Algebraic Techniques in Coding Theory: Entropy Vectors, Frames, and Constrained Coding

Thesis by

Matthew Thill

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2016

(Defended August 19, 2015)

To my parents, my brothers, and my friends.

Acknowledgements

First and foremost, I would like to warmly thank my advisor, Prof. Babak Hassibi. His constant guidance and support aside, Babak was the one who originally made me excited in this field, and who gave me my first chance to be a part of it. I have always appreciated his ability to frame problems in an understandable and approachable manner. No matter how lost or stressed I became in my work, I always had a sense of clarity after talking to him, and I cannot begin to express how much I benefited from his intuition. But most importantly, I am privileged to have been able to work with a man who connected so well with his students. Graduate school can at times be a lonely experience, but he was always by our sides when we were pushing for a late night deadline.

I would like to express my appreciation to Prof. Palghat P. Vaidyanathan, Prof. Venkat Chandrasekaran, Prof. Jehoshua Bruck, Prof. Victoria Kostina, and Dr. Dariush Divsalar for serving on my defense committee. Their interest in my work and their willingness to review my thesis is invaluable. In particular, I would like to thank Dr. Divsalar for allowing me the chance to work with him over the last few months, and to benefit from his experience.

I am indebted to my friends and labmates Wei Mao, Amin Khajehnejad, and Sormeh Shadbakht, with whom I conducted my first research in this lab and without whom I never would have found my footing (and in Sormeh's case, my car). I thank Teja Sukhavasi for driving me to Hunan Chinese restaurants late at night during my first few years. Surviving that food gave me confidence that I could survive the rest of grad school. A special thanks are due to Kishore Jaganathan and Samet Oymak for being such great travel companions to conferences, to Christos Thrampoulidis and Ramya Korlakai Vinayak for keeping our lab lively, and to Wael Halbawi for keeping me motivated in my final year of grad school, and with whom I worked on my last projects. My sincere thanks go to our group secretary Shirley Beatty for all of her help in keeping us organized and for her patience which knows no bounds. And believe me, I've tested it. To all of you, and to everyone who passed through our group, thank you for making our lab such an enjoyable environment and a fun place to spend the last few years.

Equally important have been my friends outside of my lab. My thanks are due to Patrick Wong for being an accommodating roommate and a good friend, and for lending me his ear during hard

times. To Carlos Gonzalez, thank you for helping me break out of my Caltech shell and see what Los Angeles has to offer, and for all of your advice over the years. And to Ron Appel, thank you for your sense of humor, the countless crossword puzzles, engaging conversations, and for keeping your advice to yourself. The Pasadena running community has added color to my life over the last few years. Thank you Team R.E.D. for pushing me to my limits and for taking care of me when I go beyond them. In particular Chloe Fata, without you I would most assuredly still be passed out in the middle of Millennium Park. A part of this thesis belongs to you. Better get writing.

Finally, and most importantly, I want to thank my family for being a constant source of love, support, and encouragement. To my brothers Peter, David, and Michael, and especially my wonderful parents Laura and Mark: through my undergraduate and graduate studies, I have been at Caltech for just shy of ten years, and over that time I have seen many peers and friends come and go. It's not easy to be in the same place for so long, with none of the same people. But over the past decade you have always provided me a comfortable and familiar home, and a reminder of my self-identity and my roots. Thank you for lending me to this place for so long, for investing so much in me, and for making me feel so special. I love you all.

Abstract

The study of codes, classically motivated by the need to communicate information reliably in the presence of error, has found new life in fields as diverse as network communication, distributed storage of data, and even has connections to the design of linear measurements used in compressive sensing. But in all contexts, a code typically involves exploiting the algebraic or geometric structure underlying an application. In this thesis, we examine several problems in coding theory, and try to gain some insight into the algebraic structure behind them.

The first is the study of the entropy region—the space of all possible vectors of joint entropies which can arise from a set of discrete random variables. Understanding this region is essentially the key to optimizing network codes for a given network. To this end, we employ a group-theoretic method of constructing random variables producing so-called "group-characterizable" entropy vectors, which are capable of approximating any point in the entropy region. We show how small groups can be used to produce entropy vectors which violate the Ingleton inequality, a fundamental bound on entropy vectors arising from the random variables involved in linear network codes. We discuss the suitability of these groups to design codes for networks which could potentially outperform linear coding.

The second topic we discuss is the design of frames with low coherence, closely related to finding spherical codes in which the codewords are unit vectors spaced out around the unit sphere so as to minimize the magnitudes of their mutual inner products. We show how to build frames by selecting a cleverly chosen set of representations of a finite group to produce a "group code" as described by Slepian decades ago. We go on to reinterpret our method as selecting a subset of rows of a group Fourier matrix, allowing us to study and bound our frames' coherences using character theory. We discuss the usefulness of our frames in sparse signal recovery using linear measurements.

The final problem we investigate is that of coding with constraints, most recently motivated by the demand for ways to encode large amounts of data using error-correcting codes so that any small loss can be recovered from a small set of surviving data. Most often, this involves using a systematic linear error-correcting code in which each parity symbol is constrained to be a function of some subset of the message symbols. We derive bounds on the minimum distance of such a code based on its constraints, and characterize when these bounds can be achieved using subcodes of Reed-Solomon codes.

Contents

Chapter 1

Introduction

The broad intent of this thesis is to explore a set of problems in coding theory, where the term "coding theory" is in and of itself used broadly. In the context of information theory and communications, classical coding theory is often associated with the transmission of a message in a manner which is robust to various types of corruption. For instance, if we were to encode a message as a length-n vector of zeros and ones $\mathbf{v} \in \{0,1\}^n$ which is transmitted to a receiver, the receiver would ideally be able to decipher the original message even if it did not correctly receive some of the n symbols. A '0' may have been erased in the transmission process, or may have been incorrectly interpreted as a '1' by the receiver. The classical solution to this problem is to choose n large enough so that the vectors $\mathbf{v}$ corresponding to each of the possible messages can be designed to have mutually large Hamming distance between each other, leading to the notion of an error-correcting code.

But today the term "coding theory" encompasses a wide range of problems involving both the communication and the storage of information. For instance, the Internet has demanded efficient protocols to transmit information from a set of sources to a group of receivers over a network, sparking the field of network coding. In certain signal processing examples, it is desirable to encode messages as vectors over $\mathbb{C}^n$ rather than the binaries, and to have these vectors have large *angular* separation rather than Hamming distance. This leads to the notion of a spherical code—a set of points which are well-spaced over a high-dimensional sphere. Spherical codes are closely connected to the problem of constructing sets of vectors or frames with low coherence, a field which in turn has strong connections to the construction of compressive sensing matrices for sparse signal recovery. Even classical error-correcting codes are finding new applications, now in the storage and protection of large amounts of information. Companies commonly have many file servers which are subject to crashes and require a degree of redundancy in their data. More and more, these companies are moving away from naively making multiple copies of their files in favor of encoding the data as an error-correcting code, storing each symbol of a codeword on a different server. This can significantly reduce the number of servers

needed to protect the files.

In what follows, we will study a handful of these problems in different ways, but our approach will typically involve algebraic methods. Algebra is, of course, no stranger to coding theory. Indeed, arguably one of the most famous classes of error-correcting codes is that of Reed-Solomon codes, which are elegantly constructed subspaces of vector spaces over finite fields. As a result, many popular Reed-Solomon decoders—the Berlekamp-Massey decoder, for instance—employ polynomial arithmetic to correct errors in codewords. We will encounter Reed-Solomon codes in Chapter 5. But we will venture beyond finite field arithmetic. Chapters 2, 3, and 4 will require a more general group theoretic approach (though finite fields will certainly show up). In particular, we will use tools from group action theory and representation theory to construct various codes with coveted properties, thereby giving us some algebraic structural insight into the problems at hand.

1.1 Entropy Vectors and the Ingleton Inequality

A great deal of interest has been invested in determining what types of coding schemes can achieve capacity on a network, particularly in light of the revelation by Dougherty, Freiling, and Zeger [40] that simple linear codes are insufficient in some cases. One of the most general ways to view the network coding problem is to consider each message sent over an edge $e \in \mathcal{E}$ of the network as a random variable X_e, typically taking a value over a discrete or finite set of possible messages. For any network coding protocol, we can determine the vector of joint entropies $(H(X_e, e \in \alpha))_{\emptyset \neq \alpha \subseteq \mathcal{E}} \in \mathbb{R}^{2^{|\mathcal{E}|}-1}$. This is fittingly referred to as an *entropy vector*, and many quantities of interest for the network is a function of the set of associated entropy vectors.

We call the set of all possible entropy vectors arising from n discrete random variables the *entropy region* in $\mathbb{R}^{2^n-1}$, denoted Γ_n^*. A network with n edges enforces a set of constraints on the entropy region, so given a quantity of interest (for example, the mutual information between a set of sources and receivers), we can conceivably determine the optimal network code by optimizing this quantity over the portion of Γ_n^* carved out by the network. Unfortunately, Γ_n^* has only been classified for $n \leq 3$. There has been some progress in understanding Γ_n^* for larger n. For instance, Zhang and Yeung showed that its closure is a convex cone [114], but Matúš [73] proved that $\overline{\Gamma_n^*}$ is not polymatroidal for $n \geq 4$. Many have sought inner bounds on the entropy region, such as the space of linear-representable matroidal rank functions (see Section 2.3), but even this region is only known for small values of n (though lately new linear rank inequalities have been discovered more frequently [19, 39, 41, 42, 64]). The entropy region still remains largely mysterious, however, so new ways of studying it are always of interest.

In Chapter 2, we discuss a method of constructing entropy vectors from groups. Essentially,

we fix a set of n subgroups G_i, $i = 1, ..., n$, of a finite group G from which we randomly draw an element g. For each i, we define a random variable X_i by determining the coset of G_i in which g lies. The resulting entropy vector is referred to as a *group-characterizable* entropy vector, and it has been shown by Chan and Yeung [23] that any element in the closure $\overline{\Gamma_n^*}$ can be approximated by group-characterizable entropy vectors. This gives us a springboard to study the entropy region by characterizing the types of entropy vectors which can arise from various finite groups.

By the same token, given an entropy inequality that is known to constrain many joint sets of random variables, we can endeavor to design random variables which are not limited by this inequality using group-theoretic methods. The inequality in question for us will be the *Ingleton Inequality*, which is a fundamental constraint on the dimensions of four linear subspaces of a vector space. As a result, any four random variables X_i, '$i = 1, ..., 4$, arising from a *linear* network code must satisfy the Ingleton Inequality, which can be written as

$$h_1 + h_2 + h_{34} + h_{123} + h_{124} \leq h_{12} + h_{13} + h_{14} + h_{23} + h_{24}. \tag{1.1}$$

While it is known that there are entropy vectors in Γ_4^* which violate the Ingleton Inequality, it is not immediately clear how to construct random variables that produce them, or how and when they can be incorporated into network codes.

In Chapter 2, we show how to produce Ingleton-violating group-characterizable entropy vectors using subgroups of projective linear groups $PGL(n, q)$ and general linear groups $GL(n, q)$ for certain values of n and q. Since these are matrix groups, we are able to give concrete characterizations of the elements in each of the subgroups G_i, $i = 1, ..., 4$. Using the theory of group actions, we are able to generalize our constructions to a broader class of groups, and to understand why they violate Ingleton from a more geometric perspective. Furthermore, we broach the subject of how these groups might arise in network codes, particularly in the form of the *group network codes* described in Section 2.5.

Chapter 2 is joint work with Wei Mao and Babak Hassibi, and appears in [70] and [69].

1.2   Low-Coherence Frames

In Chapters 3 and 4, we shift our focus to the construction of sets of vectors in $\mathbb{C}^m$ or $\mathbb{R}^m$, called *frames*, which have mutually small correlation between each other. The maximum magnitude of the inner product between two frame vectors is called the *coherence* of the frame, and a classic problem in frame theory is to find frames which achieve low coherence. For frames with n vectors, where $n \leq m$, this can be done by choosing a set of orthogonal vectors. Thus, in some sense, when $n > m$ a low-coherence frame is an approximation of a basis for a vector space. Today, some might view this

problem as more applicable to signal estimation than coding theory since a great deal of recent study has explored the sparse-signal recovery properties of these frames [1, 75, 80, 91]. In particular, when taken to be the columns of a matrix they tend to have good RIP constants, and hence lend themselves to the linear programming-based compressive sensing algorithms described in [12] and [11]. They also have provably good performance with the One-Step Thresholding (OST) algorithm from [1].

At its core, however, this problem has its roots in coding theory. If we normalize the frame elements, then constructing a low-coherence frame is almost equivalent to the problem of designing a spherical code—a set of n points spread around the m-dimensional unit sphere which have mutually large angular separation (see Fig. 1.1). Spherical codes have applications, for instance, when we wish to encode a message as a vector in which each entry is subject to independent Gaussian noise.

Figure 1.1: Illustration of a spherical code—a set of points spaced out around the unit sphere. Ideally, the inner product between the unit vectors corresponding to two such points should be small in magnitude, corresponding to a large angular separation between the vectors.

Furthermore, the constructions we will present in Chapters 3 and 4 are based on the concept of *group codes* presented by Slepian in the 1960s [90]. A group code is formed by taking the image of a vector $\mathbf{v} \in \mathbb{C}^m$ under a multiplicative group of unitary matrices $\mathcal{U} = \{\mathbf{U}_1, ..., \mathbf{U}_n\} \subseteq \mathbb{C}^{m \times m}$ to form the set $\{\mathbf{U}_i\mathbf{v}\}_{i=1,...,n}$. This method reduces the total number of distinct inner product magnitudes between the frame elements from a possible $\binom{n}{2}$ to a mere $n - 1$, with the inner products taking the form $\mathbf{v}^*\mathbf{U}_i\mathbf{v}$ (ignoring the inner product corresponding to the identity matrix, which is simply the inner product of $\mathbf{v}$ with itself). By choosing the group $\mathcal{U}$ and the vector $\mathbf{v}$ appropriately, we will see in Chapter 3 that the resulting frame becomes the columns of a submatrix of the $n \times n$ Discrete Fourier

Matrix, which (after normalizing the frame elements) takes the form

$$\mathbf{M} = \frac{1}{\sqrt{m}} \begin{bmatrix} 1 & \omega^{a_1} & \omega^{a_1 \cdot 2} & \ldots & \omega^{a_1 \cdot (n-1)} \\ 1 & \omega^{a_2} & \omega^{a_2 \cdot 2} & \ldots & \omega^{a_2 \cdot (n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{a_m} & \omega^{a_m \cdot 2} & \ldots & \omega^{a_m \cdot (n-1)} \end{bmatrix}, \tag{1.2}$$

where $\omega = e^{2\pi i/n}$ and $a_i \in \{0, ..., n-1\}$ for each i.

