RIGIDITY AND RANK OF GROUP-CIRCULANT MATRICES

MICHAEL YANG

Abstract. Given a finite group G, a ring Λ, and a function $f : G \to \Lambda$, a G-circulant matrix of f is a $|G| \times |G|$ matrix M with rows and columns indexed by the elements of G for which $M_{xy} = f(xy)$ for all $x, y \in G$. We study the fundamental properties of G-circulants when Λ is an algebraically closed field with characteristic coprime to $|G|$.

We begin by proving new results about the matrix rigidity of G-circulants for nonabelian G, which are the first of its kind. We show that for any sequence of finite groups G_i whose abelian normal subgroups have sufficiently small index, the family of G_i-circulants is not Valiant-rigid. Furthermore, we show that this result applies for families of groups $\{G_i\}_i$ whose representations are bounded above in degree.

Next, we exhibit a formula for the rank of any G-circulant in terms of the decomposition of its corresponding function $f : G \to \Lambda$ into the matrix coefficients of the irreducible representations of G. While this was known to Diaconis, we present a more elementary proof that avoids the full strength of Schur Orthogonality.

We then apply this formula to the case of G-circulants for cyclic G. Through this, we generalize a theorem of Chen, providing a necessary and sufficient criterion for when zero-one circulants are always nonsingular. Additionally, we answer an open problem about singular circulant digraphs posed by Lal–Reddy and give a probabilistic estimate for the regularity of zero-one singular circulant matrices.

Lastly, we investigate orthogonal representations of graphs. Given a finite, simple graph G, we provide a novel lower bound for the minimal dimension in which a faithful orthogonal representation for G exists. Furthermore, we use our bound to determine the aforementioned minimal dimension for an infinite family of Kneser graphs up to a constant factor.

1. Introduction

For an arbitrary finite group G and ring Λ, fix a function $f : G \to \Lambda$. The G-circulant matrix of f is the $|G| \times |G|$ matrix M with rows and columns indexed by the elements of G such that the entry in the row corresponding to x and the column corresponding to y is equal to $f(xy)$ for all $x, y \in G$.

The rank of G-circulant matrices has long been of interest. In 2017, Croot, Lev, and Pach [CLP17] resolved a longstanding problem by providing an upper bound on the ranks of circulant matrices for functions $f : \mathbb{F}_q^n \to \mathbb{F}_q$, using this as the key fact to prove that arithmetic-progression-avoiding subsets of $\mathbb{Z}_4^n$ are exponentially small. This approach was subsequently generalized to all abelian groups by Ellenberg and Gijswijt [EG17].

Date: December 19, 2023.

Key words and phrases. rank, circulant, matrix rigidity, representation.

Circulant ranks have also been used in the context of matrix rigidity. Matrix rigidity, introduced by Valiant in 1977 in his seminal paper [Val77], is a quantitative metric of how far a matrix is from being low rank (see Section 2 for a precise definition). It has long been an open question to find an explicit family of sufficiently rigid matrices; doing so would have significant ramifications in coding and complexity theory, proving unconditional super-linear circuit lower bounds for circuits with logarithmic depth [Val77]. In 2017, Dvir–Edelman [DE17] adapted the core result from Croot–Lev–Pach [CLP17] to establish an asymptotic bound on the matrix rigidity of $\mathbb{F}_q^n$-circulants. Recent other work on matrix rigidity has mainly focused on proving that certain families of matrices are not sufficiently rigid, most notably culminating in the work of Dvir–Liu [DL19], which showed that all families of abelian circulants are not rigid.

Up until now, both circulant ranks and matrix rigidity have primarily been studied in the context of abelian groups. We provide generalizations in the nonabelian setting and categorize our results into four main sections.

First, we demonstrate the non-rigidity of all families of G-circulants for nonabelian groups G which have a sufficiently large abelian normal subgroup. This is the first rigidity result that applies to G-circulants for nonabelian G. Specifically, we prove the following:

Theorem 1.1. *Consider an infinite family of groups $\{G_i\}_i$ and functions $\{f_i : G_i \to \Lambda\}_i$. If every sufficiently large G_i has an abelian normal subgroup with index bounded by some fixed polynomial in $\log(|G_i|)$, then the family of G-circulant matrices of the functions $\{f_i\}_i$ is not rigid over $\mathbb{C}$. In particular, this holds for any family $\{G_i\}_i$ of groups where the degrees of the irreducible representations of the G_i are bounded above by some constant.*

Group-circulants of this form admit a special structure, where we can rearrange the rows and columns to induce more regularity among the entries; we can subsequently apply the results of Dvir–Liu [DL19] to conclude that the matrices cannot be rigid. A more precise formulation and proof can be found in Theorem 2.8.

Next, we give a new proof of a result found in Diaconis 1990 [Dia90] to compute explicitly the rank of all G-circulant matrices for nonabelian G in terms of the linear representation theory of G. Our proof is more elementary in the sense that it does not rely on the full strength of the Schur Orthogonality relations. To state it, we need a definition from representation theory: given an irreducible representation ρ of a finite group G over a ring Λ, define the *(i,j)th matrix coefficient* $\rho_{ij}(x)$ to be the function $\rho_{ij} : G \to \Lambda$ such that $\rho_{ij}(g)$ for any $g \in G$ returns the (i,j)th entry of the matrix $\rho(g)$.

If Λ is an algebraically closed field whose characteristic does not divide $|G|$, then it is known that the matrix coefficients of the irreducible representations of G are a basis for the space of functions from G to Λ [Eti+11, Prop. 4.7.1][Ter99, Thm. 2(1)].

The idea of our proof is that when a function viewed as a linear combination of the matrix coefficients of the irreducible representations of G, we can take advantage of the multiplicativity of the representations to prove that an expanded version of

the linear combination is in a sense "optimal" in that the only way to reduce the number of summands is through obvious factorizations, allowing us to calculate the rank explicitly. In particular, we prove the following theorem:

Theorem 1.2. *For any group G, field Λ with characteristic coprime to $|G|$, and function $f : G \to \Lambda$, express f in the form*

$$f(x) = \sum_{\rho \in I} \left(\sum_{1 \le i,j \le \deg \rho} c_{\rho,i,j} \rho_{ij}(x) \right)$$

where I is the set of isomorphism classes of irreducible representations of G and where $c_{\rho,i,j} \in \Lambda$. Then, the rank of the $|G| \times |G|$ circulant matrix defined by $f(xy)_{ij}$ is equal to

$$\sum_{\rho \in I} \left[(\deg \rho) \ \mathit{rank} \left(\begin{bmatrix} c_{\rho,1,1} & c_{\rho,1,2} & \cdots & c_{\rho,1,N} \\ c_{\rho,2,1} & c_{\rho,2,2} & \cdots & c_{\rho,2,N} \\ \vdots & \vdots & \ddots & \vdots \\ c_{\rho,N,1} & c_{\rho,N,2} & \cdots & c_{\rho,N,N} \end{bmatrix} \right) \right].$$

Lastly, in the fourth section, we use this result to provide several applications to, and generalizations from, existing results regarding classical circulant matrices. By viewing these as G-circulant matrices in the specific context of $G = \mathbb{Z}/n\mathbb{Z}$, Theorem 3.11 provides an explicit criterion for which classical circulants are invertible. We will use this to prove the following statement, which generalizes a 2021 result of Chen [Che21]:

Theorem 1.3. *An $n \times n$ zero-one circulant with k ones in the first row is always invertible if and only if at least one of k and $n - k$ cannot be expressed as a linear combination of the prime divisors of n.*

A more precise statement and proof of this theorem is found in Theorem 4.6.

In the same section, we also resolve an open problem by Lal–Reddy regarding when certain classes of directed circulant graphs are singular. To answer this question, it suffices to determine the invertibility of a certain family of circulant matrices: we do so in Theorem 4.7. We end this section with a brief result on the probability of an $n \times n$ zero-one circulant matrix to be singular.

Our last section covers orthogonal representations of graphs. Loosely speaking, an n-dimensional orthogonal representation of a graph is an assignment of a vector to each of the graph's vertices where orthogonal vectors correspond to adjacent vertices. Orthogonal representations of graphs have many applications: they were first introduced by Lovász to determine the Shannon Capacity of a graph, are used to detect hidden quantum variables, and appear naturally in the setting of partition logics [Svo20].

In Section 5, we provide a novel lower bound of the minimal dimension for which an orthogonal representation (satisfying certain nondegeneracy properties) exists for any graph. Notating this minimal dimension as $\xi_F(G)$, we prove the following:

Theorem 1.4. *For any finite, simple, undirected graph G,*

$$\frac{n^2}{4(n^2 - 2|E(G)|)} \le \xi_F(G).$$

We then show this bound is tight up to a constant factor for certain families of Kneser graphs, building on a line of work initiated by Golovnev–Haviv [GH20]. A precise statement is given in Section 5 and Theorem 5.7.

2. Matrix Rigidity of Group-Circulants

In this section, we prove a statement about the matrix rigidity of G-circulants for special G.

2.1. Background.

The intuitive idea behind matrix rigidity is that a matrix is *rigid* if it differs from *any* low-rank matrix by a large number of entries.

More technically, fix a matrix M of dimension $n \times n$, and let $0 \leq r \leq n$ be an integer. Now, we define the quantity $R_M(r)$ as the minimum number of entries we need to change in M in order to produce a matrix with rank at most r.

Now, consider an infinite family of square matrices $\mathcal{M}$. We say that this family is *Valiant-rigid* if for all sufficiently large $M \in \mathcal{M}$, there exists a constant $\varepsilon > 0$ such that

$$\frac{N^{1+\varepsilon}}{R_M\left(\frac{N}{\log\log N}\right)}$$

is bounded uniformly in N, where $N = \dim M$.

If a family of matrices is Valiant-rigid, then it satisfies the following property:

Theorem 2.1 (Valiant 1977, Corollary 6.3)**.** *Suppose Λ is a field. If M is a Valiant-rigid $N \times N$ matrix, then the linear map corresponding to M cannot be computed by circuits of size $O(N)$ and depth $O(\log N)$.*

This is valuable in complexity theory, as it establishes unconditional superlinear lower bounds on circuits with logarithmic depth. Thus, it is of interest to determine whether or not families of matrices are Valiant-rigid.

