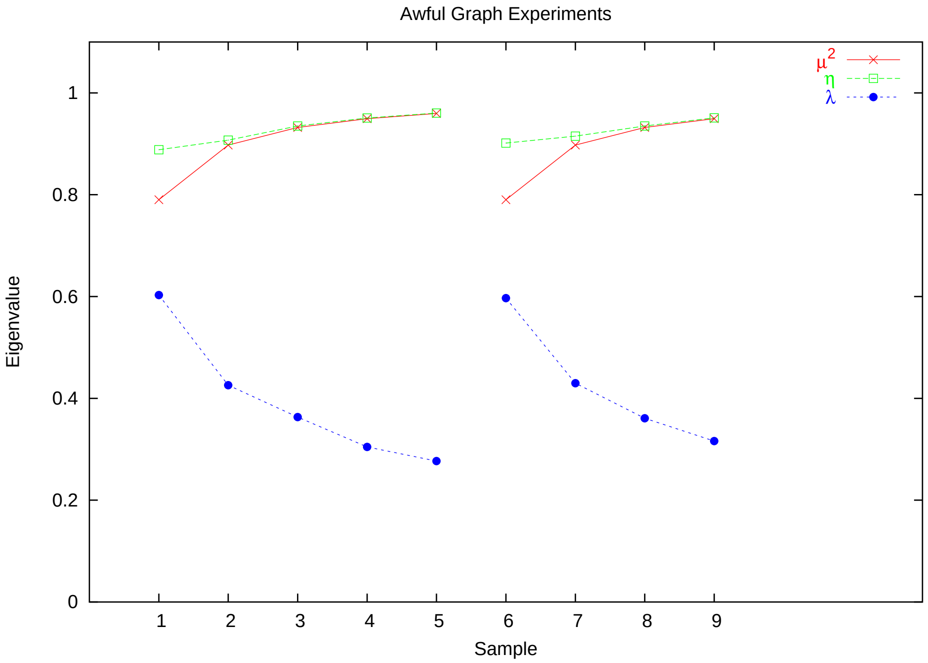

Figure 5.4: Zig-zag product of random regular graphs and awful graphs

$\lambda_3 = -\frac{1}{D}$ corresponds to y where

$$y_i = \begin{cases} 1, \, i = q \\ -\frac{1}{D}, \text{ else} \end{cases}$$

p receives $-\frac{D-1}{D^2}$ from itself and $\frac{1}{D}$ from q, so it finishes with weight $\frac{1}{D^2}$. q receives $-\frac{1}{D^2}$ from each of its D neighbors, so it ends with $-\frac{1}{D}$. All other i receive $-\frac{D-1}{D^2}$ from non-p, q vertices including themselves, and they receive $\frac{1}{D}$ from q, so they end with $\frac{1}{D^2}$.

For λ_k for $k \geq 4$, the eigenvectors are of the form $\nu^{i,j}$ for $i, j \notin \{p, q\}$ where $\nu^{i,j}$ is 1 on vertex i, -1 on vertex j, and 0 elsewhere. It is clear that the effects of the weights of i and j cancel out on all vertices. It is obvious we can take $N - 3$ independent such vectors. □

Figures 5.4 and 5.5 show the awful graph experiments. λ, μ, η are defined as before. For Samples 1-5, the large base graphs are the random regular graphs $R_{50,10}, \ldots, R_{50,50}$ and the small graphs are $\Omega_{10}, \ldots \Omega_{50}$, both in increments of 10. For Samples 6-9, the large base graphs are $R_{80,10}, \ldots R_{80,40}$ and the small graph are $\Omega_{10}, \ldots, \Omega_{40}$.

From Figure 5.4 it seems that even though the expansion of R_{N,D_1} improves, the expansion of $R_{N,D_1} ⓩ \Omega_{D_1}$ is in fact bounded by Ω_{D_1}. Zooming in, we see in Figure 5.5 that $\lambda_2(R_{N,D_1} ⓩ \Omega_{D_1}) \geq \lambda_2(\Omega_{D_1})^2$ holds even though all other eigenvalues of Ω_{D_1} approach 0 as D_1 increases. The two quantities are highly correlated, with a correlation coefficient of 0.9062 over all 9 samples. We also find that $\lambda_2(R_{N,D_1} ⓩ \Omega_{D_1})$ and $\lambda_2(R_{N,D_1})$ are anti-correlated, which implies that R_{N,D_1} is not the cause of $R_{N,D_1} ⓩ \Omega_{D_1}$'s poor expansion.