We then show for certain values of n that by taking some care in choosing the frequencies a_i, or equivalently controlling the specific forms of the matrices $\mathbf{U}_i$, we can further reduce the number of distinct inner product magnitudes down to $\frac{n-1}{\kappa}$, where κ is a divisor of $n-1$. The motivation behind this is that frames of the form (1.2) are known to achieve the lowest possible coherence for given dimensions m and n when all the mutual inner products between the frame elements have the same magnitude. Such frames are examples of what are called *Grassmanian* frames, and are very important in communications and coding theory [91]. Unfortunately, few Grassmanian frames are known, and those that take the form of (3.18) only arise when the $\{a_i\}$ form a rare collection of numbers called a *difference set* [110]. By reducing the number of distinct inner product magnitudes we essentially approximate a Grassmanian frame, and as a result we are able to prove upper bounds on the coherence of our frames which are rather tight in practice.

Our construction utilizes a group-theoretic trick to select the frequencies $\{a_i\}$, and in Chapter 4 we show that this technique can be extended to form frames by selecting a subset of rows of a generalized group Fourier matrix, which is the natural generalization of the DFT matrix. By extending our results to this context, we will open ourselves to a much richer set of frames which can be realized as group codes resulting from a broader set of groups $\mathcal{U}$. This will allow us to construct low-coherence frames achieving a wider range of dimensions $m \times n$, and to design frame vectors whose entries come from a much smaller alphabet than those of the form (3.18). In certain cases, our frames contain only ± 1 entries and are actually composed of subsets of rows from Hadamard matrices.

There are several important advantages to our frames over those constructed from popular random methods. First, the fact that they are designed from group representations allows us to analyze the inner products between frame elements in terms of the characters of the group. The algebraic manner in which we select the representations facilitates the proof of some very sharp bounds on coherence.

Furthermore, it enables us to study other aspects of our frames, including their *average coherence*. This quantity was described in [1] and [75] which showed that when conditions on both the usual coherence and the average coherence are satisfied (the so called "Coherence Property" and "Strong Coherence Property") then the matrix $\mathbf{M} \in \mathbb{C}^{m \times n}$ whose columns are the frame elements can provably

Figure 1.2: An example of a constrained code. Each code symbol c_i is a function of only the symbols m_j to which it is connected in the bipartite graph.

be used to estimate sparse vectors $\mathbf{x} \in \mathbb{C}^n$ from the set of measurements $\mathbf{y} = \mathbf{Mx}$, even in the presence of added noise $\mathbf{e} \in \mathbb{C}^n$. More importantly, this estimation can be done quickly using the aforementioned OST algorithm, a single-step process that involves using the largest entries of $\mathbf{M}^*\mathbf{y}$ to estimate the support of $\mathbf{x}$. We will see in Chapter 4 that for our frames the average coherence can be computed explicitly, allowing us to determine rather precise regimes under which the Coherence Property and Strong Coherence Property are satisfied.

The work in Chapter 3 appears in [96–100]. The results of Chapter 4 are largely from our work in [95].

1.3 Constrained Coding

In Chapter 5, we move on to discuss the subject of constrained coding. Loosely speaking, this problem deals with encoding a set of s messages and a codeword composed of n symbols over some alphabet. Each of the n code symbols is constrained to be a function of some *subset* of the s messages. In keeping with the language of classical coding theory, we typically think of our messages as a single length-s vector $\mathbf{m} = [m_1, ..., m_s]$ of symbols over some alphabet, which we encode as the vector $\mathbf{c} = [c_1, ..., c_n]$ subject to the constraints. We often represent the coding constraints via a bipartite graph $G = (\mathcal{M}, \mathcal{V}, \mathcal{E})$, where the vertices $\mathcal{M}$ represent the message symbols and those in $\mathcal{V}$ the code symbols. A code symbol $c_i \in \mathcal{V}$ is then constrained to be a function of the message symbols contained in its neighborhood in $\mathcal{M}$. (See Figure 1.2).

This problem arises in a variety of contexts, such as in the case of a sensor network where each sensor has access to a subset of measurements. We would like to arrange our sensors appropriately so that if a small number of them malfunction, we can still recover all of the measurements. In the absence of malfunctions, we would like to be able to obtain all of the information more efficiently from

a smaller number of sensors. This suggests the use of a systematic code, i.e., the message symbols m_i appear explicitly as a subset of the code symbols $\{c_j\}_{j=1,\dots,n}$. The problem also arises in the field of distributed storage of data, particularly in the recent field of locally repairable codes [24,54,58]. These are systematic codes which divide the message symbols into several "local" groups. The remaining code symbols are then designated to protect one or several of these groups. Thus if one of the systematic symbols is lost, it can be recovered just by accessing the code symbols which protect its local group (along with the remaining systematic symbols in its group). The motivation for this setup is in the situation where we would like to protect a large amount of data which is stored in a large file server or a set of hard drives. We assume some of the hard drives contain the data in its original form, and each remaining hard drive stores a function of the data contained in one local group of hard drives. Using a systematic locally repairable code ensures the security of the data in the event of several drives crashing, allows for quick download of the data in the event of no crashes in the systematic portion of the code, and necessitates only a small set of hard drives to be accessed to repair a single crashed drive in one of the local groups.

Thus in Chapter 5 we will pay particularly close attention to *systematic* constrained codes, with an eye toward analyzing the code's *minimum distance* based on the topology of the bipartite graph which constrains it (as in Figure 1.2). The minimum distance—the smallest Hamming distance between any two codewords—determines the maximum number of code symbols which could be lost or corrupted while still ensuring that the entire codeword could be correctly determined from nearest neighbor decoding (that is, selecting the valid codeword which is closest to the corrupted codeword in Hamming distance). We will primarily focus our attention on systematic linear codes, and seek subcodes of Reed-Solomon codes which meet a set of constraints. Codes of this form are desirable for their known fast decoding algorithms, e.g. [5,7,71,77]. We will derive bounds on the minimum distance of constrained codes that are reminiscent of the cut set-type bounds from [48], and we will refine these bounds in the case that we require a systematic code. For certain types of constraining graphs, we provide code constructions which achieve these bounds, utilizing Reed-Solomon and MDS codes in our designs. The results of Chapter 5 are based on work with Wael Halbawi and Babak Hassibi which appeared in [49].

Chapter 2

Violating the Ingleton Inequality Using Finite Groups

2.1 Entropy Vectors

Let $X_1, ..., X_n$ be a set of jointly-distributed discrete random variables. For any subset $\alpha \subseteq \{1, ..., n\}$ let X_α denote set of random variables indexed by α,

$$X_\alpha := \{X_i \ : \ i \in \alpha\},$$

and let h_α denote their joint entropy,

$$h_\alpha = H(X_\alpha) = H(X_i \ : \ i \in \alpha).$$

The *entropy vector* associated with the random variables is the $(2^n - 1)$-tuple consisting of the joint entropies of all nontrivial subsets of the X_i:

$$\mathbf{h} := (h_\alpha \ : \ \varnothing \neq \alpha \subseteq [n]) \in \mathbb{R}^{2^n - 1}.$$

We denote by Γ_n^* the set of all possible entropy vectors arising from n discrete random variables. Its closure $\overline{\Gamma_n^*}$ is in fact a convex cone [114], which can be shown through timesharing arguments.

2.2 Group-Characterizable Entropy Vectors

We now discuss a connection between groups and entropy vectors which will allow us to use group theoretic methods to study the entropy region. Let G be a finite group with subgroups $G_1, ..., G_n$. Let Λ be a random variable which is uniformly distributed on the elements of G, and let $X_i = \Lambda G_i$

for each $i = 1, ..., n$. That is, X_i is a random variable that takes the value of the left coset of G_i in G in which Λ lies. As such, since all these cosets are the same size, X_i is uniformly distributed over G/G_i, the $\frac{|G|}{|G_i|}$ cosets of G_i in G.

The entropy vector $\mathbf{h}$ arising from these discrete random variables X_i is called a *group-characterizable entropy vector*. From our above discussion, we see that

$$h_i = H(X_i) = \log\left(\frac{|G|}{|G_i|}\right). \tag{2.1}$$

Furthermore, for any subset $\alpha \subseteq [n]$, the cosets $X_\alpha = \{X_i : i \in \alpha\}$ are uniquely determined by the coset $\Lambda G_\alpha \in G/G_\alpha$ where G_α is the intersection $\bigcap_{i \in \alpha} G_i$, also a subgroup of G. Thus its entropy is the same as that of the random variable ΛG_α, so we will identify this with the random variable X_α and we have by a similar token that

$$h_\alpha = H(X_\alpha) = \log\left(\frac{|G|}{|G_\alpha|}\right). \tag{2.2}$$

Interestingly enough, it turns out that *any* entropy vector can be approximated by a scaled group-characterizable entropy vector [23]. The idea is as follows: suppose X is a discrete random variable taking on N possible values $1, ..., N$ with respective probabilities $p_1, ..., p_N$. If we take T independent copies of X, vectorized as $\mathbf{X} := (X^{(1)}, ..., X^{(T)})$, then $H(\mathbf{X}) = TH(X)$. The *strongly typical sequences* are those realizations of $\mathbf{X}$ where approximately $p_i T$ entries $X^{(j)}$ take on the value i, for each $i = 1, ..., N$. The number of such sequences is approximately

$$\binom{T}{p_1 T \ \ldots \ p_N T} := \frac{T!}{(p_1 T)! \ldots (p_N T)!}, \tag{2.3}$$

where we assume the quantities $p_i T$ are integers. For T large, a strongly typical sequence will occur with probability approaching 1, and each of these sequences is equiprobable. Thus, we can approximate $H(X)$ as

$$H(X) = \frac{1}{T} H(\mathbf{X}) \approx \frac{1}{T} \log\left(\frac{T!}{(p_1 T)! \ldots (p_N T)!}\right). \tag{2.4}$$

Now, consider the symmetric group $G = S_T$ of permutations on T elements, which has size $T!$. Suppose we partition these elements into subsets of size $p_1 T, ..., p_N T$, and let G_X be the subgroup of G of permutations which preserve these subsets. Then G_X has size $(p_1 T)! \ldots (p_N T)!$, and $H(X) = \frac{1}{T} \log\left(\frac{|G|}{|G_X|}\right)$.

Now suppose we have n discrete random variables, $X_1, ..., X_n$. It is not too difficult to see that we can find a different partition of T for each X_j, $j = 1, ..., n$, such that when G_j is chosen to be the set

of permutations that respects the j^{th} partition, the group-characterizable entropy vector associated to the G_j is approximately a scaled version of the entropy vector corresponding to the original X_j. A more detailed argument appears in [23] where it is rigorously proven that if Υ_n is the region of group-characterizable entropy vectors for n variables, and $\overline{\text{cone}(\Upsilon_n)}$ the closure of its convex cone, then $\overline{\Gamma_n^*} = \overline{\text{cone}(\Upsilon_n)}$.

While this process can indeed approximate any entropy vector, and can presumably be used to allow us to study the entropy region using group theoretic techniques, it often requires the set T (and consequently the permutation group G) to group very large. This begs the question of whether we can identify small groups which can yield entropy vectors with interesting properties.

2.3 Matroidal Bounds on the Entropy Region

Entropy vectors have an important connection to matroid theory. A *matroid* is a set $\mathcal{M}$ of elements together with a *rank function* $r : 2^{\mathcal{M}} \to \mathbb{Z}_{\geq 0}$ satisfying the following:

1. $r(\varnothing) = 0$, and for any $\varnothing \neq \mathcal{A} \subseteq \mathcal{M}$ we have $r(\mathcal{A}) \leq |\mathcal{A}|$.

2. $r(\cdot)$ is monotonic: If $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{M}$, then $r(\mathcal{A}) \leq r(\mathcal{B})$.

3. $r(\cdot)$ is submodular: For any subsets $\mathcal{A}$ and $\mathcal{B}$ in $\mathcal{M}$, we have

$$r(\mathcal{A} \cup \mathcal{B}) + r(\mathcal{A} \cap \mathcal{B}) \leq r(\mathcal{A}) + r(\mathcal{B}).$$

For example, any set of vectors in a vector space satisfy these conditions when $r(\cdot)$ is taken to be the usual rank function on a vector space. Thus in some sense, a matroid is a generalization of a vector space. A matroid is called *linear representable* (or just "representable") if the set $\mathcal{M}$ and the function $r(\cdot)$ can be mapped to a set of vectors in a vector space with the same ranks of corresponding subsets. For any subset $\mathcal{A} \subseteq \mathcal{M}$ we will often use the notation $r_{\mathcal{A}} := r(\mathcal{A})$, and we will speak of the *rank vector* of a matroid:

$$\mathbf{r} = (r_{\mathcal{A}} \ : \ \varnothing \neq \mathcal{A} \subseteq \mathcal{M}) \in \mathbb{Z}^{2^{|\mathcal{M}|}-1}.$$

If $\mathcal{M}$ is taken to be a set of random variables $\mathcal{M} = \{X_1, ..., X_n\}$ and we consider the entropy function $h : 2^{\mathcal{M}} \to \mathbb{R}$, $h(\mathcal{A}) = H(X_i \ : \ i \in \mathcal{A})$, then we can see that $h(\cdot)$ satisfies conditions 2 and 3. A function with these properties, together with the set $\mathcal{M}$, is called a *polymatroid*. In the context of entropy, conditions 2 and 3 together are called the *Shannon inequalities*, and they correspond to all conditions which can be expressed as the conditional mutual information of a set of random variables

being nonnegative:

$$I(X_\alpha; X_\beta | X_\gamma) = H(X_\alpha, X_\gamma) + H(X_\beta, X_\gamma) - H(X_\alpha, X_\beta, X_\gamma) - H(X_\gamma) \geq 0, \tag{2.5}$$

where α, β, and γ are subsets of $\{1, ..., n\}$, and X_α, X_β, X_γ are the sets of random variables corresponding to these subsets.