In the notation of Dvir–Liu [DL19], we now define a weaker notion of rigidity. Define the *regular rigidity* $\mathrm{r}_M(r)$ of a matrix M as the minimum number s such that it possible to change at most s entries in each row and column of M to produce a matrix of rank at most r.

Definition 2.2. We say a family $\mathcal{M}$ of matrices is *quasipolynomially non-rigid* (QNR) over a field $\mathbb{F}$ if there are constants $c_1, c_2 > 0$ such that for any $\varepsilon > 0$, all sufficiently large matrices $M \in \mathcal{M}$ satisfy

$$\mathrm{r}_M^{\mathbb{F}}\left(\frac{N}{\exp(\varepsilon^{c_1}(\log N)^{c_2})}\right) \leq N^{\varepsilon},$$

where M is an $N \times N$ matrix.

Definition 2.3. We say an $N \times N$ matrix A (or, more precisely, a family $\mathcal{M}$ of matrices in N) has *QNR rank* over a field $\mathbb{F}$ if there exist constants $c_1, c_2 > 0$ such that for all sufficiently large $M \in \mathcal{M}$,

$$\operatorname{rank}(M) \sim O\left(\frac{N}{\exp(\varepsilon^{c_1}(\log N)^{c_2})}\right).$$

Definition 2.4. Given a group G, define the *Fourier matrix* F of G to be the $|G| \times |G|$ matrix, with columns indexed by the elements of G and rows indexed by the matrix coefficients of the irreducible representations of G, such that the matrix entry in the row corresponding to ρ_{ij} and the column corresponding to $g \in G$ is equal to $\rho_{ij}(g)$.

Dvir–Liu [DL19] established that for all abelian G, any family of G-circulant matrices is quasipolynomially non-rigid.

Theorem 2.5 (Dvir–Liu 2019, Theorem 1.5)**.** *Let G be an abelian group. The family of G-circulant matrices is QNR over $\mathbb{C}$. For a finite field $\mathbb{F}_q$, if $\gcd(|G|, q) = 1$, then the family of G-circulant matrices is QNR over $\mathbb{F}_q$.*

Notice that if a family of matrices is quasipolynomially non-rigid, it is also not Valiant-rigid.

In addition, we need two other results for our proof in this section. We first consider a result from Isaacs–Passman [IP64]:

Theorem 2.6 (Isaacs–Passman 1964, Corollary 3.6)**.** *There exists a function $k : \mathbb{Z}^+ \to \mathbb{Z}^+$ with the following property. If G is any group all of whose irreducible representations have degree $\leq n$, then G has an abelian normal subgroup N of index less than or equal to $k(n)$.*

Secondly, we have the following result due to Meyer [Mey73].

Theorem 2.7 (Meyer 1973, Theorem 4.1)**.** *Consider an $m \times n$ matrix M over any field partitioned into four blocks, which we call A, C, R, D. Then, there exist matrices E_A, F_A, A^-, E_W, and F_Y for which*

$$\operatorname{rank}(M) = \operatorname{rank}(A) + \operatorname{rank}(E_A C) + \operatorname{rank}(E_W(D - RA^- C)F_Y) + \operatorname{rank}(RF_A).$$

In particular,

$$\operatorname{rank}(M) \leq 2(\operatorname{rank}(A) + \operatorname{rank}(C) + \operatorname{rank}(R) + \operatorname{rank}(D)).$$

2.2. **Quasipolynomial Nonrigidity of Group-Circulants.** In this section, we will prove the following theorem.

Theorem 2.8. *Consider an infinite family $\{G_i\}_i$ of groups and functions $\{f_i : G_i \to \Lambda\}_i$. If every sufficiently large G_i has an abelian normal subgroup with index bounded by some fixed polynomial in $\log(|G_i|)$, then the family of G-circulant matrices of the functions $\{f_i\}_i$ is quasipolynomially non-rigid over $\mathbb{C}$.*

In fact, Theorem 2.8, which is stated in terms of group theory, has the following corollary in the language of representation theory.

Corollary 2.9. *Consider again an infinite family $\{G_i\}_i$ of groups, and for each i, let F_i be the Fourier matrix of G_i: that is, the $|G_i| \times |G_i|$ matrix whose entries are the matrix coefficients of its irreducible representations. Suppose that for some constant c, we have $\deg(\rho) < c$ for all i and $\rho \in \operatorname{Irr}(G_i)$, the set of isomorphism classes of the irreducible representations of G_i. Then, the family of matrices $\{F_i\}_i$ is quasipolynomially non-rigid over $\mathbb{C}$. Furthermore, the family of G_i-circulants is quasipolynomially non-rigid as i gets large.*

First, we will prove Theorem 2.8. The main idea is the following: group-circulant matrices over groups with large abelian normal subgroups exhibit a lot of structure, and we can permute the rows and columns of the corresponding G-circulant matrix to induce more familiar structure which we can work with.

In particular, for all sufficiently large G_i, let N_i denote an abelian normal subgroup of G_i with index bounded by some fixed polynomial in $\log(|G_i|)$. We index the rows and columns of each G_i-circulant in the following way:

- The first $|N_i|$ rows/columns correspond to the elements in the coset eN_i.
- The next $|N_i|$ rows/columns correspond to another, distinct coset hN_i for $h \in G$, and so on.

In short, we index the rows and columns of each G_i-circulant by the cosets of N_i.

This indexing has several interesting properties. Specifically, consider the entry in a G_i-circulant matrix corresponding to a row element which is in the coset aN_i and a column element in the coset bN_i for elements $a, b \in G_i$. Then, that corresponding entry in the matrix is going to be of the form $f(g)$, where $f : G_i \to \Lambda$ is an arbitrary function and $g \in G_i$ is a group element in the coset abN_i (since N_i is normal in G_i).

For the sake of making this explicit, we will first establish how we are notating G-circulants. Index the group elements of G as $e, g_1, g_2, \ldots, g_{|G|-1}$, and index the rows (from top to bottom) and columns (from left to right) in that order. Each G-circulant then takes on the following form, where $f : G \to \Lambda$ is allowed to be any function:

$$\begin{bmatrix} f(e) & f(g_1) & f(g_2) & \cdots & f(g_{|G|-1}) \\ f(g_1) & f(g_1^2) & f(g_1g_2) & \cdots & f(g_1g_{|G|-1}) \\ f(g_2) & f(g_2g_1) & f(g_2^2) & \cdots & f(g_2g_{|G|-1}) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ f(g_{|G|-1}) & f(g_{|G|-1}g_1) & f(g_{|G|-1}g_2) & \cdots & f(g_{|G|-1}^2) \end{bmatrix}$$

For the sake of convenience, we will drop all the fs in the matrix, leaving just the product of elements; it is clear how to reconstruct a G-circulant from this pared-down matrix.

Now, when we index the rows and columns by normal subgroups, we are changing the orders of the rows and columns of the original G-circulant. In what follows, let $[G_i : N_i] = C$, and let $a_1, a_2, \ldots, a_C$ be group elements in G_i such that the cosets $a_1N_i, a_2N_i, \ldots, a_CN_i$ are all distinct. Then, we partition the matrix into a $C \times C$ block matrix in the following way:

$$\begin{array}{c} \\ a_1N_i \\ a_2N_i \\ a_3N_i \\ \vdots \\ a_CN_i \end{array} \begin{array}{c} \begin{array}{ccccc} a_1N_i & a_2N_i & a_3N_i & \cdots & a_CN_i \end{array} \\ \begin{pmatrix} a_1^2N_i & a_1a_2N_i & a_1a_3N_i & \cdots & a_1a_CN_i \\ a_2a_1N_i & a_2^2N_i & a_2a_3N_i & \cdots & a_2a_CN_i \\ a_3a_1N_i & a_3a_2N_i & a_3^2N_i & \cdots & a_3a_CN_i \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_Ca_1N_i & a_Ca_2N_i & a_Ca_3N_i & \cdots & a_C^2N_i \end{pmatrix} \end{array}.$$

In this partition, each block has the following structure: for some indexed list of elements $n_1, n_2, \ldots, n_{|N_i|}$ of N_i, the block corresponding to the row coset a_XN_i and the column coset a_YN_i has rows indexed $a_Xn_1, a_Xn_2, \ldots, a_Xn_{|N_i|}$ and columns

indexed $a_Y n_1, a_Y n_2, \ldots, a_Y n_{|N_i|}$. Adopting notation from our writing of circulant matrices, each block has the following form:

$$\begin{bmatrix} a_X n_1 a_Y n_1 & a_X n_1 a_Y n_2 & \cdots & a_X n_1 a_Y n_{|N_i|} \\ a_X n_2 a_Y n_1 & a_X n_2 a_Y n_2 & \cdots & a_X n_2 a_Y n_{|N_i|} \\ \vdots & \vdots & \ddots & \vdots \\ a_X n_{|N_i|} a_Y n_1 & a_X n_{|N_i|} | a_Y n_2 & \cdots & a_X n_{|N_i|} a_Y n_{|N_i|} \end{bmatrix}.$$

Notice that each of the elements above is in the coset $a_X a_Y N_i$. Actually, we claim that the above matrix is isomorphic to an N_i-circulant. By this, we mean that if we multiply every element above by $(a_X a_Y)^{-1}$ on the left, then the resulting structure will be exactly equal to the structure in an N_i-circulant.

Doing this, we see that the matrix becomes

$$\begin{bmatrix} a_Y^{-1} n_1 a_Y n_1 & a_Y^{-1} n_1 a_Y n_2 & \cdots & a_Y^{-1} n_1 a_Y n_{|N_i|} \\ a_Y^{-1} n_2 a_Y n_1 & a_Y^{-1} n_2 a_Y n_2 & \cdots & a_Y^{-1} n_2 a_Y n_{|N_i|} \\ \vdots & \vdots & \ddots & \vdots \\ a_Y^{-1} n_{|N_i|} a_Y n_1 & a_Y^{-1} n_{|N_i|} | a_Y n_2 & \cdots & a_Y^{-1} n_{|N_i|} a_Y n_{|N_i|} \end{bmatrix}.$$

But this is just an N_i-circulant with rows indexed $a_Y^{-1} n_k a_Y$ for $1 \leq k \leq |N_i|$ and columns indexed by $n_1, n_2, \ldots, n_{|N_i|}$ in that order.