Surprisingly, it seems that awful graphs are actually not as bad as cyclic graphs in this setting. This gives further evidence that Conjecture 5.1.3 holds.

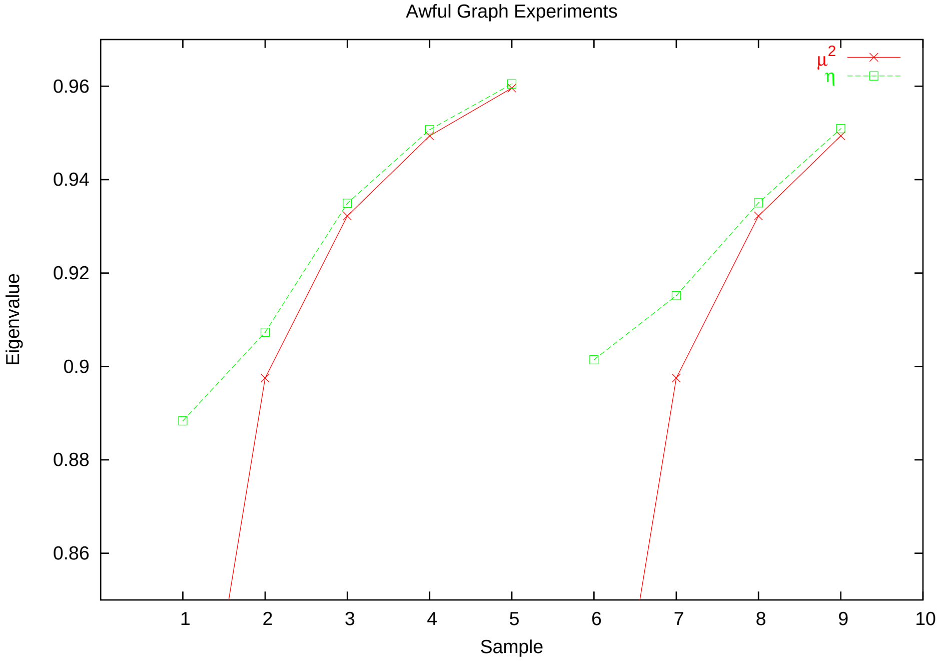

Figure 5.5: Zoomed view of Figure 5.4

One might also wonder about whether the construction-neutral bounds such as the one in Theorem 2.3.2 or Proposition 5.1.2 affect the behavior here. However, note that the degree of the awful graphs is relatively large (e.g. $D_2 = 29$), and so Theorem 2.3.2 and Proposition 5.1.2's lower bounds for the expansion of $R_{N,D_1} ⓩ \Omega_{D_1}$ (e.g. $\frac{2\sqrt{29^2-1}}{29^2} \approx 0.068925$ and $\frac{1}{\sqrt{D_2}} = 0.185695$) are much lower than the computed eigenvalues of $R_{N,D_1} ⓩ \Omega_{D_1}$, and so have little bearing on our experimental analysis.

Future work might try to prove Conjecture 5.1.3. There are various techniques to try to prove this bound in general, such as finding induced subgraphs in the zig-zag product or finding equitable partitions of the graph, which may give a quotient graph that may be analyzed. Unfortunately, at the time of this writing none of these attempts have been fruitful.

Chapter 6

Semi-Squaring

In the process of looking at lower bounds on the second largest eigenvalue of the zig-zag product, we invented the semi-square operation and noted the interesting fact from Lemma 5.3.5 that any semi-square $G^{\lfloor 2 \rceil}$ (Definition 5.3.3) of a graph G has at least as good expansion as G^2. It is therefore an interesting question to see whether for any graph G one can deterministically compute *some* semi-square with better expansion than G^2. Since taking a walk in the semi-square graph does not require additional randomness compared to the regular square because the middle involution is deterministic, this would aid recursive constructions of expanders such as those using the zig-zag product. In this chapter, we study which graphs are amenable to improved expansion by semi-squaring. The material here is entirely new.

6.1 Semi-Square Invariance

Recall that, akin to the semi-direct product, the semi-square of a graph is ambiguous unless we specify the involution σ that is used. Recall from Definition 5.3.8 that a graph G is semi-square invariant if the semi-square has expansion no better than $\lambda_2(G)^2$ no matter which involution σ we choose.