From these observations, it is clear that the Shannon inequalities form an outer bound for the space of entropic vectors. They do not, however, define the entropy region for all n, and in fact as we will discuss shortly, they do not even completely describe the closure of the cone of representable matroidal rank vectors.

2.4 The Ingleton Inequality

It turns out that the Shannon inequalities completely characterize the region of representable rank vectors for matroids $\mathcal{M}$ of up to 3 elements. But larger representable matroids must additionally satisfy the following constraint, called the *Ingleton inequality*:

Theorem 1 (Ingleton Inequality). *Let S_1, S_2, S_3, and S_4 be subspaces of a vector space, and for any $\alpha \subseteq \{1, 2, 3, 4\}$, let r_α denote the rank of the subspace generated by the S_i, $i \in \alpha$. Then*

$$r_1 + r_2 + r_{34} + r_{123} + r_{124} \leq r_{12} + r_{13} + r_{14} + r_{23} + r_{24}. \tag{2.6}$$

Proof. This was proven by Ingleton in 1971. [59] $\qquad\qquad\square$

We will speak of an entropy vector $\mathbf{h} \in \overline{\Gamma}_4^*$ satisfying the Ingleton inequality if its entries obey the relation

$$h_1 + h_2 + h_{34} + h_{123} + h_{124} \leq h_{12} + h_{13} + h_{14} + h_{23} + h_{24}. \tag{2.7}$$

The Ingleton inequality is the simplest example of a non-Shannon inequality. Together with the Shannon inequalities, it completely characterizes the region of representable matroids up to size $n = 4$, [53] though for $n \geq 5$ there are other defining inequalities which these do not imply. [19, 39, 41, 42, 64]

The reason to go into detail discussing the region of linear representable matroid rank vectors is that [53] shows that any such rank vector is indeed entropic. Thus, this region is an inner bound for the entropy region. It is, however, a proper inclusion, since there exist entropy vectors which violate the Ingleton inequality [53, 72].

Now let us return to the topic of group-characterizable entropy vectors. In this case, we have a group G and four subgroups G_1, G_2, G_3, and G_4. Instead of dealing with the entries of a rank vector r_α, $\alpha \subset [4]$, we deal with the joint entropies $h_\alpha = \log \frac{|G|}{|G_\alpha|}$, where $G_\alpha = \cap_{i \in \alpha} G_i$. Substituting these entries into (2.7), and after rearranging terms and taking the exponential of both sides, we obtain the group-based analog of the Ingleton inequality:

$$|G_{12}||G_{13}||G_{14}||G_{23}||G_{24}| \leq |G_1||G_2||G_{34}||G_{123}||G_{124}|. \tag{2.8}$$

Again, we emphasize that since group-representable entropy vectors can approximate the entire entropy region, not all such vectors need satisfy (2.8). But it turns out that many of the more "basic" groups can only produce group-characterizable entropy vectors which satisfy the Ingleton inequality, as we see from the following conditions presented in [22, 66, 70]:

Theorem 2. *Let G be a group with subgroups G_1, G_2, G_3 and G_4. Then the following conditions suffice for these subgroups satisfying the Ingleton inequality of (2.8):*

1. *G is abelian.*

2. *G_i is a normal subgroup of G for each i.*

3. *The set product $G_1 G_2 := \{g_1 g_2 \; : \; g_1 \in G_1, \; g_2 \in G_2\}$ is a subgroup of G. Equivalently, $G_1 G_2 = G_2 G_1$.*

4. *$G_i = 1$ or G for some i.*

5. *$G_i = G_j$ for some $i \neq j$.*

6. *$G_{12} = 1$.*

7. *G_i is a subgroup of G_j for some $i \neq j$.*

Proof. Condition 1 is proved in [22]. Condition 2 appears in [66]. The remaining conditions are proven in [70], which employed these conditions in a computer search to find the smallest Ingleton-violating groups. $\square$

The fact that all abelian groups must satisfy the Ingleton inequality can be seen as a generalization of the fact that linear subspaces of a vector space must satisfy it in its original form (2.7). Vector spaces are, after all, a class of abelian groups. In light of this, we wish to identify instances of small nonabelian groups which produce entropy vectors violating the Ingleton inequality. Such groups have application in network coding problems, as we will see in Section 2.5.

2.5 Group Network Codes

Before proceeding further, we briefly comment on the applications of Ingleton-violating groups to network coding. In a typical network coding scenario, a network is represented as a directed, acyclic graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with vertices $\mathcal{V}$ representing the communication nodes, and edges $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ the communication channels of the network. There is a subset $\mathcal{S} \subset \mathcal{V}$ of vertices which represent the sources of the network, and each source $s \in \mathcal{S}$ is demanded by another subset $\mathcal{D}(s) \subset \mathcal{V}$. For any $v \in \mathcal{V}$, we will let $\mathcal{I}(v)$ denote the set of incoming edges to v, together with v itself if $v \in \mathcal{S}$. That is,

$$\mathcal{I}(v) := \begin{cases} \{e \in \mathcal{E} \ : \ e = (v', v), \ v' \in \mathcal{V}\} & v \notin \mathcal{S} \\ \{e \in \mathcal{E} \ : \ e = (v', v), \ v' \in \mathcal{V}\} \cup \{v\} & v \in \mathcal{S} \end{cases} \tag{2.9}$$

For simplicity, for any edge $e = (v_1, v_2) \in \mathcal{E}$ we will abuse notation and write $\mathcal{I}(e)$ for the set of incoming edges to the tail node of e, that is, $\mathcal{I}(e) := \mathcal{I}(v_1)$. Note that $\mathcal{I}(e) \subset \mathcal{S} \cup \mathcal{E}$.

A network code formally works as follows: each source s is identified with a random variable Y_s which takes some value in a certain alphabet $\mathcal{Y}_s$. We typically assume that the Y_s are uniformly distributed over their respective alphabets, and that they independent:

$$H(Y_s \ : \ s \in \mathcal{S}) = \sum_{s \in \mathcal{S}} H(Y_s). \tag{2.10}$$

Then to each edge $e = (v_1, v_2) \in \mathcal{E}$, we associate an encoded symbol Y_e from some alphabet $\mathcal{Y}_e$ which is a function ϕ_e of the incoming edge symbols and the symbol Y_{v_1} if v_1 is a source:

$$Y_e = \phi_e(Y_{e'} \ : \ e' \in \mathcal{I}(e)). \tag{2.11}$$

Likewise, for any $u \in \mathcal{D}(s)$, the demanded symbol Y_s should be uniquely determined from the incoming edges to u (and from u itself, if u happens to also be a source):

$$Y_s = \phi_{u,s}(Y_{e'} \ : \ e' \in \mathcal{I}(u)). \tag{2.12}$$

These conditions imply that

$$H(Y_e \mid Y_{e'}, \ e' \in \mathcal{I}(e)) = 0, \ \forall e \in \mathcal{E} \tag{2.13}$$

$$H(Y_s \mid Y_{e'}, \ e' \in \mathcal{I}(u)) = 0, \ \forall s \in \mathcal{S}, \ u \in \mathcal{D}(s). \tag{2.14}$$

Note that based on these criteria, any symbol Y_e, $e \in \mathcal{E}$, can be expressed directly as a function Φ_e of

the symbols Y_s, $s \in \mathcal{S}$.

In a *group network code* [20, 21], we start with a group G and for each $s \in \mathcal{S}$ and each $e \in \mathcal{E}$ we select subgroups G_s and G_e. The alphabets $\mathcal{Y}_s$ and $\mathcal{Y}_e$ are set to be the cosets of these subgroups in G, that is, $\mathcal{Y}_s = G/G_s$ and $\mathcal{Y}_e = G/G_e$. In keeping with our previous notation, for any subset $\mathcal{W} \subset \mathcal{S} \cup \mathcal{E}$, we will define the intersection subgroup $G_{\mathcal{W}} := \cap_{w \in \mathcal{W}} G_w$.

A group network code is one in which there is an element $g \in G$ such that the symbols Y_s and Y_e are equal to the cosets gG_s and gG_e, respectively. g is assumed to be uniformly selected from the elements of G, so that Y_s and Y_e are uniform random variables on their alphabets, with entropies $H(Y_s) = \log \frac{|G|}{|G_s|}$ and $H(Y_e) = \log \frac{|G|}{|G_e|}$. Let us examine what this implies about our chosen subgroups: First, note that the vector of source symbols $Y_{\mathcal{S}} := (Y_s \ : \ s \in \mathcal{S})$ is uniquely determined by the coset $gG_{\mathcal{S}}$. If our source random variables are to be independent, we must have $H(Y_{\mathcal{S}}) = \sum_{s \in \mathcal{S}} H(Y_s)$, which means that

$$\log \frac{|G|}{|G_{\mathcal{S}}|} = \sum_{s \in \mathcal{S}} \log \frac{|G|}{|G_s|}.$$

This translates to the requirement that

$$\prod_{s \in \mathcal{S}} |G_s| = |G|^{|\mathcal{S}|-1} |G_{\mathcal{S}}|. \tag{2.15}$$

In order for Y_e to be a well-defined function of the variables $\{Y_w \ : \ w \in \mathcal{I}(e)\}$, we must have that whenever there are distinct elements g and g' in G such that $gG_w = g'G_w$, $\forall w \in \mathcal{I}(e)$, then $gG_e = g'G_e$. This means that whenever $g^{-1}g' \in G_w$, $\forall w \in \mathcal{I}(e)$, then $g^{-1}g' \in G_e$, so we have the equivalent condition that

$$G_{\mathcal{I}(e)} \leq G_e. \tag{2.16}$$

By the same token, since Y_s must be a function of the variables $\{Y_w \ : \ w \in \mathcal{I}(u)\}$ for each $u \in \mathcal{D}(s)$, we also have the condition that

$$G_{\mathcal{I}(u)} \leq G_s, \ \forall u \in \mathcal{D}(s). \tag{2.17}$$

Example: Linear Network Codes. In the case of linear network codes, source messages Y_s are typically thought of as elements of a finite field, or vectors in a vector space $V = \mathbb{F}^n$ over a finite field $\mathbb{F}$. Each edge message Y_e is a linear function of the messages $\{Y_w \ : \ w \in \mathcal{I}(e)\}$, written $Y_e = \sum_{w \in \mathcal{I}(e)} M_{e,w} Y_w$ where $M_{e,w} \in \mathbb{F}^{n \times n}$. This can be realized as a group network code by setting $G = V^{\oplus |\mathcal{S}|}$, the direct sum of $|\mathcal{S}|$ copies of V. If our sources are $\mathcal{S} = \{s_1, ..., s_m\}$, then we set $G_{s_i} := V \oplus V \oplus ... \oplus 0 \oplus ... \oplus V$, where the 0 is in the i^{th} position of the direct sum. We define the groups G_e inductively as $G_e = \cap_{w \in \mathcal{I}(e)} G_w$.

In the case $n = 1$, where message symbols Y_s and Y_e are simply elements of the finite field $\mathbb{F}$, then if we have m independent sources, all of our random variables Y_e are functions of the vector $(Y_{s_1}, ..., Y_{s_m}) \in \mathbb{F}^{\oplus m}$. If some edge e is connected only to sources 1 and 2, we have $\mathcal{I}(e) = \{s_1, s_2\}$, and $G_e = 0 \oplus 0 \oplus F^{\oplus(m-2)}$, reflecting the fact that Y_e is a function of the coset $Y_{s_1} \oplus Y_{s_2} \oplus \mathbb{F}^{\oplus(m-2)}$ of G_e.

When we consider the entropy vector associated to the random variables $\{Y_t\}_{t \in \mathcal{S} \cup \mathcal{E}}$, we see that it is the group-characterizable entropy vector associated with the groups $\{G_t\}_{t \in \mathcal{S} \cup \mathcal{E}}$. For linear codes, since the overlying group G is abelian, we know from Theorem 2 that the subvector associated to any four of these random variables must satisfy the Ingleton inequality. Thus by characterizing Ingleton-violating groups, we could potentially develop group network codes which are more powerful than linear codes in the sense that their associated random variables can achieve a larger range of the entropy region. This would build on prior results [40] that that there exist networks for which linear codes cannot achieve capacity.

One limitation on the types of networks to which we can apply Ingleton-violating group codes is that more than two independent sources will produce random variables Y_s which must satisfy the Ingleton inequality:

Lemma 1. *Let $X_1, X_2, X_3,$ and X_4 be random variables. If X_i and X_j are independent, where $\{i, j\}$ is any pair other than $\{3, 4\}$, then the entropy vector of the X_i must satisfy Ingleton's inequality (2.7).*

Proof. By symmetry in the terms of the Ingleton inequality, we need only consider the cases $\{i, j\} = \{1, 2\}$ and $\{1, 3\}$. If we assume X_1 and X_2 are independent, we have $h_{12} = h_1 + h_2$. Also by submodularity we have $h_{13} + h_{23} \geq h_3 + h_{123}$ and $h_{14} + h_{24} \geq h_4 + h_{124}$. Thus,

$$h_{12} + h_{13} + h_{14} + h_{23} + h_{24} \geq h_1 + h_2 + h_3 + h_4 + h_{123} + h_{124} \tag{2.18}$$

$$\geq h_1 + h_2 + h_{34} + h_{123} + h_{124}. \tag{2.19}$$

The proof for the case $\{i, j\} = \{1, 3\}$ is similar. $\square$

Corollary 1. *Given a set of four random variables $X_1, X_2, X_3,$ and X_4, if any three are independent, then the associated entropy vector must satisfy the Ingleton inequality (2.7).*

Proof. This follows immediately from the Lemma 1. $\square$

2.6 The Smallest Ingleton-Violating Groups: $PGL(2, p)$

As we mentioned before, we would like to identify small groups with Ingleton-violating subgroups. It turns out that the smallest such group is $PGL(2,5)$, as was discovered in [68]. This is a *projective linear group*, which is defined as follows: For a prime power q, the *general linear group* $GL(n, q)$ is the multiplicative group of invertible $n \times n$ matrices with entries in the finite field $\mathbb{F}_q$. The center of this group is the set of scalar matrices

$$Z(GL(n,q)) = \left\{ \begin{bmatrix} \alpha & \ldots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \ldots & \alpha \end{bmatrix}, \ \alpha \in \mathbb{F}_q^\times \right\}.$$

The projective group is then defined to be the quotient $PGL(n, q) = GL(n, q)/Z(GL(n, q))$, which has size

$$|PGL(n,q)| = \frac{(q^n - 1)(q^n - q) \ldots (q^n - q^{n-1})}{q - 1}.$$

Thus $|PGL(2, 5)| = 120$.