This suggests the following revisualization of an arbitrary G_i-circulant with rows indexed in the above way: view it as a $C \times C$ block matrix, where each block is a matrix defined by a function on the elements of some coset of N_i.

We will now show that these G_i-circulants are quasipolynomially non-rigid. By Theorem 2.5, we can alter at most $|N_i|^\varepsilon$ entries in each row and column in each of the individual circulants to create a matrix with QNR rank. Doing so for each of the blocks in the G_i-circulant, we get an upper bound of

$$|G_i|^2 |N_i|^{-1+\varepsilon} \sim O(|G_i|^{1+\varepsilon'})$$

entries changed. In fact, we have changed $O(|G_i|^{\varepsilon'})$ entries in every row and column.

We will now show that the rank of the entire modified G_i-circulant is QNR by using a generalization of Theorem 2.7 to partitions of greater size.

Lemma 2.10. *Let $m \geq 2$ be a positive integer, and consider an arbitrary $m \times m$ block matrix M over any field with blocks $B_1, B_2, \ldots, B_{m^2}$. Then,*

$$\operatorname{rank}(M) \leq 2m \left(\sum_{j=1}^{m^2} \operatorname{rank}(B_j) \right).$$

Proof. Consider the sequence $(k_n)_{n \geq 1}$ defined by the following recursion:

- $k_1 = 1$,
- $k_{2n} = 2k_n$ for all positive integers n, and
- $k_{2n-1} = 2k_n$ for all positive integers $n \geq 2$.

The first few terms of the sequence are $1, 2, 4, 4, 8, \ldots$. More specifically, we can verify that k_n is the largest power of two strictly less than $2n$.

To prove our lemma, we will prove the stronger statement that

$$\operatorname{rank}(M) \leq k_m \left(\sum_{j=1}^{m^2} \operatorname{rank}(B_j) \right).$$

Since $k_m \leq 2m$ for all m, this will prove the original claim.

To prove the stronger statement, we will use strong induction on m, with the base case $m = 2$ being Theorem 2.7.

For the strong inductive step, we will use casework on the parity of m. When m is even, partition the $m \times m$ block matrix into four quadrants, or $m/2 \times m/2$ "super-blocks," S_1, S_2, S_3, and S_4. Theorem 2.7 implies that

$$\operatorname{rank}(M) \leq 2(\operatorname{rank}(S_1) + \operatorname{rank}(S_2) + \operatorname{rank}(S_3) + \operatorname{rank}(S_4)).$$

Applying the strong inductive hypothesis to each of the super-blocks, we get

$$\operatorname{rank}(S_\ell) \leq 2\left(k_{m/2}\right) \left(\sum_{B \in S_\ell} \operatorname{rank}(B) \right) = k_m \left(\sum_{B \in S_\ell} \operatorname{rank}(B) \right)$$

for all super-blocks S_ℓ and blocks $B \in S_\ell$. Substituting the above equation into the inequality for $\operatorname{rank}(M)$ gives the result.

Next, we consider the case where m is odd; the base case in this scenario is the vacuous $m = 1$. In the general case, we again partition the $m \times m$ matrix into four roughly even quadrants; more precisely, we partition the matrix into super-blocks S_1, S_2, S_3, and S_4 with dimensions $\frac{m+1}{2} \times \frac{m+1}{2}$, $\frac{m+1}{2} \times \frac{m-1}{2}$, $\frac{m+1}{2} \times \frac{m-1}{2}$, and $\frac{m-1}{2} \times \frac{m-1}{2}$, respectively. Again, we have

$$\operatorname{rank}(M) \leq 2(\operatorname{rank}(S_1) + \operatorname{rank}(S_2) + \operatorname{rank}(S_3) + \operatorname{rank}(S_4)).$$

Notice that while S_2 and S_3 are not square, their ranks are going to be at most the rank of any contiguous $\frac{m-1}{2} \times \frac{m-1}{2}$ submatrix contained in the respective super-blocks. Hence, applying the strong inductive hypothesis, we have the following inequalities:

$$\operatorname{rank}(S_1) \leq k_{\frac{m+1}{2}} \left(\sum_{B \in S_1} \operatorname{rank}(B) \right),$$

$$\operatorname{rank}(S_2) \leq k_{\frac{m-1}{2}} \left(\sum_{B \in S_2} \operatorname{rank}(B) \right),$$

$$\operatorname{rank}(S_3) \leq k_{\frac{m-1}{2}} \left(\sum_{B \in S_3} \operatorname{rank}(B) \right),$$

$$\operatorname{rank}(S_4) \leq k_{\frac{m-1}{2}} \left(\sum_{B \in S_4} \operatorname{rank}(B) \right).$$

Plugging this into the rank inequality for the original matrix and using the fact that $k_{\frac{m-1}{2}} \leq k_{\frac{m+1}{2}}$, we get

$$\operatorname{rank}(M) \leq 2k_{\frac{m+1}{2}} \left(\sum_{B \in M} \operatorname{rank}(B) \right) = k_m \left(\sum_{B \in M} \operatorname{rank}(B) \right),$$

proving the result. □

Recall that we would like to show that the modified G_i-circulant is of QNR rank. To this end, consider the function $s : \mathbb{Z} \to \mathbb{Z}$ given by $s(|G_i|) = [G_i : N_i]$ for all indices i. We know that the function s is asymptotically bounded by a fixed polynomial in $\log(|G_i|)$.

The modified G_i-circulant, call it M_i, is a block matrix with dimensions $s(|G_i|) \times s(|G_i|)$, where each block $B \in M_i$ is isomorphic to an N_i-circulant. Moreover, we know that each of these B has QNR rank. Applying Lemma 2.10, we know that

$$\operatorname{rank}(M_i) \leq 2s(|G_i|) \sum_{B \in M_i} \operatorname{rank}(B).$$

We would like to show that this has QNR rank with respect to M_i.

Plugging in the ranks of the individual blocks B and using the fact that there are $s(|G_i|)^2$ such blocks, we have

$$\operatorname{rank}(M_i) \leq 2s(|G_i|)^3 O\left(\frac{|G_i|/s(|G_i|)}{e^{\log(|G_i|/s(|G_i|))^\delta}}\right)$$

for some constant $0 < \delta < 1$. To finish the proof of the theorem, we claim that this term is of the order

$$O\left(\frac{|G_i|}{e^{\log(|G_i|)^{\delta'}}}\right)$$

for a constant $0 < \delta' < 1$, which would prove that M_i has QNR rank.

To see why this is, it suffices to check that

$$s(|G_i|)^3 \cdot \left(\frac{|G_i|/s(|G_i|)}{e^{\log(|G_i|/s(|G_i|))^\delta}}\right) = \frac{|G_i| \cdot s(|G_i|)^2}{e^{\log(|G_i|/s(|G_i|))^\delta}}$$

is asymptotically at most

$$\frac{|G_i|}{e^{\log(|G_i|)^{\delta'}}}.$$

In fact, we will prove the stronger statement that for any constant $0 < \delta < 1$, it is possible to choose a constant $0 < \delta' < 1$ such that

$$\frac{|G_i| \cdot s(|G_i|)^2}{e^{\log(|G_i|)^\delta}} < \frac{|G_i|}{e^{\log(|G_i|)^{\delta'}}}.$$

Indeed, this comes down to choosing a δ' such that

$$s(|G_i|)^2 < e^{\log(|G_i|)^\delta - \log(|G_i|)^{\delta'}}.$$

This is equivalent to showing that for any δ, the function $e^{\log(|G_i|)^\delta}$ grows faster than any polynomial in $\log(|G_i|)$. In fact, we only need to show that the former function grows faster than $\log(|G_i|)^\alpha$ for all positive integers α, which is apparent by taking the logarithm of both sides. This proves Theorem 2.8.

Next, we will prove Corollary 2.9.

Firstly, in the notation of Corollary 2.9, Theorem 2.6 implies that in the family $\{G_i\}_i$, as i gets sufficiently large, there is a constant C_0 such that there is an abelian normal subgroup N_i of G_i with order at least G_i/C_0. The quasipolynomial nonrigidity of these G_i-circulants then follows from Theorem 2.8.

We have thus proved that G_i-circulant matrices are all QNR for sufficiently large indices i. Now, we will show that the family $\{F_i\}_i$ of matrices is also QNR. To do this, we will use the following claim, which closely resembles Lemma 2.21 in Dvir–Liu's paper [DL19, Lem. 2.2.1].

Claim 2.11 (Approximating Rigidity-Measuring Functions)**.** Consider square matrices A, B, C with A unitary, C block-diagonal, and $B = A^*CA$. Let k be the dimension of the largest block in C. If $r_A(r) \leq s$ for some s, then $r_B(2r) \leq ks^2$.

The proof of this claim is analogous to the proof of Lemma 2.21 in Dvir–Liu [DL19, Lem. 2.2.1]. In particular, it will rely on the following lemma, which is easy to see via expansion:

Lemma 2.12. *Consider square matrices M and N of the same dimension. If M has no more than u nonzero entries in each row and column and if N has no more than v nonzero entries in each row and column, then the matrix product MN has no more than uv nonzero entries in each row and column.*

Now, we will prove the claim.

Proof. Let E be the matrix with at most s nonzero entries in each row and column such that $\operatorname{rank}(A - E) \leq r$. Then,

$$B - E^*CE = A^*C\underbrace{(A - E)}_{\text{rank } \leq r} + \underbrace{(A^* - E^*)}_{\text{rank } \leq r}CE.$$

Hence, the rank of $B - E^*CE$ is at most $2r$. But notice that E^* and E both have at most s nonzero entries in each row and column, while C has at most k such entries. Thus, by Lemma 2.12, the matrix product has at most ks^2 nonzero entries in each row and column, and the claim follows. □

Now, notice that the matrix F_i block-diagonalizes every G_i-circulant, with each block being capped at dimensions $c \times c$ (where, in the notation of Corollary 2.9, c is the upper bound on the degrees of the irreducible representations of the G_i). In particular, since G_i-circulants are QNR, we have that

$$\mathrm{r}_M\left(O\left(\frac{|G_i|}{e^{\log(|G_i|)^c}}\right)\right) \leq |G_i|^{\varepsilon}$$

for any $\varepsilon > 0$ and where M is any G_i-circulant. By the claim above, this directly implies that for the corresponding matrix $M' = F_i$, we have

$$\mathrm{r}_{M'}\left(O\left(\frac{|G_i|}{e^{\log(|G_i|)^c}}\right)\right) \leq |G_i|^{2\varepsilon}.$$

This shows that $\{F_i\}_i$ is also QNR, as desired.