It is an interesting question to determine what kinds of graphs are provably semi-square invariant and what kinds are provably not. Below, we explicitly demonstrate classes that behave in both ways. We begin by demonstrating in Propositions 6.1.2 and 6.1.4 a class of graphs that are semi-square invariant, then we demonstrate in Proposition 6.1.6 a class of graphs that are not semi-square invariant.

Definition 6.1.1. The *generalized complete graph* $K_N^{r,s}$ is the graph on N vertices with $r > 0$ edges between each pair of vertices and $s \geq 0$ self-loops on each vertex. A graph is *square-complete* if its (regular) square is a generalized complete graph. Clearly, any generalized complete graph is also square-complete.

Proposition 6.1.2. *Let G be a connected graph on $N > 2$ vertices with spectrum $\lambda_1 \geq |\lambda_2| \geq \ldots \geq |\lambda_N|$. The following are equivalent:*

1. *G's spectrum has $\lambda_1 = 1$ and $\lambda_i = \lambda_2$ for all $i \geq 2$.*

2. *G's automorphism group is the entire symmetric group S_N.*

3. G is a generalized complete graph $K_N^{r,s}$.

Proof. We show the implications $1 \implies 2 \implies 3 \implies 1$ in order.

$1 \implies 2$: Suppose G's spectrum has $\lambda_1 = 1$ and $\lambda_i = \lambda_2$ for all $i \geq 2$. It is easy to check that $\lambda_2 < 1$ strictly, since G is connected and G cannot be bipartite with this spectrum.

We can decompose any vector x into a vector parallel to uniform $x^{\|}$ and an anti-uniform vector $x^{\perp}$. $\mathrm{Sp}(u)^{\perp}$ is invariant under any permutation σ with matrix P, because P is a rigid motion of $\mathbb{R}^N$ and thus preserves angles, and hence preserves anti-uniformity. $\mathrm{Sp}(u)^{\perp}$ is exactly the eigenspace of λ_2, so by $\mathrm{Sp}(u)^{\perp}$'s invariance under any P, for any anti-uniform $x^{\perp}$ we have $APx^{\perp} = \lambda_2 Px^{\perp}$. Therefore we have

$$P^{-1}APx = P^{-1}AP(x^{\|} + x^{\perp}) = x^{\|} + \lambda_2 x^{\perp} = Ax^{\|} + Ax^{\perp} = Ax$$

Since this holds for all x, we have $P^{-1}AP = A$. Therefore $AP = PA$ and by Remark 5.2.2 σ gives an automorphism of G. So any permutation of the vertices of G is an automorphism, so G's automorphism group is the entire symmetric group S_N.

$2 \implies 3$: Suppose G's automorphism group is the entire symmetric group, but G is not a generalized complete graph. Note that G's automorphism group contains all the transpositions. Clearly the number of self-loops on each vertex must be the same, since otherwise transposing two vertices with different numbers of self-loops would not be an automorphism. Now suppose there exist vertices i, j, k, and say n is the number of edges between i and k and say m is the number of edges between j and k. Suppose $m \neq n$, and take *w.l.o.g.* $m < n$ (where $m = 0$ possibly). Then transposing i and j is clearly not an automorphism since the number of edges to k changes, a contradiction. Therefore the number of edges between any two vertices must be the same. Therefore G is a generalized complete graph.

$3 \implies 1$: We can construct an eigenvector $\nu^{i,j}$ by picking any two vertices i and j and giving i weight 1 and giving j weight -1. Then, using the intuition from Remark 1.2.8 of each vertex "sending" its weight to its neighbors, we see that i and j's weights cancel out on all their neighbors (i.e. all other vertices), and we can compute that the net contribution to one another is $\frac{-r+s}{(N-1)r+s}$. Thus $\nu^{i,j}$ an eigenvector with eigenvalue $\frac{-r+s}{(N-1)r+s}$. Since there are $N-1$ independent such eigenvectors (for example taking all vectors $\nu^{1,j}$ for $1 < j \leq N$), the spectrum has the desired form. □

Corollary 6.1.3. *A graph G is square-complete iff its spectrum has the form $\lambda_1 = 1$ and $\lambda_i = \pm\lambda_2(G)$ for all $i \geq 2$.*

Proof. Since the square of a square-complete graph G is a generalized complete graph, applying Proposition 6.1.2 means that G's spectrum has the above form. □

The following proposition tells us that any graph that has a high multiplicity of eigenvectors for the second largest eigenvalue is semi-square invariant. This makes sense, since this means the eigenspace of $\pm\lambda_2$ has high dimension (i.e. the there are many eigenvectors of λ_2 and $-\lambda_2$), and so any involution in the middle will likely map a vector in that eigenspace back into another vector in the same eigenspace. This proposition will immediately imply that square-complete graphs are semi-square invariant.