For p a prime greater or equal to 5, let t be a primitive root in $\mathbb{F}_p$, i.e., a generator of the multiplicative group $\mathbb{F}_p^\times$, which is a cyclic group of size $p - 1$. Then $PGL(2, p)$ is generated by the matrices

$$A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \ B = \begin{bmatrix} 1 & 0 \\ 0 & t \end{bmatrix}, \ C = \begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix}.$$

Indeed, $GL(2, p)$ (and therefore $PGL(2, p)$) is generated by the elementary matrices

$$\begin{bmatrix} 1 & 0 \\ \alpha & 1 \end{bmatrix}, \ \begin{bmatrix} 1 & \beta \\ 0 & 1 \end{bmatrix}, \ \begin{bmatrix} t^i & 0 \\ 0 & 1 \end{bmatrix}, \ \begin{bmatrix} 1 & 0 \\ 0 & t^j \end{bmatrix}. \tag{2.20}$$

The matrices $\begin{bmatrix} 1 & 0 \\ \alpha & 1 \end{bmatrix}$ are the powers of A. (Note that $A^k = \begin{bmatrix} 1 & 0 \\ k & 1 \end{bmatrix}$, and any element of $\mathbb{F}_p$ is simply an integer k modulo p). Any matrix of the form $\begin{bmatrix} t^i & 0 \\ 0 & t^j \end{bmatrix}$ is a power of B multiplied by a scalar matrix: $t^i B^{j-i}$. Finally, the matrices $\begin{bmatrix} 1 & \beta \\ 0 & 1 \end{bmatrix}$ are simply the powers of A conjugated by the matrix

$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, and a quick calculation shows that if we define

$$B_1 := \begin{bmatrix} 1 & 0 \\ -2^{-1} & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 2^{-1} & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix}, \tag{2.21}$$

then $B_1 C = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.

Now consider the following subgroups of $PGL(2, p)$:

$$G_1 = \langle B_1, C \rangle,$$

$$G_2 = \langle A, B \rangle,$$

$$G_3 = \langle CB_1, A^{-1}BA \rangle,$$

$$G_4 = \langle B_1 C, B \rangle.$$

By inspection, we have $G_1 = \langle C \rangle \rtimes \langle B_1 \rangle \cong D_6$, the dihedral group with six elements. These are

$$G_1 = \{B_1^i C^j, \ 0 \le i < 2, \ 0 \le j < 3\}$$
$$= \left\{ \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix}, \begin{bmatrix} -1 & -1 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} -1 & -1 \\ 0 & 1 \end{bmatrix} \right\}. \tag{2.22}$$

G_2 is also a semidirect product, $\langle A \rangle \rtimes \langle B \rangle \cong (\mathbb{Z}/p\mathbb{Z}) \rtimes (\mathbb{Z}/(p-1)\mathbb{Z})$, with elements

$$G_2 = \{A^k B^\ell, \ 0 \le k < p, \ 0 \le \ell < p - 1\}$$
$$= \left\{ \begin{bmatrix} 1 & 0 \\ \alpha & \beta \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_p, \ \beta \in \mathbb{F}_p^\times \right\}.$$

G_3 is the dihedral group $\langle A^{-1}BA \rangle \rtimes \langle CB_1 \rangle \cong D_{2(p-1)}$, where $\langle CB_1 \rangle \cong \mathbb{Z}/2\mathbb{Z}$ and $\langle A^{-1}BA \rangle \cong \mathbb{Z}/(p-1)\mathbb{Z}$. Its elements are

$$G_3 = \left\{ (A^{-1}BA)^k = \begin{bmatrix} 1 & 0 \\ t^k - 1 & t^k \end{bmatrix}, \ (CB_1)(A^{-1}BA)^k = \begin{bmatrix} -1 & -1 \\ 1 - t^{-k} & 1 \end{bmatrix} \ \middle| \ 0 \le k < p - 1 \right\}.$$

G_4 is in fact isomorphic to G_3, with $G_4 = \langle B \rangle \rtimes \langle B_1 C \rangle \cong (\mathbb{Z}/(p-1)\mathbb{Z}) \rtimes (\mathbb{Z}/2\mathbb{Z}) \cong D_{2(p-1)}$. In

this case, its elements are

$$G_4 = \left\{ B^k = \begin{bmatrix} 1 & 0 \\ 0 & t^k \end{bmatrix}, \ (B_1 C) B^k = \begin{bmatrix} 0 & t^k \\ 1 & 0 \end{bmatrix} \ \middle| \ 0 \le k < p - 1 \right\}.$$

Now we can compute the terms of the Ingleton inequality (2.7): We have

$$G_{12} = \langle B_1 \rangle, \ G_{13} = \langle C B_1 \rangle, \ G_{14} = \langle B_1 C \rangle, \tag{2.23}$$

all isomorphic to $\mathbb{Z}/2\mathbb{Z}$. Also we have

$$G_{23} = \langle A^{-1} B A \rangle, \ G_{24} = \langle B \rangle, \tag{2.24}$$

both isomorphic to the cyclic group $\mathbb{Z}/(p-1)\mathbb{Z}$. The remaining groups G_{34}, G_{123}, and G_{124} are all trivial. The two sides of the Ingleton inequality then become

$$|G_{12}||G_{13}||G_{14}||G_{23}||G_{24}| = 8(p-1)^2,$$
$$|G_1||G_2||G_{34}||G_{123}||G_{124}| = 6p(p-1).$$

So we see that these groups do indeed violate Ingleton for $p \ge 5$.

2.6.1 Ingleton Violations in $PGL(2, q)$

Without too much difficulty, the Ingleton violation of the previous section can be generalized [70] to produce a violating set of subgroups in any projective linear group $PGL(2, q)$ for q a prime power greater than or equal to 5. Say $q = p^m$ for some prime p. The finite field $\mathbb{F}_q$ is an m-dimensional vector space over the subfield $\mathbb{F}_p$, so we may fix a basis $\{\xi_1, ..., \xi_m\}$. Instead of the matrix A defined before, we now define a *set* of matrices,

$$A_{\xi_i} = \begin{bmatrix} 1 & 0 \\ \xi_i & 1 \end{bmatrix}, \ i = 1, ..., m.$$

We may assume that $\xi_1 = 1$, in which case A_1 is identical to our matrix A from before. For any $\alpha \in F_q$, we may express α as a linear combination of the ξ_i over $\mathbb{F}_p$, say $\alpha = k_1 \xi_1 + ... + k_m \xi_m$, where each k_i is an integer between 0 and $p - 1$ corresponding to an element in $\mathbb{Z}/p\mathbb{Z} \cong (\mathbb{F}_p, +)$, the set $\mathbb{F}_p$

under addition. It is easy to verify that

$$\begin{bmatrix} 1 & 0 \\ \alpha & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \xi_1 & 1 \end{bmatrix}^{k_1} \cdot ... \cdot \begin{bmatrix} 1 & 0 \\ \xi_m & 1 \end{bmatrix}^{k_m}.$$

We now can define the subgroup

$$G_A := \langle A_{\xi_1}, ..., A_{\xi_m} \rangle = \left\{ \begin{bmatrix} 1 & 0 \\ \alpha & 1 \end{bmatrix}, \ \alpha \in \mathbb{F}_q \right\}.$$

Note that this is actually a direct product $G_A \cong \langle A_{\xi_1} \rangle \times ... \times \langle A_{\xi_m} \rangle \cong (\mathbb{Z}/p\mathbb{Z})^m \cong (\mathbb{F}_q, +)$. The matrix B from the $PGL(2, p)$ case essentially remains the same, except now we take t to be a primitive element of $\mathbb{F}_q$ and define $B = \begin{bmatrix} 1 & 0 \\ 0 & t \end{bmatrix}$. The matrix C from the previous section is exactly the same: $C = \begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix}$. When $p \neq 2$, the matrix B_1 can again be defined exactly as in (2.21). If $p = 2$, we will simply define

$$B_1 := A_1 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \equiv \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix},$$

so that it has the same matrix form as before. One can verify (as before) that the matrices A_{ξ_i}, B, and C generate the elementary matrices from (2.20), and hence generate $PGL(2, q)$.

Our new subgroups of $PGL(2, q)$ can now be explicitly written as

$$G_1 = \langle B_1, C \rangle,$$
$$G_2 = \langle G_A, B \rangle,$$
$$G_3 = \langle CB_1, A_1^{-1}BA_1 \rangle,$$
$$G_4 = \langle B_1C, B \rangle.$$

It is easy to draw a parallel with the corresponding subgroups of the last section.

It remains true that $G_1 = \langle C \rangle \rtimes \langle B_1 \rangle \cong D_6$ of size 6, consisting of the same matrices as in (2.22).

It is easy to verify that G_2 is again the subgroup of lower triangular matrices in $PGL(2, q)$, and that the subgroup G_A is normal in G_2. Since G_A has trivial intersection with $\langle B \rangle$, the subgroup of diagonal matrices, and we in fact have $G_2 = G_A \rtimes \langle B \rangle \cong (\mathbb{Z}/p\mathbb{Z})^m \rtimes (\mathbb{Z}/(q-1)\mathbb{Z})$ of size $p^m(p^m - 1) = q(q-1)$.

G_3 and G_4 have essentially the same structure and matrices as before:

$$G_3 = \left\{ (A_1^{-1}BA_1)^k = \begin{bmatrix} 1 & 0 \\ t^k - 1 & t^k \end{bmatrix}, \ (CB_1)(A_1^{-1}BA_1)^k = \begin{bmatrix} -1 & -1 \\ 1 - t^{-k} & 1 \end{bmatrix} \middle| 0 \le k < q - 1 \right\}$$

$$= \langle A_1^{-1}BA_1 \rangle \rtimes \langle CB_1 \rangle$$

$$\cong D_{2(q-1)},$$

$$G_4 = \left\{ B^k = \begin{bmatrix} 1 & 0 \\ 0 & t^k \end{bmatrix}, \ (B_1C)B^k = \begin{bmatrix} 0 & t^k \\ 1 & 0 \end{bmatrix} \middle| 0 \le k < q - 1 \right\}$$

$$= \langle B \rangle \rtimes \langle B_1 C \rangle$$

$$\cong D_{2(q-1)}.$$

Both G_3 and G_4 have size $2(q-1)$.

Examining the other terms of the Ingleton inequality, the intersections G_{12}, G_{13} and G_{14} take the same forms as in (2.23), and all have size 2. The groups G_{23} and G_{24} also take the same forms as before (from equation (2.24)), but now they are isomorphic to the $\mathbb{Z}/(q-1)\mathbb{Z}$, with size $q-1$. As before, the intersections G_{34}, G_{123}, and G_{124} are trivial. The two sides of the Ingleton inequality (2.7) now become

$$|G_{12}||G_{13}||G_{14}||G_{23}||G_{24}| = 8(q-1)^2,$$

$$|G_1||G_2||G_{34}||G_{123}||G_{124}| = 6q(q-1),$$

and again we have a violation whenever $q \ge 5$.

Our next task will be to extend this example to a broader class of groups, and to explore the structural reason that these groups violate the Ingleton inequality.

2.7   Ingleton Violations in $GL(2,q)$

Since $PGL(2,q)$ is a quotient of the general linear group $GL(2,q)$, it is not surprising that we would find Ingleton violations in this group as well. In fact, it is a simple exercise to show the following simple result:

Lemma 2. *Let G be a group with $N \trianglelefteq G$. Let $H = G/N$ be the quotient group, and H_1, H_2, H_3, and H_4 be subgroups of H with preimages $G_i := \{g \in G \ : \ gN \in H_i\}$. If the H_i violate the Ingleton inequality (2.7), then so do the preimages G_i.*

Proof. It is apparent that each preimage G_i is a subgroup of G containing N. For any subset $\alpha \subseteq \{1, 2, 3, 4\}$, we have from the Lattice Isomorphism Theorem that

$$G_\alpha/N := \left(\bigcap_{i\in\alpha} G_i\right)/N = \bigcap_{i\in\alpha}(G_i/N) = \bigcap_{i\in\alpha} H_i =: H_\alpha.$$

Thus $|G_\alpha| = |H_\alpha||N|$. Since (2.7) has the same number of terms on both sides of the inequality, it is now clear that the H_i produce an Ingleton violation if and only if the G_i do as well. $\square$

We obtain the projective linear group from the quotient of $GL(2, q)$ by its center $Z(GL(2, q))$, which is the group of scalar matrices generated by $tI = \begin{bmatrix} t & 0 \\ 0 & t \end{bmatrix}$, where $t \in \mathbb{F}_q$ is a primitive element of order $q-1$. Thus, we can obtain an Ingleton-violating set of subgroups from Lemma 2 by appending tI to the list of generators for each of the subgroups in section 2.6.1 (taking the original generators now to be matrices in $GL(2, q)$ rather than $PGL(2, q)$). In fact, a computer search in $GL(2, 5)$ produces 15 sets of Ingleton-violating subgroups, up to subscript symmetries in the Ingleton inequality (for example, swapping G_1 and G_2) and conjugations of all four groups (that is, performing a change of basis on $GL(2, 5)$ to transform all the G_i) [69, 70]. These sets of Ingleton violations generalize to $GL(2, q)$ for certain values of q, which we will explore in some detail now.