3. The Ranks of Circulant Matrices

We first remark upon an equivalent reformulation for the rank of circulant matrices over Λ, which we call the *circulant rank*.

Definition 3.1. The *rank* of a function $F : G \times G \to \Lambda$ is the smallest integer $r = r(F)$ such that it is possible to write

$$F(x, y) = \sum_{1 \le i \le r} f_i(x) g_i(y)$$

where f_i, g_i are functions from G to Λ. The *circulant rank* of a function $f : G \to \Lambda$ is the rank of the function $F(x, y) = f(xy)$.

We can verify that the circulant rank of a function f is equal to the rank of its G-circulant matrix over Λ (by using Lemma 3.10).

In this section, we give an explicit formula for the rank of all G-circulant matrices for any finite group G and function $f : G \to \Lambda$, where Λ is a sufficiently nice field. This result was known to Diaconis [Dia90], but we give a different and more elementary proof.

3.1. **Rank Reduction.** In this section, we first give an explicit decomposition of any function $f : G \to \Lambda$ of the form $\sum f_i(x) g_i(y)$ for functions f_i and g_i. It is clear that we only need to obtain explicit decompositions of matrix coefficients.

Claim 3.2. The circulant rank of the function $\rho_{ij}(x)$ for any valid choice of indices i and j is less than or equal to $\deg \rho$.

Proof. Since ρ is a homomorphism, $\rho(xy) = \rho(x)\rho(y)$. This implies

$$\rho_{ij}(xy) = \sum_k \rho_{ik}(x) \rho_{kj}(y)$$

upon carrying out the matrix multiplication. The result follows from Definition 3.1. □

Once we have an explicit decomposition of any function $f : G \to \Lambda$, we would like to show that it is "optimal." In other words, if we take an arbitrary function and expand it into summands of the form $\rho_{ik}(x)\rho_{kj}(y)$ as above, the resulting decomposition should have the least number of summands possible after obvious factorizations. The following work makes this idea explicit.

Theorem 3.3 (Rank Reduction)**.** *Let A be a finite set of elements, and let m be a positive integer. Say we are given functions $f_i, g_i : A \to \Lambda$ for $1 \le i \le m$ such that*

$$\sum_{k=1}^{m-1} f_k(x) g_k(y) = f_m(x) g_m(y)$$

for all $x, y \in A$. Assume further that g_m is not the zero function. Then, f_m is identically equal to a linear combination of $f_1, f_2, \ldots, f_{m-1}$.

Proof. Let $|A| = N$, and index the members of A as $a_1, a_2, \ldots, a_N$. By the condition, there exists some index j with $1 \le j \le N$ such that $g_m(a_j) \neq 0$. Plugging in $y = a_j$ yields the result. □

Theorem 3.4 (General Rank Reduction)**.** *In the same terminology as Theorem 3.3, say we are given the equation*

$$\sum_{k=1}^{m-1} f_k(x) g_k(y) = \sum_{j=m}^{m+n} f_j(x) g_j(y)$$

for some nonnegative integer n. If we assume that the set $\{g_\ell \mid m \leq \ell \leq m+n\}$ is a linearly independent set of functions, then $f_m, f_{m+1}, \ldots, f_{m+n}$ are all linear combinations of $f_1, f_2, \ldots, f_{m-1}$.

Proof. We induct on n, with the base case $n = 0$ being Theorem 3.3. Henceforth, assume $n \geq 1$ and that the statement holds for $n = v$ for v a positive integer. To prove the statement for $n = v + 1$, we first rearrange the given equation as

$$\sum_{k=1}^{m-1} f_k(x)g_k(y) - \sum_{j=m}^{m+v} f_j(x)g_j(y) = f_{m+v+1}(x)g_{m+v+1}(y) \ (\star).$$

By the linear independence assumption, we know that g_{m+v+1} is not the zero function; hence, by Theorem 3.3, we know that f_{m+v+1} is a linear combination of the f_i for $1 \leq i \leq m+v$. To this end, write

$$f_{m+v+1}(x) = \sum_{i=1}^{m+v} c_i f_i(x).$$

Plugging this back into equation $(\star)$ and rearranging gives the equation

$$\sum_{k=1}^{m-1} f_k(x)(g_k(y) - c_k g_{m+v+1}(y)) = \sum_{j=m}^{m+v} f_j(x)(g_j(y) + c_j g_{j+v+1}(y)).$$

Now, since the functions $\{g_\ell \mid m \leq \ell \leq m+v+1\}$ are linearly independent by assumption, it is easy to see that the functions $\{g_j(y) + c_j g_{j+v+1}(y) \mid m \leq j \leq m+v\}$ is linearly independent as well. We can now apply the inductive hypothesis for $n = v$ to obtain the result. □

3.2. **Generalizations: Technical Setup.** In this section, we show that Theorem 3.4 applies when the f_i and g_i are taken to be matrix coefficients.

Claim 3.5. Consider a function $F : G \times G \to \Lambda$. Let $r = r(F)$, and write

$$F(x, y) = \sum_{i=1}^{r} f_i(x)g_i(y),$$

where $f_i, g_i : G \to \Lambda$. Then, the sets $\{f_i \mid 1 \leq i \leq r\}$ and $\{g_i \mid 1 \leq i \leq r\}$ are linearly independent as functions over G.

Proof. Assume without loss of generality that $f_1(x)$ is a linear combination of the other f_i. Then, write

$$f_1(x) = \sum_{i=2}^{r} c_i f_i(x).$$

We can then write $F(x, y)$ as

$$F(x, y) = \sum_{i=2}^{r} f_i(x)(g_i(y) + c_i g_1(y)),$$

so F actually has rank $r - 1$. This contradicts the minimality of r. □

Recall that over any field Λ, the matrix coefficients of the irreducible representations of G over Λ are linearly independent.

Claim 3.6. For any finite group G, the $|G|^4$ functions $\rho_{ij}(x)\pi_{rs}(y) : G \times G \to \Lambda$, where ρ and π are (not necessarily distinct) irreducible representations of G and $1 \leq i, j, r, s \leq N$, are linearly independent as functions on $G \times G$.

Proof. Index the irreducible representations of G as $\rho_1, \rho_2, \ldots, \rho_m$ for some positive integer m.

For a fixed representation ρ_k such that $\deg \rho_k = N$, label each ordered pair (i, j) where $1 \leq i, j \leq N$ with a unique number from 1 to N^2. Now, any linear combination of the matrix coefficients of ρ_k can be expressed as

$$c_{k,1,1}\rho_{k,1}(x)\rho_{k,1}(y) + c_{k,1,2}\rho_{k,1}(x)\rho_{k,2}(y) + \cdots + c_{k,N^2,N^2}\rho_{k,N^2}(x)\rho_{k,N^2}(y)$$

for constants $c_{k,1,1}, c_{k,1,2}, \ldots, c_{k,N^2,N^2} \in \Lambda$.

Now, assume that some linear combination of the matrix coefficients of G equals the zero function. In particular, assume that

$$\sum_{k=1}^{m} \left(\sum_{u=1}^{\deg \rho_k} \sum_{v=1}^{\deg \rho_k} [c_{k,u,v}\rho_{k,u}(x)\rho_{k,v}(y)] \right) = 0$$

as a function from $G \times G \to \Lambda$ for some choice of indices $c_{k,u,v}$. Grouping the terms by their matrix coefficient in y, we get an expression of the form

$$\sum_{v} \left(\sum_{k} \left[\sum_{u} c_{k,u,v}\rho_{k,u}(x) \right] \rho_{k,v}(y) \right) = 0,$$

where the bounds of the summations are the same as above. Since the $\rho_i(y)$ are linearly independent as functions over G, the $\sum_u c_{k,u,v}\rho_{k,u}(x)$ must equal zero for each u, v, and k. But since the $\rho_{k,u}(x)$ are also linearly independent functions over G, it follows that each of the $c_{k,u,v}$ must be zero, establishing the independence. □

3.3. **Circulant Ranks of Functions.** In this section, we synthesize our work and arrive at a new proof for the general formula for the ranks of circulant matrices. The approach we take is fundamentally different from, and uses less technology than, the one in Diaconis [Dia90]: more specifically, we only rely on the result that matrix coefficients of the irreducible representations of G over Λ form a basis for all functions $f : G \to \Lambda$ and do not need to appeal to the full results of Schur Orthogonality. Henceforth, assume that Λ is a field with characteristic coprime to $|G|$.

First, given a function $f : G \to \Lambda$, define the *parent function* $F : G \times G \to \Lambda$ of f as $F(x, y) = f(xy)$ for all $x, y \in G$. We define the *standard decomposition* of its parent function as follows:

Proposition 3.7 (Standard Decomposition)**.** *Given a function $f : G \to \Lambda$, the standard decomposition of its parent function F is the unique decomposition of $F(x, y)$ into the zero-one linear combination of the $|G|^4$ functions $\rho_{ij}(x)\pi_{rs}(y)$, where ρ and π are two (not necessarily distinct) irreducible representations of G.*

The standard decomposition of the parent function of a function $f : G \to \Lambda$ is given by "expanding" the matrix coefficients of f as per Claim 3.2 and combining like terms.

We now introduce the concept of the *standard decomposition (SD) matrix* of a parent function F of f. First, say there are m irreducible representations of G, and order the irreducible representations ρ of G as $\rho_1, \rho_2, \ldots, \rho_m$ such that

$$\deg \rho_1 \leq \deg \rho_2 \leq \cdots \leq \deg \rho_m.$$

We define the SD matrix to be a $|G|^2 \times |G|^2$ matrix with each row and column indexed by a unique irreducible representation of G in the following way:

- Assign the topmost $1 = (\deg \rho_1)^2$ row and the leftmost $(\deg \rho_1)^2$ column to ρ_1.
- Assign the next topmost $(\deg \rho_2)^2$ rows and the leftmost $(\deg \rho_2)^2$ columns to ρ_2, and so on.