Proposition 6.1.4. *Let G be a graph on N vertices with second largest eigenvalue λ_2. Suppose that either the multiplicity of λ_2 or the multiplicity of $-\lambda_2$ is equal to $k > \lfloor \frac{N-1}{2} \rfloor$. Then G is semi-square invariant.*

Proof. Suppose *w.l.o.g.* that the multiplicity of λ_2 is $k > \lfloor \frac{N-1}{2} \rfloor$. Let $v_1, \dots, v_k$ be the eigenvectors corresponding to λ_2, and let $V = \mathrm{Sp}(v_1, \dots, v_k)$ of dimension k. PV also has dimension k since P is a rigid motion. The space of all anti-uniform vectors is invariant under P, so V and PV must have a non-empty intersection. Let $v \in V$ be a vector such that $Pv \in V$, then $Av = \lambda_2 v$ and $APv = \lambda_2 v$. Letting A be the normalized adjacency matrix of G, the eigenvalue bound of Lemma 2.1.2 gives us

$$\max_{x \perp u} \frac{\|APAx\|}{\|x\|} \geq \frac{\|APAv\|}{\|v\|} = \lambda_2^2$$

It is an equality because we know from Lemma 5.3.5 that $\lambda_2(G^{\lfloor 2 \rceil}) \leq \lambda_2^2$ always. □

Corollary 6.1.5. *Any generalized complete graph is semi-square invariant. Any even-sized square-complete graph is semi-square invariant.*

Proof. This corollary is obvious from Proposition 6.1.2, Corollary 6.1.3 and Proposition 6.1.4. □

We can easily show that any graph with a distinct second largest eigenvalue is not semi-square invariant.

Proposition 6.1.6. *Let G be a graph on $N > 2$ vertices, with spectrum $1 = \lambda_1 \geq |\lambda_2| > |\lambda_3| \geq \ldots \geq |\lambda_N|$, i.e. the separation $|\lambda_2| > |\lambda_3|$ is strict. Then G is not semi-square invariant.*

We make use of the following simple lemma in the proof of the proposition.

Lemma 6.1.7. *Let A be the normalized adjacency matrix of a graph G with second largest eigenvalue λ_2 and let P be the matrix of an arbitrary involution. Then $\|APAx\| = \lambda_2^2 \|x\|$ if and only if x and PAx are eigenvectors of G with eigenvalue $\pm\lambda_2$.*

Proof. If x and PAx are eigenvectors with eigenvalue $\pm\lambda_2$, clearly $\|APAx\| = \lambda_2^2 \|x\|$.

Suppose $\|APAx\| = \lambda_2^2 \|x\|$. Recalling the analysis of Lemma 2.1.2, we know that

$$\|APAx\| \leq \lambda_2 \|PAx\| \leq \lambda_2^2 \|x\|$$

since P is length-preserving. Furthermore, the analysis in the proof of Lemma 2.1.2 tells us that equality occurs exactly when PAx and x are eigenvectors of A with eigenvalues $\pm\lambda_2$. □

Proof of Proposition 6.1.6. Let $v \in \mathbb{R}^N$ be the eigenvector of λ_2. We claim that for any anti-uniform $v \perp u$, there exists an involution σ with matrix P such that $Pv \notin \mathrm{Sp}(v)$. Suppose for the sake of contradiction that this were not the case, then since any transposition is also an involution, it would imply that v is an eigenvector of all transpositions. Since the eigenvalues of a transposition are ± 1, this means that any transposition of v is either itself or its negative. This is impossible when $N > 2$ for any $v \perp u$ where $v \neq 0$:

Case 1 If the entries of v take on at least three distinct values, say v_i, v_j, v_k, then we may say *w.l.o.g* that $v_i \neq -v_j$. Letting P be the matrix of the transposition (i, j) would then give us $Pv \neq v$ and $Pv \neq -v$, so $Pv \notin \mathrm{Sp}(v)$.