2.7.1 Instance 1: The Preimage Subgroups

To obtain the preimage subgroups predicted by Lemma 2, we consider the generators of the Ingleton-violating subgroups from Section 2.6.1 as matrices in $GL(2, q)$, and add on the generator tI. For example, the first subgroup becomes $G_1 = \langle tI, B_1, C \rangle$, where $B_1 = \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix} \in GL(2, q)$ and $C = \begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix} \in GL(2, q)$. Note that the subgroup $\langle B_1, C \rangle$ is still isomorphic to D_6 in $GL(2, q)$, and since $\langle tI \rangle$ has trivial intersection with this subgroup, G_1 is actually the direct product

$$G_1 = \langle tI \rangle \times \langle B_1, C \rangle \cong (\mathbb{Z}/(q-1)\mathbb{Z}) \times D_6.$$

The second subgroup takes the form $G_2 = \langle tI, A_{\xi_1}, ..., A_{\xi_m}, B \rangle = \langle tI, G_A, B \rangle$, where again we define $A_{\xi_i} = \begin{bmatrix} 1 & 0 \\ \xi_i & 1 \end{bmatrix}$, $B = \begin{bmatrix} 1 & 0 \\ 0 & t \end{bmatrix}$, and $G_A = \langle A_{\xi_1}, ..., A_{\xi_m} \rangle$, all matrices and subgroups of $GL(2, q)$. Since $\langle tI \rangle$ has trivial intersection with $\langle G_A, B \rangle$, we have

$$G_2 = \langle tI \rangle \times \langle G_A, B \rangle \cong (\mathbb{Z}/(q-1)\mathbb{Z}) \times ((\mathbb{Z}/p\mathbb{Z})^m \rtimes (\mathbb{Z}/(q-1)\mathbb{Z})),$$

which is the group of all lower triangular matrices in $GL(2,q)$ (analogous to the $PGL(2,q)$ scenario). In $G_3 = \langle tI, CB_1, A_1^{-1}BA_1 \rangle$, where again we have taken $\xi_1 = 1$ and $A_{\xi_1} = A_1$, we can easily verify that the entire subgroup $\langle tI, A_1^{-1}BA_1 \rangle \cong \langle tI \rangle \times \langle A_1^{-1}BA_1 \rangle$ is normal in G_3, and has trivial intersection with $\langle CB_1 \rangle$, hence

$$G_3 = (\langle tI \rangle \times \langle A_1^{-1}BA_1 \rangle) \rtimes \langle CB_1 \rangle \cong ((\mathbb{Z}/(q-1)\mathbb{Z} \times \mathbb{Z}/(q-1)\mathbb{Z}) \rtimes \mathbb{Z}.2\mathbb{Z},$$

which has elements $\left\{ t^k \begin{bmatrix} 1 & 0 \\ t^\ell - 1 & t^\ell \end{bmatrix}, \; t^{k+\ell} \begin{bmatrix} -1 & -1 \\ 1 - t^{-\ell} & 1 \end{bmatrix} \;\middle|\; k, \ell \in [q-1] \right\}.$

Finally, the group $G_4 = \langle tI, B_1C, B \rangle$ contains the subgroup $\langle tI, B \rangle \cong \langle tI \rangle \times \langle B \rangle$, which is the subgroup of all diagonal matrices $\left\{ \begin{bmatrix} \alpha & 0 \\ 0 & \beta \end{bmatrix} \;\middle|\; \alpha, \beta \in \mathbb{F}_q^\times \right\}$, and this subgroup is normal in G_4 and intersects the subgroup $\langle B_1C \rangle$ trivially. Thus,

$$G_4 = (\langle tI \rangle \times \langle B \rangle) \rtimes \langle B_1C \rangle \cong ((\mathbb{Z}/(q-1)\mathbb{Z} \times \mathbb{Z}/(q-1)\mathbb{Z}) \rtimes \mathbb{Z}/2\mathbb{Z}$$

$$= \left\{ \begin{bmatrix} \alpha & 0 \\ 0 & \beta \end{bmatrix}, \; \begin{bmatrix} 0 & \alpha \\ \beta & 0 \end{bmatrix} \;\middle|\; \alpha, \beta \in \mathbb{F}_q^\times \right\},$$

which is the group of all diagonal and antidiagonal matrices (as in the $PGL(2,q)$ case).

Computing the subgroup intersections in the Ingleton inequality, we have $G_{12} = \langle tI \rangle \times \langle B_1 \rangle$, $G_{13} = \langle tI \rangle \times \langle CB_1 \rangle$, and $G_{14} = \langle tI \rangle \times \langle B_1C \rangle$, all isomorphic to $(\mathbb{Z}/(q-1)\mathbb{Z}) \times \mathbb{Z}/2\mathbb{Z})$. Also $G_{23} = \langle tI \rangle \times \langle A_1^{-1}BA_1 \rangle$ and $G_{24} = \langle tI \rangle \times \langle B \rangle$, both isomorphic to $(\mathbb{Z}/(q-1)\mathbb{Z}) \times (\mathbb{Z}/(q-1)\mathbb{Z})$. Finally, $G_{34} = G_{123} = G_{124} = \langle tI \rangle \cong \mathbb{Z}/(q-1)\mathbb{Z}$, leading to the sides of the Ingleton inequality (2.7) taking the forms

$$|G_{12}||G_{13}||G_{14}||G_{23}||G_{24}| = 8(q-1)^7,$$

$$|G_1||G_2||G_{34}||G_{123}||G_{124}| = 6q(q-1)^6,$$

whereby the inequality is again violated when $q \geq 5$. Note that the sizes of these intersections, as well as the final form of the sides of the Ingleton inequality, are aptly predicted by the proof of Lemma 2.

Next, we will discuss the remaining Ingleton-violating instances. We will divide them into groups, the first two of which can be obtained by respectively tweaking G_1 and G_2 in the preimage subgroups.

2.7.2 Variants of the Preimage Subgroups with Different G_1

The first class of Ingleton violating sets of subgroups maintains the forms of G_2, G_3, and G_4 from Section 2.7.1, but changes G_1. In each of these instances, G_1 will now be a *subgroup* of the original

group $\langle tI \rangle \times \langle B_1, C \rangle$. We will briefly describe how this changes the intersections of the G_i, but we mention upfront that each of these variants produces Ingleton violations for $q \geq 5$, and when F_q has characteristic $p = 2$, some of these instances will overlap.

1. $G_1 = \langle B_1, C \rangle \cong D_6$:

   $G_{12} = \langle B_1 \rangle$, $G_{13} = \langle CB_1 \rangle$, and $G_{14} = \langle B_1 C \rangle$ are isomorphic to $\mathbb{Z}/2\mathbb{Z}$; $G_{123} = G_{124} = 1$.

2. $G_1 = \langle -I \rangle \times \langle B_1, C \rangle \cong (\mathbb{Z}/2\mathbb{Z}) \times D_6 \cong D_{12}$, $(p \neq 2)$:

   $G_{12} = \cong \langle -I \rangle \times \langle B_1 \rangle$, $G_{13} = \langle -I \rangle \times \langle CB_1 \rangle$, and $G_{14} = \langle -I \rangle \times \langle B_1 C \rangle$, all isomorphic to $(\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})$. Now, $G_{123} = G_{124} = \langle -I \rangle \cong \mathbb{Z}/2\mathbb{Z}$.

3. $G_1 = \langle -B_1, C \rangle \cong D_6$, $(p \neq 2)$:

   $G_{12} = \langle -B_1 \rangle$, $G_{13} = \langle -CB_1 \rangle$, and $G_{14} = \langle -B_1 C \rangle$ are all isomorphic to $\mathbb{Z}/2\mathbb{Z}$. G_{123} and G_{124} are trivial.

4. $G_1 = \langle C, tB_1 \rangle = \langle C \rangle \rtimes \langle tB_1 \rangle \cong (\mathbb{Z}/3\mathbb{Z}) \rtimes (\mathbb{Z}/(q-1)\mathbb{Z})$, $(p \neq 2)$:

   Note that the requirement that p be odd comes from the fact that if q is even then since B_1 is an element of order two, then $(tB_1)^q = tI$, and this instance collapses to the original preimage subgroup in which $G_1 = \langle tI, B_1, C \rangle$.

   When $p \neq 2$, the intersection subgroups now become $G_{12} = \langle tB_1 \rangle$, $G_{13} = \langle tCB_1 \rangle$, $G_{14} = \langle tB_1 C \rangle$ (all isomorphic to $\mathbb{Z}/(q-1)\mathbb{Z}$), and $G_{123} = G_{124} = \langle t^2 I \rangle \cong \mathbb{Z}/ \left(\frac{q-1}{2} \right) \mathbb{Z}$.

2.7.3 Variants of the Preimage Subgroups with Different G_2

In the sets of Ingleton-violating subgroups in this section, G_1, G_3, and G_4 take the same forms as in Section 2.7.1. G_2 will now be a *proper* subgroup of $\langle tI \rangle \times \langle G_A, B \rangle$. Several of these cases are equivalent when $p = 2$, and we will point these out. While most of these sets will violate Ingleton for all $q \geq 5$, several of the cases will additionally require that $\frac{q-1}{2}$ be even. We will address these on a case by case basis.

To facilitate our discussion, we will define the matrices

$$B' = \begin{bmatrix} -1 & 0 \\ 0 & t \end{bmatrix}, \quad P = \begin{bmatrix} t & 0 \\ 0 & 1 \end{bmatrix}, \quad P' = \begin{bmatrix} t & 0 \\ 0 & -1 \end{bmatrix}.$$

Note that $P = tB^{-1}$, which has order $q - 1$. B' and P' are equal to B and P, respectively, when q is even $(p = 2)$. When q is odd, then $t^{\frac{q-1}{2}} = -1$, and we see that $B' = t^{\frac{q-1}{2}} B^{\frac{q+1}{2}}$ and $P' = t^{\frac{q-1}{2}} P^{\frac{q+1}{2}}$. Both B' and P' are elements of order $q - 1$.

1. $G_2 = \langle G_A, B \rangle = \langle G_A \rangle \rtimes \langle B \rangle \cong (\mathbb{Z}/p\mathbb{Z})^m) \rtimes (\mathbb{Z}/(q-1)\mathbb{Z})$:

   Here, G_2 is actually the group of matrices taking the form

   $$\langle G_A, B \rangle = \left\{ \begin{bmatrix} 1 & 0 \\ \alpha & \beta \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q, \ \beta \in \mathbb{F}_q^\times \right\}.$$

   Now, $G_{12} = \langle B_1 \rangle \cong \mathbb{Z}/2\mathbb{Z}$. $G_{23} = \langle A_1^{-1} B A_1 \rangle$ and $G_{24} = \langle B \rangle$ are both isomorphic to $\mathbb{Z}/(q-1)\mathbb{Z}$. $G_{123} = G_{124} = 1$. In this case, the Ingleton inequality is violated for all prime powers $q \geq 5$.

2. $G_2 = \langle G_A, P \rangle = \langle G_A \rangle \rtimes \langle P \rangle \cong (\mathbb{Z}/p\mathbb{Z})^m) \rtimes (\mathbb{Z}/(q-1)\mathbb{Z})$:

   In this case, G_2 is the group of lower triangular matrices in $GL(2, q)$ with a '1' in the lower-right corner:

   $$\langle G_A, P \rangle = \left\{ \begin{bmatrix} \beta & 0 \\ \alpha & 1 \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q, \ \beta \in \mathbb{F}_q^\times \right\}.$$

   In this case, $G_{12} = \langle -B_1 \rangle \cong \mathbb{Z}/2\mathbb{Z}$. $G_{23} = \langle t^{-1} A_1^{-1} B A_1 \rangle$ and $G_{24} = \langle t B^{-1} \rangle$ are both isomorphic to $\mathbb{Z}/(q-1)\mathbb{Z}$. G_{123} and G_{124} are both 1, the trivial group. These subgroups violate the Ingleton inequality for all finite field sizes $q \geq 5$.

3. $G_2 = \langle G_A, B' \rangle = \langle G_A \rangle \rtimes \langle B' \rangle \cong (\mathbb{Z}/p\mathbb{Z})^m \rtimes (\mathbb{Z}/(q-1)\mathbb{Z}), \ (p \neq 2)$:

   It is simple to verify that

   $$\langle G_A, B' \rangle = \left\{ \begin{bmatrix} (-1)^k & 0 \\ \alpha & t^k \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q, \ k \in \{0, ..., q-2\} \right\}.$$

   Note that $t^{\frac{q-1}{2}} = -1$, so if $\frac{q-1}{2}$ is even we can see that G_2 contains the matrices $\left\{ \begin{bmatrix} 1 & 0 \\ \alpha & -1 \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q \right\}$.

   If $\frac{q-1}{2}$ is odd, G_2 instead contains the matrices $\left\{ \begin{bmatrix} -1 & 0 \\ \alpha & -1 \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q \right\}$. This gives us the intersections

   $$G_{12} = \begin{cases} \langle B_1 \rangle \cong \mathbb{Z}/2\mathbb{Z} & \text{if } \frac{q-1}{2} \text{ is even} \\ \langle -I \rangle \cong \mathbb{Z}/2\mathbb{Z} & \text{otherwise} \end{cases}$$

   $$G_{123} = G_{124} = \begin{cases} 1 & \text{if } \frac{q-1}{2} \text{ is even} \\ \langle -I \rangle \cong \mathbb{Z}/2\mathbb{Z} & \text{otherwise} \end{cases}.$$

   In either case we have $G_{23} = \langle -(A_1^{-1} B A_1)^{\frac{q+1}{2}} \rangle$ and $G_{24} = \langle B' \rangle$, both isomorphic to the cyclic

group $\mathbb{Z}/(q-1)\mathbb{Z}$. When $q = p^m$ is odd, greater than or equal to 5, this case only violates the Ingleton inequality when $\frac{q-1}{2}$ is even.