Now, we will complete the indexing of the rows and columns by assigning them each a matrix coefficient. For a particular irreducible representation ρ of G with $\deg \rho = N$, we assign matrix coefficients to the rows and columns corresponding to that representation in the following way:

- The rows will be indexed by $\rho_{ij}(x)$, with the first N rows corresponding to $\rho_{11}(x), \rho_{12}(x), \ldots, \rho_{1N}(x)$. The second N rows correspond to $\rho_{21}(x), \rho_{22}(x), \ldots, \rho_{2N}(x)$, and so on.
- The columns will be indexed by $\rho_{rs}(y)$. The leftmost N columns correspond to $\rho_{11}(y), \rho_{21}(y), \ldots, \rho_{N1}(y)$. The second N columns correspond to $\rho_{12}(y), \rho_{22}(y), \ldots, \rho_{N2}(y)$, and so on.

With this setup, we can define the standard decomposition matrix.

Definition 3.8 (Standard Decomposition Matrix)**.** The entry of the standard decomposition matrix in the column corresponding to $\pi_{rs}(y)$ and the row corresponding to $\rho_{ij}(x)$ is the coefficient of $\rho_{ij}(x)\pi_{rs}(y)$ in the standard decomposition of F.

Theorem 3.9 (Rank of SD Matrix Equals Circulant Rank of Function)**.** *Fix a finite group G, and assume that Λ is a field with characteristic coprime to $|G|$. For a function $f : G \to \Lambda$ which factors through a representation ρ, its circulant rank is equal to the rank of the standard decomposition matrix of the parent function of f.*

To prove Theorem 3.9, we need the following lemma.

Lemma 3.10 (Rank and Matrix Products)**.** *For any positive integers n and $r \leq n$, an $n \times n$ matrix has rank r if and only if it can be written as the product of an $n \times r$ and an $r \times n$ matrix, both of rank r.*

Proof of Lemma. First, if a matrix can be written as the product of an $n \times r$ and an $r \times n$ matrix, it clearly has rank at most r since every column of the resulting $n \times n$ matrix is a linear combination of the r column vectors in the $n \times r$ matrix. (Alternatively, we can use the fact that the rank of the product of two matrices is at most the minimum rank of the individual matrices.) The fact that the matrix has rank exactly r follows from the $r \times n$ matrix having rank r, implying that it has r independent columns.

Now, if an $n \times n$ matrix has rank r, taken the r pivot columns of the matrix and use that as the $n \times r$ matrix. Then, fill in the corresponding linear combinations of the other columns in terms of the pivot columns for the $r \times n$ matrix. For the same reasons as in the previous case, these matrices both have rank r. □

Now, we will prove Theorem 3.9.

Proof. Say the rank of the standard decomposition matrix of the parent function of f equals R. By Lemma 3.10, it is then possible to write this matrix as the product of an $|G|^2 \times R$ and an $R \times |G|^2$ matrix. Now, we claim that it is possible to express F as the sum of R products $f_i(x)g_i(y)$.

This can be done in the following way: the rows of the $|G|^2 \times R$ matrix are indexed by the matrix coefficients of the irreducible representations of G. Now, for all $1 \leq i \leq R$, take $f_i(x)$ as the function defined by taking the linear combination of these matrix coefficients, with each coefficient scaled by its corresponding entry of the ith column in that matrix. Similarly, $g_i(y)$ for all $1 \leq i \leq R$ is the function formed by taking the analogous linear combination of the ith *row* with respect to the basis indexed by the columns in the $R \times |G|^2$ matrix. By the minimality of matrix rank, we can ensure that R is the minimal number of summands required.

Conversely, say the rank of the parent function of some $f : G \to \Lambda$ was exactly R. Decompose f into R products $f_i(x)g_i(y)$, and apply the inverse of the corresponding process above to construct an $|G|^2 \times R$ and an $R \times |G|^2$ matrix based on the coefficients of the functions. It is straightforward to show that these multiply to form the standard decomposition matrix, which has rank R by Lemma 3.10. □

The reason we stipulated the indexing scheme for the standard decomposition matrix is because of the following structure: when given a function

$$f(x) = \sum_{\rho \in I} \left(\sum_{1 \leq i,j \leq \deg \rho} c_{\rho,i,j} \rho_{ij}(x) \right)$$

where I is the set of irreducible representations of G and where $c_{\rho,i,j} \in \Lambda$, the standard decomposition matrix of the parent function F of f is a block diagonal matrix with $|I|$ blocks, each indexed by an irreducible representation of the group. Now, we claim that each block is of the form

$$\begin{bmatrix} c_{\rho,1,1}I_N & c_{\rho,1,2}I_N & \cdots & c_{\rho,1,N}I_N \\ c_{\rho,2,1}I_N & c_{\rho,2,2}I_N & \cdots & c_{\rho,2,N}I_N \\ \vdots & \vdots & \ddots & \vdots \\ c_{\rho,N,1}I_N & c_{\rho,N,2}I_N & \cdots & c_{\rho,N,N}I_N \end{bmatrix}$$

for some fixed $\rho \in I$, where $N = \deg \rho$ and I_N denotes the $N \times N$ identity matrix.

To see why this is, note that we can expand $\rho_{ij}(xy)$ as

$$\sum_k c_{\rho,i,j} \rho_{i,k}(x) \rho_{k,j}(y).$$

The key is that the column index of the first term is equal to the row index of the second; because of our indexing scheme of the standard decomposition matrix, we know that with respect to the diagonal block in the matrix corresponding to ρ,

$\rho_{i,k}(x)$ corresponds with row $iN+k$ and $\rho_{k,j}(y)$ corresponds with column $jN+k$. The entry with respect to this row and column – where we index our row and column numbers with respect to the subblock corresponding to ρ in the SD matrix – equals $c_{\rho,i,j}$.

Since

$$iN+k \equiv jN+k \pmod{N},$$

the $N^2 \times N^2$ subblock corresponding to ρ is an $N \times N$ block matrix. Then, each appearance of a different $c_{\rho,i,j}$ in the expansion of $\rho(xy)$ appears as a copy of $c_{\rho,i,j}I_n$ in the (i,j)th position of this block matrix, as desired.

In particular, we have the following result.

Theorem 3.11 (Ranks of Circulant Matrices)**.** *For any group G, field Λ with characteristic coprime to $|G|$, and function $f : G \to \Lambda$, express f in the form*

$$f(x) = \sum_{\rho \in I} \left(\sum_{1 \le i,j \le \deg \rho} c_{\rho,i,j}\rho_{ij}(x) \right)$$

where I is the set of isomorphism classes of irreducible representations of G and where $c_{\rho,i,j} \in \Lambda$. Then, the rank of the G-circulant matrix defined by $f(xy)_{ij}$ is equal to

$$\sum_{\rho \in I} \left[(\deg \rho)\ \mathit{rank} \left(\begin{bmatrix} c_{\rho,1,1} & c_{\rho,1,2} & \cdots & c_{\rho,1,N} \\ c_{\rho,2,1} & c_{\rho,2,2} & \cdots & c_{\rho,2,N} \\ \vdots & \vdots & \ddots & \vdots \\ c_{\rho,N,1} & c_{\rho,N,2} & \cdots & c_{\rho,N,N} \end{bmatrix} \right) \right].$$

Corollary 3.12. *Every matrix coefficient has circulant rank exactly N, the degree of its corresponding representation, over all fields Λ with characteristic coprime to $|G|$.*

Corollary 3.13. *Over any group G, almost all circulant matrices are invertible if Λ is a sufficiently large finite field with characteristic coprime to $|G|$ or if char $\Lambda = 0$.*

4. Applications of Circulant Rank

In this section, we apply Theorem 3.11 in the context of classical circulant matrices. More specifically, we offer a novel way to study the invertibility of circulants through vanishing sums of roots of unity, which allows us to more quickly deduce existing results, further generalize them, and resolve an open problem in graph theory.

4.1. Background.

Notice in the case where $G = \mathbb{Z}/n\mathbb{Z}$, all G-circulant matrices take the form

$$\begin{bmatrix} a_1 & a_2 & a_3 & \cdots & a_n \\ a_n & a_1 & a_2 & \cdots & a_{n-1} \\ a_{n-1} & a_n & a_1 & \cdots & a_{n-2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_2 & a_3 & a_4 & \cdots & a_1 \end{bmatrix},$$

where each row is the cyclically permuted version of the one above. In what follows, we will call these *circulant matrices*.

Circulant matrices have many applications in signal processing and appear prominently in algorithms such the Discrete Fourier Transform; as such, they have been extensively studied [Che21] [Ter99]. Moreover, the question of the invertibility of circulant matrices has previously been of interest.

In 2021, Chen proved the following theorem:

Theorem 4.1 (Chen 2021, Theorem 1.7–1.9)**.** *For any odd positive integer n, $\mathbb{Z}/n\mathbb{Z}$-circulant matrices with $\frac{n-1}{2}$ ones and $\frac{n+1}{2}$ zeros in each row are always nonsingular over $\mathbb{C}$ if and only if $n = p^k$ or pq for distinct primes p, q and positive integers k.*

We will exhibit a new proof of this theorem in this section. To do so, we will need the following results, due respectively to Lam–Leung [LL00] and Sivek [Siv10]:

Theorem 4.2 (Lam–Leung 2000, Corollary 3.4)**.** *Let $m = p^a q^b$, where p and q are primes. Then, up to rotation, the only minimal vanishing sums of mth roots of unity are $1 + \zeta_p + \cdots + \zeta_p^{p-1} = 0$ and $1 + \zeta_q + \cdots + \zeta_q^{q-1} = 0$.*

Theorem 4.3 (Sivek 2010, Theorem 2)**.** *For any positive integer n and integer $0 \leq k \leq n$, there exist k distinct nth roots of unity with sum zero if and only if both k and $n - k$ are expressible as linear combinations of prime factors of n with nonnegative coefficients.*

4.2. **Invertibility of Circulant Matrices.** Consider Theorem 3.11 in the case where G is abelian. In this case, all the degrees of the irreducible representations of G are equal to 1. In particular, we have the following corollary:

Corollary 4.4. *Let Λ be a field with characteristic coprime to $|G|$. If G is abelian, the circulant rank of any function $f : G \to \Lambda$ is equal to the number of distinct matrix coefficients that appear in the expansion of f.*

Since the group $\mathbb{Z}/n\mathbb{Z}$ is abelian, Corollary 4.4 applies in the context of classical circulant matrices. Moreover, we know the precise matrix coefficients of the cyclic groups, allowing us to explicitly compute the rank and determine the invertibility of circulant matrices. Specifically, notice first that the matrix coefficients of $\mathbb{Z}/n\mathbb{Z}$ are the homomorphisms sending a fixed generator of $\mathbb{Z}/n\mathbb{Z}$ to a nth root of unity. Denote these n matrix coefficients as $\rho_1, \rho_2, \ldots, \rho_n$ in some order.