Case 2 If the entries of v have only two distinct values, say $v_i \neq v_j$, then since $v \perp u$ and $v \neq 0$ it must be that $v_i \neq 0$ and $v_j \neq 0$. So let P be the matrix of the transposition (i, j), then clearly $Pv \neq v$. Also, $Pv \neq -v$ because P fixes some other non-zero entries (since $N > 2$). So $Pv \notin \mathrm{Sp}(v)$.

Let A be the normalized adjacency matrix of G and let P be an involution such that $Pv \notin \mathrm{Sp}(v)$. Now suppose there exists $x \perp u$ such that $\|APAx\| = \lambda_2^2\|x\|$. Lemma 6.1.7 tells us that equality occurs exactly when both PAx and x are eigenvectors of A corresponding to λ_2. Since λ_2 occurs with multiplicity 1, it must be that both x and PAx are multiples of v. But this means $Px \in \mathrm{Sp}(x) \implies Pv \in \mathrm{Sp}(v)$ which contradicts our choice of P. So what we supposed is impossible, and $\|APAx\| < \lambda_2^2\|x\|$ strictly. □

An interesting line of future work would be to prove (or disprove) this proposition when G has second largest eigenvalue with multiplicity ≥ 2. It would be most interesting if one could show some result for any multiplicity $\leq \lfloor \frac{N-1}{2} \rfloor$. A different question would be to get some kind of guaranteed bound on the improvement in eigenvalue. For example, is there some $\gamma < 1$ such that for any G with $\lambda_2(G)$ of multiplicity 1, there exists an involution σ such that $\lambda_2(G^{\lfloor 2 \rceil}) < \gamma\lambda_2(G)^2$?

The next logical step would be to find a deterministic algorithm that can compute a suitable semi-square for a given graph. We refrain from that task here because it seems that it will probably be unproductive. It seems that any algorithm that computes a semi-square with better expansion than $\lambda_2(G)^2$ must take time polynomial in the size of the graph, since the obvious way to do so would be to compute the eigenvectors of G and then find an involution P that maps the $\pm\lambda_2$-eigenspace outside itself. This takes time polynomial in N since computing the eigenvectors takes time polynomial in the size of the normalized adjacency matrix. This is infeasible because, as we mentioned in Subection 3.2.2, in most cases we work with graphs with size exponential in the length of their representations, e.g. a rotation map based on some algebraic or combinatorial construction. Future work might strive to prove that computing a good semi-square from a short representation is difficult.

6.2 Strong Semi-Square Invariance

This section elaborates on an interesting mathematical feature of semi-squaring. Unfortunately this property has no computational importance because of inefficiency concerns.

We saw in Proposition 6.1.4 that any graph G with second largest eigenvalue of high multiplicity is semi-square invariant. However, if G is not square-complete then there always exists a sequence of involutions such that repeatedly semi-squaring the graph according to this sequence will give us a graph with better expansion than by repeatedly squaring. We formalize this below.

Definition 6.2.1. The *n'th semi-square* of a graph G is defined recursively as follows. Fix a sequence of involutions $\sigma_1, \ldots \sigma_n$. The 1'st semi-square is simply $G^{\lfloor 2 \rceil}$ according to σ_1. Then for all $2 \leq i \leq n$, the i'th semi-square $G^{\lfloor 2^i \rceil}$ is the semi-square of $G^{\lfloor 2^{i-1} \rceil}$ according to σ_i.

Definition 6.2.2. A graph G is *strongly semi-square invariant* if for any integer $n \geq 1$ and for any sequence of involutions $\sigma_1, \ldots, \sigma_n$, taking the n'th semi-square $G^{\lfloor 2^n \rceil}$ with respect to that sequence gives a graph with no better expansion than the n'th square. That is, $\lambda_2(G^{\lfloor 2^n \rceil}) = \lambda_2(G)^{2^n}$ for all $n \geq 1$.

Theorem 6.2.3. *A graph is strongly semi-square invariant graphs iff it is square-complete.*

The structure of the proof is as follows. That any square-complete graph is strongly semi-square invariant is easy to show. In the reverse direction, we know that any other kind of graph has an eigenspace corresponding to the second largest eigenvalue of dimension $k < N - 1$. Lemma 6.2.4 will show that any such eigenspace cannot be invariant for all involutions, and therefore we can apply a semi-square such that the dimension of the eigenspace of the second largest eigenvalue in the semi-square graph has dimension $\leq k - 1$. Repeating this process until the eigenspace has dimension 1 allows us to apply Proposition 6.1.4.