4. $G_2 = \langle G_A, P' \rangle = \langle G_A \rangle \rtimes \langle P' \rangle \cong (\mathbb{Z}/p\mathbb{Z})^m \rtimes (\mathbb{Z}/(q-1)\mathbb{Z})$, $(p \neq 2)$:

In this case, we have

$$\langle G_A, P' \rangle = \left\{ \begin{bmatrix} t^k & 0 \\ \alpha & (-1)^k \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q, \ k \in \{0, ..., q-2\} \right\},$$

which contains the matrices $\left\{ \begin{bmatrix} -1 & 0 \\ \alpha & 1 \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q \right\}$ if $\frac{q-1}{2}$ is even and $\left\{ \begin{bmatrix} -1 & 0 \\ \alpha & -1 \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q \right\}$ if $\frac{q-1}{2}$ is odd. Our intersection subgroups become

$$G_{12} = \begin{cases} \langle -B_1 \rangle \cong \mathbb{Z}/2\mathbb{Z} & \text{if } \frac{q-1}{2} \text{ is even} \\ \langle -I \rangle \cong \mathbb{Z}/2\mathbb{Z} & \text{otherwise} \end{cases}$$

$$G_{123} = G_{124} = \begin{cases} 1 & \text{if } \frac{q-1}{2} \text{ is even} \\ \langle -I \rangle \cong \mathbb{Z}/2\mathbb{Z} & \text{otherwise} \end{cases}.$$

$G_{23} = \langle t(A_1^{-1}BA_1)^{\frac{q-3}{2}} \rangle$ and $G_{24} = \langle P' \rangle$, both isomorphic to $\mathbb{Z}/(q-1)\mathbb{Z}$. As in the previous case, when $q \geq 5$ is odd, this case only violates Ingleton when $\frac{q-1}{2}$ is even.

5. $G_2 = \langle -I, G_A, B \rangle = \langle -I \rangle \times \langle G_A, B \rangle = (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/p\mathbb{Z})^m \rtimes (\mathbb{Z}/(q-1)\mathbb{Z})$, $(p \neq 2)$: $G_{12} = \langle -I \rangle \times \langle B_1 \rangle \cong (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})$. $G_{23} = \langle -I \rangle \times \langle A_1^{-1}BA_1 \rangle$ and $G_{24} = \langle -I \rangle \times \langle B \rangle$ are both isomorphic to $(\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/(q-1)\mathbb{Z})$. $G_{123} = G_{124} = \langle -I \rangle \cong \mathbb{Z}/2\mathbb{Z}$. This produces an Ingleton violation for any odd prime power $q \geq 5$.

6. $G_2 = \langle -I, G_A, P \rangle = \langle -I \rangle \times \langle G_A, P \rangle = (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/p\mathbb{Z})^m \rtimes (\mathbb{Z}/(q-1)\mathbb{Z})$, $(p \neq 2)$:

$G_{12} = \langle -I \rangle \times \langle B_1 \rangle \cong (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})$. $G_{23} = \langle -I \rangle \times \langle t^{-1}A_1^{-1}BA_1 \rangle$ and $G_{24} = \langle -I \rangle \times \langle P \rangle$, both isomorphic to $(\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/(q-1)\mathbb{Z})$. As in the previous case, $G_{123} = G_{124} = \langle -I \rangle \cong \mathbb{Z}/2\mathbb{Z}$. This violates the Ingleton inequality whenever $q \geq 5$.

2.7.4 The Final Four Ingleton Violations

In the four remaining sets of Ingleton-violating subgroups which occur in $GL(2, 5)$ and generalize to other general linear groups, we always have $G_1 = \langle B_1, C \rangle \cong \langle C \rangle \rtimes \langle B_1 \rangle \cong D_6$. The remaining subgroups are all equal or conjugate to one of $\langle G_A, B \rangle, \langle G_A, B' \rangle, \langle G_A, P \rangle$, or $\langle G_A, P' \rangle$. As we have already described, each of these is a semidirect product of the normal subgroup G_A by the cyclic group

generated by the one of the matrices B, B', P, and P', and is isomorphic to $(\mathbb{Z}/p\mathbb{Z})^m \rtimes (\mathbb{Z}/(q-1)\mathbb{Z})$. The conjugators of the above groups will take the form

$$
E = \begin{bmatrix} -1 & 1 \\ 1 & 0 \end{bmatrix}, \quad Q = \begin{bmatrix} 2 & 1 \\ 1 & 0 \end{bmatrix}, \quad W = \begin{bmatrix} 0 & 1 \\ -1 & 1 \end{bmatrix}.
$$

For a subgroup H and a matrix M, we will use the common notation H^M to describe the conjugated subgroup $M^{-1}HM := \{M^{-1}XM \ : \ X \in H\}$ which is itself a subgroup isomorphic to H. As before, there will be conditions on $q = p^m$ which must be met in order for these sets of subgroups to produce Ingleton violations, but we will address these as we come to them.

1. $G_2 = \langle G_A, B \rangle$, $G_3 = \langle G_A, P \rangle^E$, $G_4 = \langle G_A, P \rangle^Q$, $(p \neq 3)$:

    Note that in this case, the requirement that $p \neq 3$ stems from the fact that otherwise, $E \equiv Q$, and the groups G_3 and G_4 are identical. Thus the Ingleton inequality cannot be violated by Theorem 2. We can also verify that G_3 and G_4 respectively become

$$
\langle G_A, P \rangle^E = \left\{ \begin{bmatrix} 1-\alpha & \alpha \\ 1-\alpha-\beta & \alpha+\beta \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q, \ \beta \in \mathbb{F}_q^\times \right\},
$$

$$
\langle G_A, P \rangle^Q = \left\{ \begin{bmatrix} 1+2\alpha & \alpha \\ 2(\beta-2\alpha-1) & \beta-2\alpha \end{bmatrix} \ \middle| \ \alpha \in \mathbb{F}_q, \ \beta \in \mathbb{F}_q^\times \right\}. \tag{2.25}
$$

    These groups intersect with G_1 as $G_{12} = \langle B_1 \rangle$, $G_{13} = \langle B_1 C \rangle$, and $G_{14} = \langle C B_1 \rangle$, which are isomorphic to $\mathbb{Z}/2\mathbb{Z}$. We also have

$$
G_{23} = \langle P \rangle^E = \left\{ \begin{bmatrix} 1 & 0 \\ 1-t^k & t^k \end{bmatrix} \ \middle| \ k \in \{0, ..., q-2\} \right\}, \tag{2.26}
$$

$$
G_{24} = \langle P \rangle^Q = \left\{ \begin{bmatrix} 1 & 0 \\ 2(t^k-1) & t^k \end{bmatrix} \ \middle| \ k \in \{0, ..., q-2\} \right\}, \tag{2.27}
$$

    both of which are isomorphic to $\mathbb{Z}/(q-1)\mathbb{Z}$. The remaining intersections in the Ingleton inequality (2.7) are trivial. This violates Ingleton for $q \geq 5$.

2. $G_2 = \langle G_A, B' \rangle$, $G_3 = \langle G_A, P' \rangle^E$, $G_4 = \langle G_A, P' \rangle^Q$ $(p \neq 2, 3)$: In this case, we must have $p \neq 3$ for the same reason as in the previous case, and based on the forms of B' and P' we see that unless $p \neq 2$, these groups will be identical to those in the previous case. We have already

described the form of the matrices in G_2, and those in G_3 and G_4 will take the form

$$\langle G_A, P' \rangle^E = \left\{ \begin{bmatrix} (-1)^k - \alpha & \alpha \\ (-1)^k - t^k - \alpha & t^k + \alpha \end{bmatrix} \middle| \alpha \in \mathbb{F}_q, \ k \in \{0, ..., q-2\} \right\},$$

$$\langle G_A, P' \rangle^Q = \left\{ \begin{bmatrix} (-1)^k + 2\alpha & \alpha \\ 2(t^k - 2\alpha - (-1)^k) & t^k - 2\alpha \end{bmatrix} \middle| \alpha \in \mathbb{F}_q, \ k \in \{0, ..., q-2\} \right\}.$$

Since $t^{\frac{q-1}{2}} = -1$, we can see that if $\frac{q-1}{2}$ is even, G_2 will contain the matrices $\left\{ \begin{bmatrix} 1 & 0 \\ \alpha & -1 \end{bmatrix} \right\}$,

while G_3 will contain those of the form $\left\{ \begin{bmatrix} 1-\alpha & \alpha \\ 2-\alpha & \alpha-1 \end{bmatrix} \right\}$ and G_4 will include the matrices

$\left\{ \begin{bmatrix} 1+2\alpha & \alpha \\ -2^2(1+\alpha) & -1-2\alpha \end{bmatrix} \right\}$. On the other hand, when $\frac{q-1}{2}$ is odd, G_2 includes the matrices

$\left\{ \begin{bmatrix} -1 & 0 \\ \alpha & -1 \end{bmatrix} \right\}$, G_3 contains $\left\{ \begin{bmatrix} -1-\alpha & \alpha \\ -\alpha & \alpha-1 \end{bmatrix} \right\}$ and G_4 will contain the set $\left\{ \begin{bmatrix} -1+2\alpha & \alpha \\ -2^2\alpha & -1-2\alpha \end{bmatrix} \right\}$.

Thus, when $\frac{q-1}{2}$ even, we have as in the previous case: $G_{12} = \langle B_1 \rangle$, $G_{13} = \langle B_1 C \rangle$, and $G_{14} = \langle C B_1 \rangle$, all isomorphic to $\mathbb{Z}/2\mathbb{Z}$, and $G_{34} = 1$. When $\frac{q-1}{2}$ is odd, G_{12}, G_{13}, and G_{14} become trivial, and $G_{34} = \langle -I \rangle \cong \mathbb{Z}/2\mathbb{Z}$. In either case,

$$G_{23} = \langle P' \rangle^E = \left\{ \begin{bmatrix} (-1)^k & 0 \\ (-1)^k - t^k & t^k \end{bmatrix} \middle| k \in \{0, ..., q-2\} \right\}, \tag{2.28}$$

$$G_{24} = \langle P' \rangle^Q = \left\{ \begin{bmatrix} (-1)^k & 0 \\ 2(t^k - (-1)^k) & t^k \end{bmatrix} \middle| k \in \{0, ..., q-2\} \right\}, \tag{2.29}$$

both isomorphic to $\mathbb{Z}/(q-1)\mathbb{Z}$, and $G_{123} = G_{124} = 1$. This set of subgroups produces an Ingleton violation for $q \geq 5$ and $\frac{q-1}{2}$ even.

3. $G_2 = \langle G_A, P \rangle^E$, $G_3 = \langle G_A, B \rangle$, $G_4 = \langle G_A, B \rangle^W$ ($p \neq 3$):

   In this case, if $p = 3$ we will have $2 \equiv -1$ in $\mathbb{F}_q$, and we can see that the matrices B_1 and C will actually be elements of $\langle G_A, P \rangle^E$, and hence G_1 will be a subgroup of G_2 and Ingleton cannot be violated by Theorem 2.

   We have described the form of the elements in each of the subgroups G_1, G_2, and G_3, and the

elements of G_4 are actually

$$\langle G_A, B\rangle^W = \left\{ \begin{bmatrix} \beta & \alpha \\ 0 & 1 \end{bmatrix} \ \middle|\ \alpha \in \mathbb{F}_q, \ \beta \in \mathbb{F}_q^\times \right\}$$
$$= \{ X^T \ : \ X \in \langle G_A, P\rangle \}.$$

Now we have $G_{12} = \langle B_1 C\rangle$, $G_{13} = \langle B_1\rangle$, and $G_{14} = \langle CB_1\rangle$, all isomorphic to $\mathbb{Z}/2\mathbb{Z}$. G_{23} is the same as in (2.26), while G_{24} becomes

$$G_{24} = \langle B\rangle^W = \left\{ \begin{bmatrix} t^k & 1 - t^k \\ 0 & 1 \end{bmatrix} \ \middle|\ k \in \{0, ..., q-2\} \right\} \cong \mathbb{Z}/(q-1)\mathbb{Z}. \qquad (2.30)$$

Also, we have $G_{34} = G_{123} = G_{124} = 1$. This instance produces an Ingleton violation for any $q = p^m \geq 5$, provided $p \neq 3$.

4. $G_2 = \langle G_A, P'\rangle^E$, $G_3 = \langle G_A, B'\rangle$, $G_4 = \langle G_A, B'\rangle^W$ $(p \neq 2, 3)$:

Again, we cannot have $p = 2$ because the groups in this case will become equal to those of the previous one, and when $p = 3$ (so that $2 \equiv -1$) we can verify that B_1 and C are elements $\langle G_A, P'\rangle^E$, so that $G_1 \leq G_2$ and Ingleton cannot be satisfied by Theorem 2.

We have described the forms of all the matrices in these subgroups except for those of G_4, which are

$$\langle G_A, B'\rangle^W = \left\{ \begin{bmatrix} t^k & \alpha \\ 0 & (-1)^k \end{bmatrix} \ \middle|\ \alpha \in \mathbb{F}_q, \ k \in \{0, ..., q-2\} \right\}$$
$$= \{ X^T \ : \ X \in \langle G_A, P'\rangle \}.$$

Note that G_2 and G_3 are the same groups (but swapped) from case 2 in this section, in which we discussed the differences in the matrices they contain depending on whether $\frac{q-1}{2}$ is even or odd. Similarly, we see that when $\frac{q-1}{2}$ is even G_4 contains the matrices $\left\{ \begin{bmatrix} -1 & \alpha \\ 0 & 1 \end{bmatrix} \right\}$, and when $\frac{q-1}{2}$ is odd G_4 contains $\left\{ \begin{bmatrix} -1 & \alpha \\ 0 & -1 \end{bmatrix} \right\}$ for $\alpha \in \mathbb{F}_q$. We can thus verify that when $\frac{q-1}{2}$ is even, $G_{12} = \langle B_1 C\rangle$, $G_{13} = \langle B_1\rangle$, and $G_{14} = \langle CB_1\rangle$, all isomorphic to $\mathbb{Z}/2\mathbb{Z}$, and G_{34} is trivial. When $\frac{q-1}{2}$ is odd, we have that $G_{12} = G_{13} = G_{14} = 1$ and $G_{34} = \langle -I\rangle \cong \mathbb{Z}/2\mathbb{Z}$. In either case, G_{123}

and G_{124} are trivial, G_{23} is the same as in (2.28), and G_{24} becomes

$$G_{24} = \langle B' \rangle^W = \left\{ \begin{bmatrix} t^k & (-1)^k - t^k \\ 0 & (-1)^k \end{bmatrix} \ \middle| \ k \in \{0, ..., q-2\} \right\} \cong \mathbb{Z}/(q-1)\mathbb{Z}.$$

This instance produces an Ingleton violation when $q \geq 5$ and $\frac{q-1}{2}$ is even.