Consider a function $f : \mathbb{Z}/n\mathbb{Z} \to \mathbb{C}$ and the circulant matrix with respect to f. Let the $\mathbb{C}$-linear combination of f in terms of the ρ_i be written as

$$f(x) := \sum_{i=1}^{n} a_i \rho_i(x)$$

for all $x \in \mathbb{Z}/n\mathbb{Z}$, where $a_i \in \mathbb{C}$. Corollary 4.4 implies that the circulant matrix of f is not invertible if and only if at least one of the a_i is equal to zero.

Notice that the equation above must hold for all $x \in \mathbb{Z}/n\mathbb{Z}$. Since we know the values of $f(x)$ and $\rho_i(x)$ for all such x, letting x vary gives a system of n equations we can use to solve for the a_i. Combining the resulting system in matrix form, we

get the equation

$$\begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \cdots & \omega^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{n-1} & \omega^{2n-2} & \cdots & \omega^{(n-1)^2} \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix} = \vec{v},$$

where $\omega = e^{2\pi i/n}$ and $\vec{v}$ is the n-dimensional vector such that the ith coordinate of $\vec{v}$ is $f(i)$ (when we zero-index the coordinates). Alternatively, $\vec{v}$ is the transpose of the topmost row in the circulant matrix. To isolate the a_i, we can invert the discrete Fourier transform matrix on the left; doing so gives

$$\begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix} = \frac{1}{n} \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega^{-1} & \omega^{-2} & \cdots & \omega^{-(n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{-(n-1)} & \omega^{-(2n-2)} & \cdots & \omega^{-(n-1)^2} \end{bmatrix} \vec{v}.$$

Thus, we have the following result:

Lemma 4.5. *An $n \times n$ circulant matrix with first row $\langle c_1, c_2, \ldots, c_n \rangle$ is invertible if and only if the vector*

$$\begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix} = \frac{1}{n} \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega^{-1} & \omega^{-2} & \cdots & \omega^{-(n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{-(n-1)} & \omega^{-(2n-2)} & \cdots & \omega^{-(n-1)^2} \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{bmatrix}$$

has all coordinates nonzero.

Lemma 4.5 provides a tractable criterion from which to determine whether or not a circulant matrix is invertible. In particular, we will use it to provide a new proof of Theorem 4.1.

Proof. For a circulant matrix to be nonsingular, it must have full rank. By Corollary 4.4, this means for all $n = p^k$ and pq, any function $f : \mathbb{Z}/n\mathbb{Z} \to \mathbb{C}$ which outputs the value one $\frac{n-1}{2}$ times and zero the other $\frac{n+1}{2}$ times must have all coefficients nonzero in its $\mathbb{C}$-linear combination of matrix coefficients.

In the notation of Lemma 4.5, our goal is to show that for $n = p^k$ or pq, none of $a_1, a_2, \ldots, a_n$ are zero. In fact, we need to show that for every row, it is impossible to pick $\frac{n-1}{2}$ entries that sum to zero.

We will now break the problem up into two cases. The first case is where $n = p^k$, where each row is either a complete collection of the p^kth roots of unity or a multiset of p^ℓth roots of unity for some $\ell < k$. Theorem 4.2 implies that the only minimal vanishing sums of roots of unity in any of these cases is the sum of the pth roots of unity, which means that any vanishing sums of p^ith roots of unity must have a number of summands which is a multiple of p. However, $\frac{n-1}{2} = \frac{p^k-1}{2}$ is not a multiple of p, so any sum of $\frac{n-1}{2}$ p^ith roots of unity does not vanish and hence the corresponding circulant matrix is nonsingular.

The other case is when $n = pq$, in which case we know that the only minimal vanishing sums of nth roots of unity are that of the pth and qth roots of unity. Like before, we can characterize each row of the matrix: it is either a multiset of pth or qth roots of unity, or a complete set of pqth roots of unity. Similarly as before, any sum of $\frac{n-1}{2}$ pth or qth roots of unity will not vanish since $\frac{n-1}{2}$ is not a multiple of p or of q.

For the rows which contain pqth roots of unity, it follows from a similar argument as before that if some set of $\frac{n-1}{2} = \frac{pq-1}{2}$ roots of unity has sum zero, then $\frac{n-1}{2}$ must be a linear combination of p and q. In particular, since $\frac{n-1}{2}$ is neither a multiple of p or of q, it must be expressible as $rp + sq$ for some $r, s \geq 1$. In other words, we want to partition the set of pqth roots of unity into r sets of p evenly-spaced roots and s sets of q evenly-spaced roots, all of which are disjoint. However, this is impossible: by the Chinese Remainder Theorem, *any* set of p evenly-spaced roots will have at least one root that overlaps with *any* set of q evenly-spaced roots. Thus, it is impossible to have a vanishing sum of $\frac{n-1}{2}$ pqth roots of unity, and hence all the corresponding circulant matrices are nonsingular.

Now, it remains to show that all other odd numbers n do produce singular matrices. It is possible to write any n in this form as $n = pqr$, where p and q are the two smallest distinct primes dividing n and where r is an odd integer. We will show that there exist $\frac{pqr-1}{2}$ pqrth roots of unity which have vanishing sum.

By Theorem 4.3, it suffices to show that both $\frac{pqr-1}{2}$ and $\frac{pqr+1}{2}$ can be written as a linear combination of p, q, and the prime divisors of r. In fact, notice that since

$$\frac{pqr+1}{2}, \frac{pqr-1}{2} > pq > pq - p - q,$$

we can write $\frac{pqr-1}{2}$ as a linear combination of p and q alone, showing that there exist singular matrices in this case. □

Moreover, the technique we used in our reproof of Theorem 4.1 allows us to substantially generalize the result and to provide an arithmetic criterion on when an arbitrary $n \times n$ zero-one circulant is invertible, regardless of the number of ones in the first row.

Theorem 4.6. *For any positive integer n and integer $0 \leq k \leq n$, an $n \times n$ circulant matrix with k ones and $n - k$ zeros in each row is always nonsingular if and only if at least one of k and $n - k$ cannot be expressed as a $\mathbb{Z}_{\geq 0}$-linear combination of the distinct prime divisors of n.*

Proof. First, notice that by our proof of Theorem 4.1, the result is equivalent to proving that if k and $n - k$ cannot be expressed as a $\mathbb{Z}_{\geq 0}$-linear combination of the distinct prime divisors of n, then it is impossible to pick k roots of unity out of any fixed row of the $n \times n$ Fourier matrix with sum zero.

We will begin by showing that if both k and $n - k$ are such linear combinations, then it is possible to construct a singular matrix. This follows from Theorem 4.3.

Next, assume that we have a singular zero-one circulant matrix with k ones in each row; we will prove that both k and $n - k$ can be expressible in the above form. From the matrix singularity, we know that there exists a multiset of k (not necessarily distinct) roots of unity with sum equal to zero. In fact, this set of roots

of unity is going to be a multiset of ℓth roots of unity for some $\ell \mid n$ and $\ell > 1$. In particular, Theorem 4.2 tells us that k must be a $\mathbb{Z}_{\geq 0}$-linear combination of the prime factors of ℓ, and thus of n. To show that $n-k$ must also be such a combination, we note that the sum of the roots of unity in the entire row is also zero; hence, the matrix by replacing 1s with 0s (and vice versa) will also have that row sum equal to zero, and $n-k$ is thus a $\mathbb{Z}_{\geq 0}$ linear combination of the prime factors of ℓ (and n) by applying the same result.

Thus, if at least one of k and $n-k$ are not linear combinations of the prime divisors of n, we know that no row can have vanishing sum (because every row of vanishing sum with k summands will necessarily have both of k and $n-k$ be such a combination). It follows that the matrix cannot be singular. □

Theorem 4.6 gives another proof for the nonsingularity of the $N = p^k, pq$ cases in Theorem 4.1. For the $N = p^k$ case, we need to show that $\frac{p^k-1}{2}$ is not a multiple of p, which is true. For $N = pq$, assume that both $\frac{pq-1}{2}$ and $\frac{pq+1}{2}$ can be written as $\mathbb{Z}_{\geq 0}$-combinations of p and q. Since neither of them are multiples of p or of q, they must both be written of the form $rp + sq$ for some $r, s \geq 1$. Now, if we add them together, we obtain the equation

$$pq = r'p + s'q$$

for some $r', s' \geq 2$, and subtracting $p + q$ yields

$$pq - p - q = (r' - 1)\, p + (s' - 1)\, q,$$

a $\mathbb{Z}^+$-linear combination of p and q. This directly contradicts the Chicken McNugget Theorem, so all matrices in this case are nonsingular.

Next, we use Lemma 4.5 to answer an open problem posed by Lal–Reddy in [LR11], who ask for a set of necessary and sufficient conditions for (r, s, t) circulant digraphs to be singular.

An (r, s, t) circulant digraph, where $r, s, t \in \mathbb{Z}_{\geq 0}$, is a directed graph whose adjacency matrix an $n \times n$ circulant with first row of the following form:

$$\langle \underbrace{1, 1, \ldots, 1}_{r}, \underbrace{0, 0, \ldots, 0}_{t}, \underbrace{1, 1, \ldots, 1}_{s}, \underbrace{0, 0, \ldots, 0}_{n-(r+t+s)} \rangle.$$

Thus, it suffices to determine when these matrices are singular.

Theorem 4.7. *An (r, s, t) circulant digraph is singular if and only if at least one of the following four conditions hold:*

- $r = t = 0$,
- $1 < \gcd(r, s) \mid n$,
- $\gcd(r - s, n)$ *is even and*

$$r + t = \frac{(2k+1)n}{2c}$$

  for some nonnegative integer k and positive integer c, or
- $1 < t \mid \gcd(r + s, n)$.