Lemma 6.2.4. *Let $V \subsetneq \mathbb{R}^N$ be a non-zero anti-uniform subspace of $\mathbb{R}^N$ of dimension $k < N-1$. Then there exists an involution σ such that V is not invariant under σ.*

Proof. It is sufficient to show that V is not invariant under all the transpositions since the transpositions are themselves involutions.

Suppose there exists a non-zero anti-uniform subspace $V \subset \mathbb{R}^N$ of dimension $k < N - 1$ that is invariant under all transpositions. It is clear that $N > 2$ in order for there to exist such a V. Let $W = V^\perp$ be its orthogonal complement. It is clear that if V is invariant under a transposition σ whose matrix is P, then so is W. This is because any such P is a rigid motion of $\mathbb{R}^N$: if $w \perp v$ then $Pw \perp Pv$.

Let P be the matrix of some transposition (i, j). Since $P^2 = I$ and P is symmetric, it acts on $\mathbb{R}^N$ as a reflection across an invariant hyper-plane R of dimension $N - 1$. The normal vector of this hyper-plane is exactly the (-1)-eigenvector of P, which is 1 on i, -1 on j, and 0 elsewhere. Let us call this (-1)-eigenvector $\nu^{i,j}$.

Claim. V is invariant under the transposition (i, j) iff either V or W is orthogonal to $\nu^{i,j}$.

Suppose both V and W are not orthogonal to $\nu^{i,j}$. It follows that neither is contained in R. Choose $v \in V$ and $w \in W$ that are not contained in R. Then $v - Pv$ and $w - Pw$ are non-zero vectors parallel to $\nu^{i,j}$. They are non-zero because v, w are not in R, and they are parallel to $\nu^{i,j}$ because they are reflected by P:

$$P(v - Pv) = -(v - Pv) \qquad\qquad P(w - Pw) = -(w - Pw)$$

So they are parallel to each other. Thus we have

$$\langle v - Pv, w - Pw\rangle \neq 0 \implies \langle v, w\rangle - 2\langle v, Pw\rangle + \langle Pv, Pw\rangle = -2\langle v, Pw\rangle \neq 0$$

and therefore $PW \neq W$ since we have exhibited $Pw \in PW$ that is not orthogonal to V, and so V and W are not invariant under P.

The reverse direction of the claim is obvious.

We can show that there is no non-zero anti-uniform subspace V of dimension $k < N - 1$ such that for any $\nu^{i,j}$, either V is orthogonal to $\nu^{i,j}$ or $W = V^\perp$ is orthogonal to $\nu^{i,j}$. Suppose for the sake of contradiction that there were such a V. There are two cases.

Case 1 If V were orthogonal to all $\nu^{i,j}$, then any non-zero vector in V would have the same entries on vertices i, j for all i, j, and so it would be uniform. So that subspace is simply the uniform subspace $\mathrm{Sp}(u)$. Its orthogonal complement is the $(N-1)$-dimension subspace of all anti-uniform vectors. Neither of these are anti-uniform subspaces of dimension $k < N - 1$.

Case 2 If V is invariant under all transpositions but is not orthogonal to all $\nu^{i,j}$, then we claim that $W = V^\perp$ must be orthogonal to all $\nu^{i,j}$. It is clear that if V is orthogonal to both $\nu^{i,j}$

and $\nu^{j,k}$, this means all non-zero $v \in V$ have $v_i = v_j$ and $v_j = v_k$, and therefore it follows that $v_i = v_k$ and so V is also orthogonal to $\nu^{i,k}$. Because of this "transitivity" property, we can partition the indices $1 \dots N$ into index classes $C_1, C_2, \dots C_m$, where each $C_r \subset [N]$ and their disjoint union is $[N]$. V is orthogonal to $\nu^{i,j}$ if and only if i, j belong to the same index class C_r.

Now if V is not orthogonal to some $\nu^{i,j}$, then by the above claim it must be that $W = V^\perp$ is orthogonal to $\nu^{i,j}$. So for any $\nu^{i,j}$ such that i, j are in different index classes, it must be that $\forall w \in W,\ w \neq 0$ that $w_i = w_j$. But this immediately implies that $w_i = w_j \forall i, j$ by transitivity: for any $i, j \in C_r$ and $k \in C_s$ where $r \neq s$, we have $w_i = w_k$ and $w_j = w_k$, hence $w_i = w_j$. So W is orthogonal to all $\nu^{i,j}$, and we are just back in the first case.