2.8 Interpreting the Ingleton Violations Using Group Actions

The projective linear group $PGL(n, q)$ can be interpreted as the set of linear transformations acting on the projective linear space $PG(\mathbb{F}_q^n)$, in which points are defined as the lines in $\mathbb{F}_q^n$. In other words, $PG(\mathbb{F}_q^n)$ is the set $\{\mathbf{v} \in \mathbb{F}_q^n \setminus \mathbf{0}\}$ under the relation $\lambda \mathbf{v} \equiv \mathbf{v}, \forall \lambda \in \mathbb{F}_q, \lambda \neq 0$. The number of points in $PG(\mathbb{F}_q^n)$ is accordingly $\frac{q^n-1}{q-1} = q^{n-1} + q^{n-2} + ... + 1$.

Recall that a *group action* of a group G on a set S is a homomorphism from G to $\text{Perm}(S)$, the permutation group of S. If $g \in G$ and $s \in S$, we will denote by gs the image of s under the permutation associated to g. The set $Gs := \{gs \ : \ g \in G\}$ is called the *orbit* of the element s under G. For a subset $S' \subseteq S$, gS' is the set of images $\{gs \ : \ s \in S'\}$, and the *stabilizer* of S' is the set $\text{Stab}_G(S') = \{g \in G \ : \ gs \in S', \ \forall s \in S'\}$, also denoted $G(S')$. We will use the notation $G_{S'}$ to denote the *pointwise stabilizer* of S', the subgroup of $G(S')$ defined as $G_{S'} := \{g \in G \ : \ gs = s, \ \forall s \in S'\}$. It should be clear that for a single element $s \in S$, we have $G(s) = G_s$.

We say a group action is *transitive* if for any $s_1, s_2 \in S$, there is a $g \in G$ such that $gs_1 = s_2$. We say the action is r-transitive if it is transitive on the set of ordered r-tuples of distinct points in S. That is, if $\mathbf{s} = (s_1, ..., s_r)$ and $\mathbf{s}' = (s_1', ..., s_r')$ are points in S^r where no two s_i (and no two s_i') are equal, then there is a $g \in G$ such that $g\mathbf{s} := (gs_1, ..., gs_r) = \mathbf{s}'$. G is *sharply* r-transitive on S if it is r-transitive and, in addition, only the identity element $1 \in G$ fixes any r points in S.

An important classical result which we will use is called the *orbit stabilizer theorem*:

Theorem 3 (Orbit Stabilizer Theorem). *Let G be a group acting on a set S. For any $s \in S$, the size of the orbit of s under G is $|Gs| = \frac{|G|}{|G_s|}$.*

Proof. Consider the set of left cosets G/G_s. Every element in the coset gG_s maps s to the same element in S, namely gs. Thus there are *at most* $|G/G_s| = \frac{|G|}{|G_s|}$ elements in the orbit of s. Furthermore, for any two left cosets g_1G_s and g_2G_s, if $g_1s = g_2s$, then $g_1^{-1}g_2$ fixes s and hence is an element of G_s, so the cosets are the same. It follows that $|Gs|$ is *exactly* the number of cosets, $\frac{|G|}{|G_s|}$. $\square$

Using this terminology, it is easy to see that $PGL(2, q)$ is 2-transitive on the projective space $PG(\mathbb{F}_q^2)$, but the following is also true:

Lemma 3. *The group action of $PGL(2,q)$ on the projective space $PG(\mathbb{F}_q^2)$ is sharply 3-transitive.*

Proof. Fix two ordered triples (s_1, s_2, s_3) and (s_1', s_2', s_3') in $PG(\mathbb{F}_q^2)$. Since any two distinct points in $PG(\mathbb{F}_q^2)$ must be linearly independent in $\mathbb{F}_q^2$ (because they generate two different lines), we may write $s_3 \equiv s_1 + \alpha s_2$ and $s_3' \equiv s_1' + \beta s_2'$ for some $\alpha, \beta \in \mathbb{F}_q$. (Note that we may assume the coefficients of s_1 and s_1' are nonzero, since otherwise s_3 and s_3' would be equivalent to s_2 and s_2' respectively in the projective space.

Now let $X \in PGL(2,q)$ be such that $Xs_1 = s_1'$ and $Xs_2 = \alpha^{-1}\beta s_2' \equiv s_2'$. Then

$$Xs_3 \equiv Xs_1 + \alpha Xs_2 \equiv s_1' + \beta s_2' \equiv s_3,$$

and we have 3-transitivity.

Now, if $s_i = s_i'$ for each $i = 1, 2, 3$, so that X fixes the triple (s_1, s_2, s_3). Then considering X as a matrix in $GL(2,q)$ and the points s_i as typical points in the vector space $\mathbb{F}_q^2$, we must have $Xs_1 = \lambda_1 s_1$ and $Xs_2 = \lambda_2 s_2$ for some $\lambda_1, \lambda_2 \in \mathbb{F}_q$. Writing $s_3 = s_1 + \alpha s_2$, we also must have

$$\lambda_1 s_1 + \alpha\lambda_2 s_2 = Xs_3 = \lambda_3 s_3 = \lambda_3 s_1 + \alpha\lambda_3 s_2,$$

so both λ_1 and λ_2 must be equal to λ_3, hence X is a scalar transformation, $X = \lambda_3 I$, which is equivalent to the identity transformation in $PGL(2,q)$. This gives us sharpness. $\qquad\square$

Let $G = PGL(2,q)$. Now consider the standard basis for $\mathbb{F}_q^2$, $\left\{ e_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, e_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$, each element giving rise to a point in the projective space: $x_1 = \{\lambda e_1\}$ and $x_2 = \{\lambda e_2\}$. Recall our Ingleton-violating subgroups in $PGL(2,q)$ from Section 2.6.1:

$$G_1 = \left\{ I, \begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix}, \begin{bmatrix} -1 & -1 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} -1 & -1 \\ 0 & 1 \end{bmatrix} \right\}$$

$$G_2 = \left\{ \begin{bmatrix} 1 & 0 \\ \alpha & \beta \end{bmatrix} \,\middle|\, \alpha \in \mathbb{F}_q, \ \beta \in \mathbb{F}_q^\times \right\}$$

$$G_3 = \left\{ \begin{bmatrix} 1 & 0 \\ \beta - 1 & \beta \end{bmatrix}, \begin{bmatrix} -1 & -1 \\ 1 - \beta & 1 \end{bmatrix} \,\middle|\, \beta \in \mathbb{F}_q^\times \right\}$$

$$G_4 = \left\{ \begin{bmatrix} 1 & 0 \\ 0 & \beta \end{bmatrix}, \begin{bmatrix} 0 & \beta \\ 1 & 0 \end{bmatrix} \,\middle|\, \beta \in \mathbb{F}_q^\times \right\}$$

The group G_2 is the set of lower-triangular matrices in $PGL(2,q)$, which is in fact the stabilizer of

the line x_2, so we may write $G_2 = G_{x_2}$. G_3 and G_4 are both isomorphic to $D_{2(q-1)}$, with normal subgroups G_{23} and G_{24} (both isomorphic to $\mathbb{Z}/(q-1)\mathbb{Z}$ which intersect trivially. This means that we have the quotient groups G_3/G_{23} and G_4/G_{24} isomorphic to $\mathbb{Z}/2\mathbb{Z}$. If we let $g \in G_3$ be a generator of G_3/G_{23}, and let $x_3 := gx_2$, then we claim G_3 is the stabilizer of the set $\{x_2, x_3\}$. Indeed, this set is fixed by both g and G_{23}, and the size of $\text{Stab}_G(\{x_2, x_3\})$ can be computed by the orbit-stabilizer theorem (considering that G acts transitively on the *pairs* of points of the projective plane) to be

$$|\text{Stab}_G(\{x_2, x_3\})| = \frac{|G|}{|G\{x_2, x_3\}|} = \frac{|PGL(2,q)|}{\#\{\{x_i, x_j\} \ : \ x_i, x_j \in PG(\mathbb{F}_q^2), \ x_i \neq x_j\}}$$
$$= \frac{(q^2-1)(q^2-q)/(q-1)}{\binom{q+1}{2}}$$
$$= 2(q-1).$$

Thus, by size considerations G_3 is the entire stabilizer group of $\{x_2, x_3\}$. We can easily compute x_3 to be $\begin{pmatrix} -1 \\ 1 \end{pmatrix}$. A similar argument shows that G_4 is the stabilizer of the set $\{x_2, x_4\}$, where $x_4 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$.

Finally, we claim $G_1 \cong D_6$ is the stabilizer of the set $\{x_2, x_3, x_4\}$. Indeed, we can easily verify that G_1 stabilizes this set, and since $PGL(2, q)$ is sharply 3-transitive by Lemma 3, we know that $\text{Stab}_G(\{x_2, x_3, x_4\}) \cong S_3$. But this is isomorphic to D_6, so G_1 is the full stabilizer.

It is not too difficult to see that the choice of the points x_2, x_3, and x_4 was arbitrary, provided that they are distinct. Indeed, using the relation

$$|H_1 H_2| = \frac{|H_1||H_2|}{|H_{12}|}, \tag{2.31}$$

which is true for any subgroups H_1 and H_2, we can rewrite the Ingleton inequality of (2.7) in the form

$$\frac{|G_1||G_2|}{|G_{12}|} \geq \frac{|G_{13}G_{23}||G_{14}G_{24}|}{|G_{34}|}. \tag{2.32}$$

Then, fixing points α, β and γ, and setting $G_1 = G(\{\alpha, \beta, \gamma\})$, $G_2 = G_\alpha$, $G_3 = G(\{\alpha, \beta\})$ and $G_4 = G(\{\alpha, \gamma\})$, we actually have $G_{13}G_{23} = G_3$ and $G_{14}G_{24} = G_4$. To see this, note that G_{13} is the subgroup of G_1 which fixes γ and permutes $\{\alpha, \beta\}$, and G_{23} is simply the point wise stabilizer of α and β, i.e. $G_{\alpha, \beta}$. Now since G is sharply 3-transitive, any $g \in G$ is uniquely determined by its values on three distinct points of S. Thus if $h \in G_3$ and we consider the three points $\alpha' := h^{-1}\alpha$, $\beta' := h^{-1}\beta$, and $\gamma' := h^{-1}\gamma$, there is a unique g which fixes each of α' and β' while satisfying $g\gamma' = \gamma$. By our choice of h, we know that $\{\alpha', \beta'\} = \{\alpha, \beta\}$, so $g \in G_{23}$. Sharp 3-transitivity also gives us a group element g' which fixes γ and satisfies $g'\alpha' = \alpha$ and $g'\beta' = \beta$. This g' is in G_{13}. The composition $g'g$ satisfies $g'g\alpha' = \alpha$, $g'g\beta' = \beta$ and $g'g\gamma' = \gamma$, so by sharp 3-transitivity we must have $g'g = h$. We can

now deduce that $G_3 = G_{13}G_{23}$, and a similar argument shows that $G_4 = G_{14}G_{24}$. Note that we only needed sharp 3-transitivity to prove this.

We can now rewrite (2.32) as

$$\frac{|G_1||G_2|}{|G_{12}|} \geq \frac{|G_3||G_4|}{|G_{34}|}. \tag{2.33}$$

We can use the orbit stabilizer theorem to compute the quantities $|G_1|$, $|G_2|$, and $|G_3|$ as we did in the above calculation of $|\text{Stab}_G(\{x_2, x_3\})|$, using only knowledge of the size of G, the fact that G is 3-transitive, and how each of the subgroups acts on the points α, β and γ. We find that

$$|G_1| = 6 \tag{2.34}$$

$$|G_2| = q(q - 1) \tag{2.35}$$

$$|G_3| = |G_4| = 2(q - 1). \tag{2.36}$$

We can compute $|G_{12}|$ using the orbit stabilizer theorem by considering G_{12} as the subgroup of G_γ which fixes the set $\{\alpha, \beta\}$:

$$\begin{aligned}
|\text{Stab}_{G_\gamma}(\{\alpha, \beta\})| &= \frac{|G_\gamma|}{|G_\gamma\{\alpha, \beta\}|} \\
&= \frac{q(q-1)}{\binom{q}{2}} \\
&= 2.
\end{aligned}$$

Finally, we compute $|G_{34}|$ by noting that G_{34} is the pointwise stabilizer of the set $\{\alpha, \beta, \gamma\}$, i.e. $G_{\alpha,\beta,\gamma}$. We consider the action of this group on the ordered triples of distinct points in the projective space, and again apply the orbit stabilizer theorem:

$$\begin{aligned}
|G_{\alpha,\beta,\gamma}| &= \frac{|G|}{\#\{(x_1, x_2, x_3) \in S \times S \times S \ : \ x_i \neq x_j, \ \forall i \neq j\}} \\
&= \frac{(q^2 - 1)(q^2 - q)/(q - 1)}{(q + 1)q(q - 1)} \\
&= 1.
\end{aligned}$$

Plugging these quantities into (2.33), we see that as expected the Ingleton inequality is violated. This example is not only gives rise to a whole class of Ingleton-violating subgroups in $PGL(2, q)$, but it gives a structural interpretation for the fundamental reason *why* they violate Ingleton. More importantly, though, is the fact that it reveals that we really only needed G to be a sharply 3-transitive group of

a particular size.

2.8.1 Ingleton Violations in More General 2-Transitive Groups

It turns out that even requiring G to be strictly 3-transitive is more than we need demand. Instead, let us just require that G act 2-transitively on a set S of size at least 3. Suppose further that for some subset $\{\alpha, \beta, \gamma\} \subseteq S$, G acts as the symmetric group S_3. That is, for every permutation σ of the elements $\{\alpha, \beta, \gamma\}$, there is some $g \in G$ such that $g\alpha = \sigma\alpha$, $g\beta = \sigma\beta$, and $g\gamma = \sigma\gamma$. This condition is clearly if G acts 3-transitively on S, as in the case of $PGL(2,q)$ acting on $PG(\mathbb{F}_q^2)$.