Proof. By Lemma 4.5, we know that this is nonsingular if and only if

$$\omega^0 + \omega^1 + \cdots + \omega^{r-1} + \omega^{r+t} + \omega^{r+t+1} + \cdots + \omega^{r+t+s-1} \neq 0$$

for all $\omega \in \mathbb{C}$ such that $\omega^n = 1$. First, if $\omega = 1$, this implies that $r+s > 0$. Otherwise, $r \neq 1$ and the above expression can be written in the equivalent form

$$\frac{\omega^r - 1}{\omega - 1} + \left(\omega^{r+t}\right)\left(\frac{\omega^s - 1}{\omega - 1}\right) \neq 0 \iff (\omega^r - 1) + \left(\omega^{r+t}\right)(\omega^s - 1) \neq 0.$$

Now, we will investigate what happens in the case where one of these expressions equals zero, i.e. the (r, s, t) digraph is singular. In this case, we know that for some nth root of unity ω, we must have

$$(\omega^r - 1) + \left(\omega^{r+t}\right)(\omega^s - 1) = 0 \iff \omega^{r+t} = -\frac{\omega^r - 1}{\omega^s - 1}.$$

Now, we will prove that in order for

$$\omega^{r+t} = -\frac{\omega^r - 1}{\omega^s - 1},$$

then ω^r and ω^s must be equal or conjugate. Notice that since $|\omega^{r+t}| = 1$, then we must have

$$\left|-\frac{\omega^r - 1}{\omega^s - 1}\right| = 1 \implies |\omega^r - 1| = |\omega^s - 1|.$$

Since

$$|\omega^r - 1| = (\omega^r - 1)\left(\overline{\omega^r} - 1\right) = (\omega^r - 1)\left(\frac{1}{\omega^r} - 1\right) = 2 - \omega^r - \omega^{-r} = 2 - 2\mathrm{Re}(\omega^r),$$

it becomes clear that

$$|\omega^r - 1| = |\omega^s - 1| \iff \mathrm{Re}(\omega^r) = \mathrm{Re}(\omega^s),$$

implying that ω^r and ω^s are either equal or conjugates.

Now, we have two more cases to consider. If ω^r and ω^s are conjugates, then we claim that in order for

$$(\omega^r - 1) + \left(\omega^{r+t}\right)(\omega^s - 1) = 0,$$

we must have either $\omega^r = \omega^s = 1$ or $\omega^t = 1$. To see why, plugging in $\omega^s = \frac{1}{\omega^r}$ gives

$$(\omega^r - 1) + \left(\omega^t\right)(1 - \omega^r) = 0 \iff \left(1 - \omega^t\right)(\omega^r - 1) = 0,$$

implying the conclusion.

Otherwise, if $\omega^r = \omega^s$, then

$$(\omega^r - 1) + \left(\omega^{r+t}\right)(\omega^s - 1) = 0 \iff (\omega^r - 1)\left(1 + \omega^{r+t}\right),$$

so we need either $\omega^r = \omega^s = 1$ or $\omega^{r+t} = -1$.

Collating our results, we have the following criteria for when the (r, s, t) digraph can be singular:

- $r = t = 0$, in which case the entire matrix is the zero matrix,
- There exists an nth root of unity $\omega \neq 1$ (not necessarily primitive) such that $\omega^r = \omega^s = 1$,

- There exists an nth root of unity $\omega \neq 1$ such that $\omega^r = \omega^s$ and such that $\omega^{r+t} = -1$, or
- There exists an nth root of unity $\omega \neq 1$ such that ω^r and ω^s are complex conjugates and $\omega^t = 1$.

We will now determine the arithmetic properties r, s, and t must satisfy in terms of n in order for these criteria to hold.

For the second criterion to hold, we would like an nth root of unity $\omega \neq 1$ such that $\omega^r = \omega^s = 1$. This occurs if and only if $\gcd(r, s) > 1$ and $\gcd(r, s) \mid n$ (where we define $\gcd(m, 0) = m$ for any positive integer m).

For the third criterion to hold, we need $\gcd(r - s, n) \neq 1$ in order for $\omega^r = \omega^s$. Furthermore, we need n even and $2 \mid \gcd(r - s, n)$, as we are given $\omega^{r+t} = -1$. More precisely, we need

$$r + t = \frac{(2k + 1)n}{2c}$$

for some nonnegative integer k and positive integer $1 \leq c \leq n - 1$ such that $c \mid (2k + 1)n$. In particular, one special case of this worth mentioning is when $\omega = -1$, or equivalently $c = \frac{n}{2}$; in this case, it is necessary and sufficient to have $r \equiv s \not\equiv t \pmod 2$.

For the last criterion, ω^r and ω^s being complex conjugates tells us that we must have $\gcd(r + s, n) > 1$. Furthermore, since $\omega^t = 1$ for this value of ω, we need $t \mid \gcd(r + s, n)$.

Aggregating these four yields the criteria in the statement. □

We end this section by noting that Lemma 4.5 also allows us to exhibit a lower bound on the number of singular circulant matrices of a given dimension. In particular, this implies a lower bound on the probability that a randomly chosen $n \times n$ zero-one circulant matrix is singular.

Proposition 4.8. *For sufficiently large positive integers n, the probability that a randomly-chosen zero-one $\mathbb{Z}/n\mathbb{Z}$-circulant matrix is singular is, asymptotically, at least* $\frac{1}{\sqrt{p}} \left(\frac{2p}{\pi n} \right)^{\frac{p-1}{2}}$, *where p is the smallest prime dividing n.*

Proof. Consider the $n \times n$ discrete Fourier matrix. If p is the smallest prime dividing n, notice that the inverse of this matrix has a row consisting of $\frac{n}{p}$ copies of the list

$$\langle 1, \zeta^{-1}, \zeta^{-2}, \ldots, \zeta^{-p+1} \rangle,$$

where $\zeta = e^{2\pi i/p}$ is a primitive pth root of unity. Call this row vector $\vec{u}$. Now, a zero-one circulant is singular if the n-dimensional vector $\vec{v}$ consisting of the entries from the first row satisfies $\vec{v} \cdot \vec{u} = \vec{0}$.

Now, we will exhibit a lower bound on the number of such vectors $\vec{v}$. To do this, we can see that one way to create such a vector $\vec{v}$ is by letting the sum $\vec{v} \cdot \vec{u}$ consist only of complete sets of roots of unity modulo p. Say we have k such complete sets for some positive integer k; this implies that each power ζ^ℓ of ζ must appear exactly k times in the dot product expansion of $\vec{v} \cdot \vec{u}$.

To count this, we see that there are $\binom{n/p}{k}$ ways to choose the k copies of ζ^ℓ for any power ℓ; since there are p such powers, a lower bound for the number of

possible vectors $\vec{v}$ is

$$\sum_{k=0}^{n/p} \binom{n/p}{k}^p.$$

It is well known from [PS+72] that this is asymptotically approximated by the expression

$$\frac{2^n}{\sqrt{p}} \left(\frac{2p}{\pi n}\right)^{\frac{p-1}{2}},$$

and the result follows upon dividing by the 2^n total $\mathbb{Z}/n\mathbb{Z}$ zero-one circulants. □

5. The Minrank and Faithful Orthogonality Dimension of Graphs

5.1. **Background.** One avenue in the literature for attempting problems regarding matrix rigidity has been to recast the rigidity question through the lens of graph theory.

More concretely, fix a finite simple graph G with n vertices. In the notation of Golovnev–Haviv [GH20], we that an $n \times n$ matrix M over a field $\mathbb{F}$ *represents* G if its diagonal entries are all nonzero and $M_{ij} = 0$ for only the $i, j \in V$ such that $(i, j) \notin E(G)$. Define the *minrank* of G as

$$\text{minrk}(G) = \{\min(\text{rank}(M)) \mid M \text{ represents } G\}.$$

Golovnev–Haviv [GH20] exhibited a connection between the minrank and what first appears to be an unrelated concept: orthogonal representations of graphs. A t-dimensional *orthogonal representation* of a graph G with respect to a field $\mathbb{F}$ as an assignment of a nonzero vector $v \in \mathbb{F}^t$ to each vertex of G, such that any two vectors corresponding to adjacent vertices in G are orthogonal. If an orthogonal representation has two vectors orthogonal *if and only if* their corresponding vertices are adjacent, we call the orthogonal representation *faithful*. We call the minimal dimension t for which such an orthogonal (resp. faithful) representation exists the *orthogonality dimension* (resp. *faithful orthogonality dimension*) of G.

In this section, we provide a novel technique and lower bound for the faithful orthogonality dimension of a graph through the intermediary use of minrank. To do so, we will need a result from linear algebra, due to Golovnev–Regev–Weinstein [GRW18].

Lemma 5.1 (Golovnev–Regev–Weinstein 2018, Lemma 3)**.** *For an $n \times n$ matrix M, define $s(M)$ to be the number of nonzero entries in M. If M has nonzero entries on the main diagonal,*

$$\text{rank}(M) \geq \frac{n^2}{4s(M)}.$$

In Section 5.3, we will also use the following two results, the first of which is due to Mackay [Mac03] and the second of which is due to Rödl [LSS89].

Proposition 5.2 (Mackay 2003, Equation 1.17)**.** *For integers N and r sufficiently large,*

$$\log \binom{N}{r} \simeq (N - r) \log \frac{N}{N - r} + r \log \frac{N}{r}.$$

Theorem 5.3 (Rödl 1987)**.** *For a finite, simple, undirected graph G with n vertices, its faithful orthogonality dimension is at most $2(n-D-1)$, where D is the maximum degree among all vertices in G.*

5.2. **Minrank and Faithful Orthogonality Dimension.** We will correct a typo in Golovnev–Haviv [GH20], who posit that the orthogonality dimension is an upper bound for the minrank. The authors should have been tracking the *faithful orthogonality dimension* (notated by ξ_F), which provides the requisite nondegeneracy parameters. In particular:

Lemma 5.4. *For a finite, simple, undirected graph G, $minrk(G) \leq \xi_F(\overline{G})$.*

Proof. We follow the idea and notation of the third footnote of Golovnev–Haviv [GH20]. Consider a t-dimensional faithful orthogonal representation of an n-vertex graph G over some field $\mathbb{F}$. Now, consider B, a $n \times t$ matrix over $\mathbb{F}$ whose rows are the vectors corresponding to the faithful orthogonal representation of G. We claim that the $n \times n$ matrix $M := B \cdot B^T$ represents $\overline{G}$.