So $\mathrm{Sp}(u)$ and $\mathrm{Sp}(u)^\perp$ are the unique subspaces invariant under all transpositions, and so there is no anti-uniform subspace of dimension $k < N - 1$ that is invariant under all transpositions. □

Proof of Theorem 6.2.3. To show that if G is square-complete then it is strongly semi-square invariant, it suffices to show that if a graph G is square-complete graph, then any semi-square of G is still square-complete. By Corollary 6.1.3, the spectrum of G has $\lambda_i = \pm\lambda_2$ for all $i \geq 2$, so the eigenspace of $\pm\lambda_2$ is the entire space of anti-uniform vectors, and hence is invariant under any involution. So with A the normalized adjacency matrix of G and P the matrix of any involution, we have that $\|APAx\| = \lambda_2^2 \|x\|$ for any $x \perp u$, and so the spectrum of any $G^{\lfloor 2 \rceil}$ has $\lambda_i = \pm\lambda_2^2$ for all $i \geq 2$. Therefore the second-largest eigenvalue always decreases by a square each time, and therefore $\lambda_2(G^{\lfloor 2^n \rceil}) = \lambda_2(G)^{2^n}$ for all $n \geq 1$.

In the reverse direction, Corollary 6.1.3 tells us that if G is not square-complete then it has eigenvalues smaller in absolute value than the second largest. Let $v_1, \dots, v_k$ be the eigenvectors corresponding to $\pm\lambda_2$ where $k < N - 1$, and set $V = \mathrm{Sp}(v_1, \dots, v_k)$, a subspace of dimension k.

Claim. There exists an involution σ such that the dimension of the eigenspace of $\pm\lambda_2^2$ for the semi-square $G^{\lfloor 2 \rceil}$ with respect to σ is strictly less than k.

From Lemma 6.2.4 we know that for any non-zero anti-uniform subspace with dimension less than $N - 1$, there exists an involution σ on which V is not invariant. Let P be the matrix of such a σ for V, which must exist since the dimension of V is $k < N - 1$.

For the sake of contradiction, suppose that the eigenspace of $\pm\lambda_2^2$ for the semi-square $G^{\lfloor 2 \rceil}$ with respect to our chosen σ has dimension k. Then there exists orthogonal eigenvectors $v'_1, \dots v'_k$ such that $\|APAv'_i\| = \lambda_2^2 \|v'_i\|$. Lemma 6.1.7 says that the v'_i are eigenvectors of A corresponding to $\pm\lambda_2$, and therefore they are a basis for V. We said V is not an invariant subspace of P, so it is impossible that all the $Pv'_i \in V$. But if for some i we have $Pv'_i \notin V$, then Pv'_i is not an eigenvector with eigenvalue $\pm\lambda_2$, and therefore $\|APAv'_i\| < \lambda_2^2 \|v'_i\|$ strictly, a contradiction. So what we supposed is false and the dimension of the eigenspace of $\pm\lambda_2^2$ for the semi-square $G^{\lfloor 2 \rceil}$ according to σ has dimension $< k$.

Thus, we may repeatedly semi-square G, each time choosing an appropriate involution, until we get that the dimension of the eigenspace of $\pm\lambda_2^{2^n}$ is 1 for some repeated semi-square $G^{\lfloor 2^n \rceil}$, at which point we can apply Proposition 6.1.6 to show that $\lambda_2(G^{\lfloor 2^{n+1} \rceil}) < \lambda_2(G)^{2^{n+1}}$. □

Chapter 7

Conclusion

In this thesis we have presented the development of expander graphs, a class of regular undirected multi-graphs that have properties useful in a wide variety of fields within computer science and discrete mathematics. We have traced their development from their algebraic roots to the current work on combinatorial constructions. In the process, we have posed several new questions, and invented and analyzed the semi-square operation, an operation inspired by the analysis of lower bounds for the zig-zag product.

There are still many questions open for future investigation. Primary among them is whether one can use the zig-zag product or some similar combinatorial construction to create infinite families of Ramanujan graphs. Some more specific questions we pose in this thesis include whether one can show that the expansion of zig-zag product necessarily depends on the expansion of both the graphs of which it is composed, and specifically whether that dependence satisfies Conjecture 5.1.3. We also pose several questions about the semi-square, such as proving uniform bounds on the eigenvalue of the semi-square, and determining the complexity of computing the semi-square of a graph, which we believe to be difficult given a small representation of the graph. The area of combinatorial expander construction is still in its nascent stages, and we hope in the future it will blossom into an area with constructions as interesting and efficient as the algebraic constructions we already know.

Bibliography

[1] Agrawal, M., Saxena, N., and Kayal, N. PRIMES is in P. Preprint.

[2] Ajtai, M., Komlós, J., and Szemerédi, E. Sorting in $c \log n$ parallel steps. *Combinatorica 3(1)* (1983), 1–19.

[3] Alon, N. Eigenvalues and expanders. *Combinatorica 6(2)* (1986), 83–96.

[4] Alon, N., Lubotzky, A., and Wigderson, A. Semi-direct product in groups and zig-zag product in graphs: connections and applications. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science* (2001), IEEE, Ed., pp. 630–637.

[5] Alon, N., and Milman, V. D. λ_1, isopermitric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory, Series B 38(1)* (1985), 73–88.

[6] Artin, M. *Algebra.* Prentice Hall, 1991.

[7] Blum, M., Karp, R., Vornberger, O., Papadimitriou, C., and Yannakais, M. The complexity of testing whether a graph is a superconcentrator. *Inform. Process. Letters 13* (1981), 164–167.

[8] Chung, F. On concentrators, superconcentrators, generalizers and nonblocking networks. *Bell Sys. Tech. Journal 58* (1978), 1765–1777.

[9] Chung, F. *Spectral Graph Theory*, vol. 92 of *CBMS Regional Conference Series in Mathematics.* AMS, 1994.

[10] Courant, R., and Hilbert, D. *Methoden der mathematischen Physik*, vol. 1. Springer Verlag, 1924.

[11] Friedman, J. A proof of Alon's second eigenvalue conjecture. In *STOC '03* (2003). To Appear.

[12] Gabber, O., and Galil, Z. Explicit constructions of linear-sized superconcentrators. *J. of Comp. and Sys. Sci. 22(3)* (1981), 407–420.

[13] Gillman, D. A Chernoff bound for random walks on expander graphs. In *IEEE Symposium on Foundations of Computer Science* (1993), pp. 680–691.

[14] Godsil, C., and Royle, G. *Algebraic Graph Theory*, vol. 207 of *Graduate Texts in Mathematics.* Spring Verlag, 2001.

[15] Goldreich, O. A sample of samplers - a computational perspective on sampling (survey). *Electronic Colloquium on Computational Complexity (ECCC) 4(020)* (1997).

[16] Impagliazzo, R., and Zuckerman, D. How to recycle random bits. In *IEEE Symposium on Foundations of Computer Science* (1989), pp. 248–253.

[17] Jimbo, S., and Maruoka, A. Expanders obtained from affine transformations. *Combinatorica 7(4)* (1987), 343–355.

[18] Kahale, N. Eigenvalues and expansion of regular graphs. *Journal of the ACM 42(5)* (1995), 1091–1106.

[19] Leven, D., and Galil, Z. NP-completeness of finding the chromatic index of regular graphs. *J. Alg. 4* (1983), 35–44.

[20] Lubotzky, A., Phillips, R., and Sarnak, P. Ramanujan graphs. *Combinatorica 8(3)* (1988), 261–277.

[21] Margulis, G. A. Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and superconcentrators. *Problems of Information Transmission 24(1)* (1988), 39–46. Translated from the Russian *Problemy Peredachi Informatsii.*

[22] Nilli, A. On the second eigenvalue of a graph. *Discrete Mathematics 91(2)* (1991), 207–210.

[23] Pinsker, M. On the complexity of a concentrator. In *7th International Teletraffic Conference* (1973), pp. 318/1–318/4.

[24] Rabin, M. Probabilistic algorithm for testing primality. *Journal of Number Theory 12* (1980), 128–138.

[25] Reingold, O., Vadhan, S., and Wigderson, A. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In *IEEE Symposium on Foundations of Computer Science* (2000), pp. 3–13.

[26] Sipser, M., and Spielman, D. Expander codes. In *IEEE Symposium on Foundations of Computer Science* (1994), pp. 566–576.

[27] Tanner, M. Explicit construction of concentrators from generalized N-gons. *SIAM Journal on Algebraic Discrete Methods 5(3)* (1984), 287–293.

[28] Valiant, L. Graph theoretic properties in computational complexity. *J. Comp. and Sys. Sci. 13* (1976), 278–285.

[29] Wigner, E. On the distribution of the roots of certain symmetric matrices. *Annals of Mathematics 67* (1958), 325–327.