To see why, consider the entry M_{uw}, the entry corresponding to the row of vertex u and the column of vertex w. This is the dot product of the vectors in row u of B and column w of B^T. In particular, this is zero if and only if u and w are *not* adjacent in $\overline{G}$. This implies that they must be adjacent in G, as desired. □

Now, for a graph G with n vertices, notice that the number of nonzero entries in any matrix which represents G is equal to $n + 2|E(G)|$. We obtain the following result:

Corollary 5.5. *For a finite, undirected graph G with n vertices, we have*

$$minrk(G) \geq \frac{n^2}{4(n + 2|E(G)|)}.$$

Combining the results of Lemma 5.4 and Corollary 5.5, we get the following result:

Theorem 5.6. *For any finite, undirected graph G with n vertices, we have*

$$\frac{n^2}{4(n + 2|E(G)|)} \leq minrk(G) \leq \xi_F(\overline{G}).$$

The bound consisting of the first and third quantities in Theorem 5.6 is novel, and it allows for easy computation of a lower bound for the faithful orthogonality dimension. Equivalently, we can rewrite the inequality as

$$\frac{n^2}{4(n^2 - 2|E(G)|)} \leq \xi_F(G).$$

5.3. **Faithful Orthogonality Dimension of Kneser Graphs.** In this section, we will use Theorem 5.6 to prove the following theorem about the value of the ξ_F of Kneser graphs. A *Kneser graph* $K(n, k)$ is a graph with vertices indexed by the k-element subsets of $\{1, 2, \ldots, n\}$, and where two sets are connected if and only if they are disjoint. The next result shows that for certain families of these graphs, we can bound the faithful orthogonality dimension up to a constant factor.

Theorem 5.7. *There exist two fixed constants c_1 and c_2 such that the following statement holds: for every sufficiently small positive constant $\varepsilon \ll 1$, there exists a positive integer n such that*

$$c_1 n^{\frac{\varepsilon n}{2}} \leq \xi_F(K(n, \varepsilon n)) \leq c_2 n^{\frac{\varepsilon n}{2}}.$$

Proof. First, consider the graph-theoretic properties of $K(n, \varepsilon n)$. This is a regular graph with $\binom{n}{\varepsilon n}$ vertices and where each vertex has degree $\binom{n-\varepsilon n}{\varepsilon n}$. In particular, Theorem 5.4 implies that

$$\frac{\binom{n}{\varepsilon n}^2}{4\left(\binom{n}{\varepsilon n}^2 - \binom{n}{\varepsilon n}\binom{n-\varepsilon n}{\varepsilon n}\right)} \leq \xi_F(K(n, \varepsilon n)) \leq 2\left(\binom{n}{\varepsilon n} - \binom{n-\varepsilon n}{\varepsilon n} - 1\right).$$

Simplifying, we obtain

$$\frac{\binom{n}{\varepsilon n}}{4\left(\binom{n}{\varepsilon n} - \binom{n-\varepsilon n}{\varepsilon n}\right)} \leq \xi_F(K(n, \varepsilon n)) \leq 2\left(\binom{n}{\varepsilon n} - \binom{n-\varepsilon n}{\varepsilon n} - 1\right).$$

Thus, it suffices to show that for sufficiently small ε, there exists an integer n such that $\binom{n}{\varepsilon n} - \binom{n-\varepsilon n}{\varepsilon n} \sim n^{\frac{\varepsilon n}{2}}$.

To do this, we will use Proposition 5.2. Doing so, we obtain

$$\binom{n}{\varepsilon n} \simeq e^{n(-((1-\varepsilon)\ln(1-\varepsilon)+\varepsilon\ln\varepsilon))}$$

and

$$\binom{n-\varepsilon n}{\varepsilon n} \simeq e^{(n-\varepsilon n)\left(-\left(\left(1-\frac{\varepsilon}{1-\varepsilon}\right)\ln\left(1-\frac{\varepsilon}{1-\varepsilon}\right)+\frac{\varepsilon}{1-\varepsilon}\ln\left(\frac{\varepsilon}{1-\varepsilon}\right)\right)\right)}.$$

The result is equivalent to showing that for sufficiently small ε, the equation

$$e^{n(-((1-\varepsilon)\ln(1-\varepsilon)+\varepsilon\ln\varepsilon))} - e^{(n-\varepsilon n)\left(-\left(\left(1-\frac{\varepsilon}{1-\varepsilon}\right)\ln\left(1-\frac{\varepsilon}{1-\varepsilon}\right)+\frac{\varepsilon}{1-\varepsilon}\ln\left(\frac{\varepsilon}{1-\varepsilon}\right)\right)\right)} = n^{\frac{\varepsilon n}{2}}$$

has arbitrarily large solutions.

To do this, we will first show that when $n = \frac{2}{\varepsilon}$, the left-hand side is at least the right. Indeed, we would like to prove the inequality

$$e^{\frac{2}{\varepsilon}(-((1-\varepsilon)\ln(1-\varepsilon)+\varepsilon\ln\varepsilon))} - e^{\left(\frac{2}{\varepsilon}-2\right)\left(-\left(\left(1-\frac{\varepsilon}{1-\varepsilon}\right)\ln\left(1-\frac{\varepsilon}{1-\varepsilon}\right)+\frac{\varepsilon}{1-\varepsilon}\ln\left(\frac{\varepsilon}{1-\varepsilon}\right)\right)\right)} \geq \frac{2}{\varepsilon}.$$

This can be checked for sufficently small values of ε.

Next, to show that the original equation has arbitrarily large solutions, notice that as $n \to \infty$, we have

$$e^{n(-((1-\varepsilon)\ln(1-\varepsilon)+\varepsilon\ln\varepsilon))} - e^{(n-\varepsilon n)\left(-\left(\left(1-\frac{\varepsilon}{1-\varepsilon}\right)\ln\left(1-\frac{\varepsilon}{1-\varepsilon}\right)+\frac{\varepsilon}{1-\varepsilon}\ln\left(\frac{\varepsilon}{1-\varepsilon}\right)\right)\right)} \ll n^{\frac{\varepsilon n}{2}}.$$

Since the opposite inequality holds for $n = \frac{2}{\varepsilon}$, there must be a solution in the interval $\left[\frac{2}{\varepsilon}, \infty\right)$. This implies that there are arbitrarily large solutions to the equation, as desired; taking n to be the integer closest to one of these solutions will give the desired value of n. □

6. Acknowledgements

Thank you to my mentor, Dr. Minh-Tâm Trinh, for guiding me to the discovery of the above results while also taking the time to review this paper. Thank you also to Prof. Zeev Dvir of Princeton University for his insights about matrix rigidity, to Prof. Matt Baker of the Georgia Institute of Technology for his context on vanishing sums of roots of unity, and to Prof. Alexander Razborov for his information about the faithful orthogonality dimension of graphs. Thank you to Dr. Tanya Khovanova and Dr. Felix Gotti, who provided valuable feedback on this paper. Lastly, thank you to the PRIMES-USA program for the amazing opportunity to conduct this research.

References

[Che21] Zhangchi Chen. "On nonsingularity of circulant matrices". In: *Linear Algebra and its Applications* 612 (2021), pp. 162–176.

[CLP17] Ernie Croot, Vsevolod F Lev, and Péter Pál Pach. "Progression-free sets in are exponentially small". In: *Annals of Mathematics* (2017), pp. 331–337.

[DE17] Zeev Dvir and Benjamin Edelman. "Matrix rigidity and the Croot-Lev-Pach lemma". In: *arXiv preprint arXiv:1708.01646* (2017).

[Dia90] Persi Diaconis. "Patterned matrices". In: *Proc. of Symposia in Applied Mathematics*. Vol. 40. 1990, pp. 37–58.

[DL19] Zeev Dvir and Allen Liu. "Fourier and circulant matrices are not rigid". In: *arXiv preprint arXiv:1902.07334* (2019).

[EG17] Jordan S Ellenberg and Dion Gijswijt. "On large subsets of with no three-term arithmetic progression". In: *Annals of Mathematics* (2017), pp. 339–343.

[Eti+11] Pavel I Etingof et al. *Introduction to representation theory*. Vol. 59. American Mathematical Soc., 2011.

[GH20] Alexander Golovnev and Ishay Haviv. "The (generalized) orthogonality dimension of (generalized) Kneser graphs: Bounds and applications". In: *arXiv preprint arXiv:2002.08580* (2020).

[GRW18] Alexander Golovnev, Oded Regev, and Omri Weinstein. "The minrank of random graphs". In: *IEEE Transactions on Information Theory* 64.11 (2018), pp. 6990–6995.

[IP64] I Martin Isaacs and DS Passman. "Groups with representations of bounded degree". In: *Canadian Journal of Mathematics* 16 (1964), pp. 299–309.

[LL00] Tsit Yuen Lam and Ka Hin Leung. "On vanishing sums of roots of unity". In: *Journal of algebra* 224.1 (2000), pp. 91–109.

[LR11] AK Lal and A Satyanarayana Reddy. "Non-singular circulant graphs and digraphs". In: *arXiv preprint arXiv:1106.0809* (2011).

[LSS89] László Lovász, Michael Saks, and Alexander Schrijver. "Orthogonal representations and connectivity of graphs". In: *Linear Algebra and its applications* 114 (1989), pp. 439–454.

[Mac03] David JC MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.

[Mey73] Carl D Meyer Jr. "Generalized inverses and ranks of block matrices". In: *SIAM Journal on Applied Mathematics* 25.4 (1973), pp. 597–602.

[PS+72] George Pólya, Gabor Szegö, et al. *Problems and Theorems in Analysis: Series, integral calculus, theory of functions*. Springer, 1972.

[Siv10] Gary Sivek. "On vanishing sums of distinct roots of unity". In: (2010).

[Svo20] Karl Svozil. "Faithful orthogonal representations of graphs from partition logics". In: *Soft Computing* 24.14 (2020), pp. 10239–10245.

[Ter99] Audrey Terras. *Fourier analysis on finite groups and applications*. 43. Cambridge University Press, 1999.

[Val77] Leslie G Valiant. "Graph-theoretic arguments in low-level complexity". In: *International Symposium on Mathematical Foundations of Computer Science*. Springer. 1977, pp. 162–176.