Now, again define the subgroups:

$$
\begin{aligned}
G_1 &= G(\{\alpha, \beta, \gamma\}), \\
G_2 &= G_\alpha, \\
G_3 &= G(\{\alpha, \beta\}), \\
G_4 &= G(\{\alpha, \gamma\}).
\end{aligned}
\tag{2.37}
$$

Lemma 4. *Let the group G act 2-transitively on a set S, and act as the symmetric group on a subset $\{\alpha, \beta, \gamma\} \subseteq S$. Then if G_1, G_2, G_3, and G_4 are defined as in (2.37), we have $G_3 = G_{13}G_{23}$ and $G_4 = G_{14}G_{24}$.*

Proof. We prove the equality for G_3, and the proof for G_4 is analogous. Since the group product $G_{13}G_{23}$ is a subset of G_3, it suffices show that $|G_{13}G_{23}| = |G_3|$. But from (2.31) we have that $|G_{13}G_{23}| = |G_{13}||G_{23}|/|G_{123}|$. We claim that

$$
\frac{|G_{13}|}{|G_{123}|} = \frac{|G_3|}{|G_{23}|}.
$$

Indeed, G_{123} is the subgroup of G_{13} which fixes α, and likewise G_{23} is the stabilizer of α in G_3. Thus from the orbit-stabilizer theorem, $\frac{|G_{13}|}{|G_{123}|}$ is the size of the orbit of α under the action of G_{13}, and $\frac{|G_3|}{|G_{23}|}$ is the size of the orbit of α under the group G_3. But by hypothesis, both of these orbits are equal to the set $\{\alpha, \beta\}$, so we are done. $\qquad\square$

This lemma shows that the Ingleton inequality again takes the form of (2.33). Now computing the remaining terms, we note that G_{12} is the subgroup of G_1 which fixes α, thus by the orbit stabilizer theorem, $\frac{|G_1|}{|G_{12}|}$ is equal to the size of the orbit of α under G_1 which is the set $\{\alpha, \beta, \gamma\}$, and hence $\frac{|G_1|}{|G_{12}|} = 3$. Also by the orbit stabilizer theorem and transitivity, $|G_2| = |G|/|G\alpha| = |G|/|S|$.

Let us define $\tau := |G_3|/|G_{34}|$. Since G_{34} is the subgroup of G_3 which fixes α, β and γ (which can equivalently be interpreted as either $G_3(\{\alpha, \gamma\})$, $(G_3)_{\alpha,\gamma}$, or $(G_3)_{\beta,\gamma}$), we see that τ depends on how G

acts on these three elements. By the orbit stabilizer theorem, τ is equal to the size of the orbit of the ordered triple (α, β, γ) under G_3. This is at least 2, since G_3 acts as the symmetric group on the set $\{\alpha, \beta\}$ by assumption. On the other hand, for each permutation of $\{\alpha, \beta\}$, G_3 can potentially map γ to any of the elements $S \setminus \{\alpha, \beta\}$, so we have $2 \leq \tau \leq 2(|S| - 2)$. τ achieves the lower bound when, for example, G_3 is a subgroup of G_γ. It achieves the upper bound in the case where G is 3-transitive, as in our previous example with the projective linear group. To simplify notation, we will let $\tau' = \tau/2$, so that $\frac{|G_3|}{|G_{34}|} = 2\tau'$ and $1 \leq \tau' \leq |S| - 2$.

Finally, by 2-transitivity and the orbit stabilizer theorem, $|G_4|$ is equal to $|G|$ divided by the number of pairs of elements in S. Thus, $|G_4| = \frac{|G|}{\binom{|S|}{2}} = \frac{2|G|}{|S|(|S|-1)}$.

Now, examining the Ingleton inequality (2.7), we see that Ingleton is violated if $\frac{|G_{12}||G_{13}||G_{14}||G_{23}||G_{24}|}{|G_1||G_2||G_{34}||G_{123}||G_{124}|} > 1$. For a set of subgroups $\mu = (G_1, G_2, G_3, G_4)$, we define the *Ingleton ratio* to be the quantity

$$r(\mu) = \frac{|G_{12}||G_{13}||G_{14}||G_{23}||G_{24}|}{|G_1||G_2||G_{34}||G_{123}||G_{124}|}. \tag{2.38}$$

In the case of the groups in (2.37), the Ingleton ratio becomes

$$r(\mu) = \frac{|G_{12}||G_3||G_4|}{|G_1||G_2||G_{34}|} = \frac{4\tau'}{3(|S| - 1)}. \tag{2.39}$$

If G is chosen so that τ' is close to $|S| - 2$, then as $|S|$ becomes large the Ingleton ratio approaches $\frac{4}{3}$, producing an Ingleton violation.

Chapter 3

Group Frames with Few Distinct Inner Products and Low Coherence

A frame is the following generalization for the basis of a vector space:

Definition 1. *Let $\mathcal{V}$ be a vector space equipped with an inner product $\langle \cdot, \cdot \rangle$ (or more specifically, a separable Hilbert space). A set of elements $\{f_k\}_{k \in \mathcal{I}}$, where $\mathcal{I}$ is a countable index set, is a frame for $\mathcal{V}$ if there exist positive constants A and B such that*

$$A||f||_2^2 \leq \sum_{k \in \mathcal{I}} |\langle f, f_k \rangle|^2 \leq B||f||_2^2, \tag{3.1}$$

for all $f \in \mathcal{V}$. A frame is called tight if $A = B$ in this definition, and unit norm if $||f_k||_2 = 1, \forall k \in \mathcal{I}$.

Most often we will consider our frame vectors to be the columns $\{\mathbf{m}_i\}_{i=1}^n$ of a matrix $\mathbf{M} = [\mathbf{m}_1, \mathbf{m}_2, \ldots, \mathbf{m}_n] \in \mathbb{C}^{m \times n}$. We will speak of the *coherence of* $\mathbf{M}$ to be the coherence of the frame $\{\mathbf{m}_i\}$. If $\mathbf{f} \in \mathbb{C}^{m \times 1}$ is any vector, the sum in (3.1) takes the form $\sum_{k=1}^n |\langle \mathbf{f}, \mathbf{m}_k \rangle|^2 = \mathbf{f}^* \mathbf{M} \mathbf{M}^* \mathbf{f}$. By examining the singular value decomposition of $\mathbf{M}\mathbf{M}^*$, we can see that the frame is tight if and only if $\mathbf{M}\mathbf{M}^* = \lambda \mathbf{I}_m$ where $\mathbf{I}_m$ is the $m \times m$ identity matrix and $\lambda = A = B$ in (3.1). If the columns $\{\mathbf{m}_i\}$ form a unit-norm tight frame, then we have the relation

$$m\lambda = \text{Tr}(\mathbf{M}\mathbf{M}^*) = \text{Tr}(\mathbf{M}^*\mathbf{M}) = \sum_{i=1}^n ||\mathbf{m}_i||_2^2 = n, \tag{3.2}$$

from which we see that $\lambda = \frac{n}{m}$. We will typically restrict our attention to unit-norm frames.

Of particular interest in frame design are the magnitudes of the inner products between distinct frame elements, $|\langle f_i, f_j \rangle|$, $i \neq j$. A unit-norm frame is called *equiangular* if all of these magnitudes are equal: $|\langle f_i, f_j \rangle| = \alpha$, $\forall i \neq j$, for some constant α. In general, we would like all of the inner

product magnitudes to be as small as possible so that the frame vectors are well-spaced about the m-dimensional unit sphere. We define the *coherence* μ of the frame to be the largest of these magnitudes:

$$\mu = \max_{i \neq j} \frac{|\langle f_i, f_j \rangle|}{||f_i||_2 \cdot ||f_j||_2}.$$

Designing frames with low coherence is a problem that has connections to a wide range of fields, including compressive sensing [10–12, 37, 38, 101], spherical codes [31, 91], LDPC codes [44], MIMO communications [56, 57], quantum measurements [43, 84, 85], etc. Frame theory has also made its mark as an interesting field in its own right, with a great collection of recent work by Casazza, Kutyniok, Fickus, Mixon, and many others [13–16, 18, 45].

The following is a classical lower bound on the coherence due to Welch [108]:

Theorem 4. *Let $\{f_i\}_{i=1}^n$ be a unit-norm frame in $\mathbb{C}^m$ or $\mathbb{R}^m$. The coherence $\mu := \max_{i \neq j} |\langle f_i, f_j \rangle|$ satisfies*

$$\mu \geq \sqrt{\frac{n-m}{m(n-1)}}, \tag{3.3}$$

with equality if and only $\{f_i\}$ is both tight and equiangular. In this case, the frame is called Grassmanian.

Proof. The bound in (3.3) is one of a more general set of bounds originally derived by Welch in [108]. This version of the theorem is typically proven (e.g. in [91]) by considering the eigenvalues of the Gram matrix $\mathbf{G} := [\langle f_i, f_j \rangle]$. $\mathbf{G}$ is positive semidefinite of rank at most m. If $\mathbf{G}$ has eigenvalues $\lambda_1, ..., \lambda_n$, which are necessarily real and nonnegative, we may assume that at most the first m of these are nonzero. The Frobenius norm gives us

$$\sum_{i=1}^n \lambda_i^2 = \text{Tr}(\mathbf{G}^2) = \sum_{i,j} |\langle f_i, f_j \rangle|^2. \tag{3.4}$$

The right side of (3.4) becomes

$$\sum_{i,j} |\langle f_i, f_j \rangle|^2 = n + \sum_{i \neq j} |\langle f_i, f_j \rangle|^2 \leq n + \frac{n(n-1)}{2}\mu^2, \tag{3.5}$$

where $\mu = \max_{i,j} |\langle f_i, f_j \rangle|$. Here, equality is achieved if and only if the frame is equiangular.

On the other hand, from the Cauchy-Schwartz inequality we have

$$n^2 = \text{Tr}(\mathbf{G})^2 = \left(\sum_{i=1}^m \lambda_i\right)^2 \leq m \sum_{i=1}^m \lambda_i^2 = m \sum_{i=1}^n \lambda_i^2. \tag{3.6}$$

Here we have equality if and only if the λ_i are all equal, which is the case if and only if the frame is tight.

Finally, combining (3.4), (3.5), and (3.6), we obtain the result.

$\square$

Thus, we would like to identify tight, equiangular frames for use in constructing matrices which achieve this lower bound. This problem arises in various contexts, for example line packing problems [27]. It should be emphasized that such frames do not exist for all values of m and n, so in general, we would also like to find ways to optimize the coherence by choosing $\mathbf{M}$ cleverly from an appropriate class of matrices. Our approach will be to use the group frame construction proposed by Slepian [90] in the 1960s. Group frames have received a great deal of attention in recent years, notably in the substantial collection of work by Vale, Waldron, and others [25, 55, 102–104, 106]. We will discuss them in some detail shortly, but an excellent review of the work in group frames can be found in [17].

On one final note before proceeding, a common approach to produce a set of vectors with low correlation is to construct a set of Mutually Unbiased Bases (MUBs). Two bases $\{e_1, ..., e_m\}$ and $\{e'_1, ..., e'_m\}$ for $\mathbb{C}^d$ are mutually unbiased if each is orthonormal, and $|\langle e_i, e'_j \rangle| = \frac{1}{\sqrt{m}}$ for any i and j. Algebraic constructions of up to $m + 1$ MUBs are known in prime-power dimensions m, allowing for a number of vectors at most $m^2 + m$ [2, 65, 109]. The frame constructions presented in this chapter will at times outperform this coherence, though typically with a smaller number of vectors. More importantly, though, our frames will not require m to be prime.

3.1 Reducing the Number of Distinct Inner Products in Tight Frames

In practice, constructing frames which are both tight and equiangular can prove difficult. It turns out, however, that we can expect reasonably low coherence from tight frames if we just require that the inner products between frame elements take on few distinct values, provided that each of these values arises the same number of times. We begin with the following generalization of the Welch bound:

Lemma 5. *Let $\{f_i\}_{i=1}^n$ be a unit-norm frame in $\mathbb{C}^m$ or $\mathbb{R}^m$. Then the mean value of the $n(n-1)$ squared inner product norms $\{|\langle f_i, f_j \rangle|^2\}_{i \neq j}$ satisfies*

$$\frac{1}{n(n-1)} \sum_{i \neq j} |\langle f_i, f_j \rangle|^2 \geq \frac{n-m}{m(n-1)}, \tag{3.7}$$

with equality if and only if $\{f_i\}$ is a tight frame.

Proof. The quantity $\frac{1}{n(n-1)} \sum_{i \neq j} |\langle f_i, f_j \rangle|^2$ is very closely related to the *frame potential* defined in [4], and (3.7) follows from Theorem 6.2 in that work. In the interest of being self-contained, we remark that the proof essentially follows the second half of the proof of the Welch Bound (Theorem 4). From equations (3.4) and (3.6), we obtain

$$n^2 \leq m \sum_{i,j} |\langle f_i, f_j \rangle|^2 = m \left(n + \sum_{i \neq j} |\langle f_i, f_j \rangle|^2 \right), \tag{3.8}$$

and rearranging terms we get (3.7). Since (3.6) holds with equality if and only if the frame is tight, the same is true of (3.7), and we are done. □

Using Lemma 5, we can obtain upper bounds on the coherence of tight frames which become particularly effective when there are few distinct inner product values $\{|\langle f_i, f_j \rangle|\}_{i \neq j}$, with each value arising the same number of times in this set.

Lemma 6. *Let $\{f_i\}_{i=1}^n$ be a unit-norm tight frame in $\mathbb{C}^m$ or $\mathbb{R}^m$, such that the inner product norms $\{|\langle f_i, f_j \rangle|\}_{i \neq j}$ take on κ distinct values, with each value arising the same number of times as such an inner product norm. Then the coherence μ of the frame is at most a factor of $\sqrt{\kappa}$ greater than the Welch bound:*

$$\mu \leq \sqrt{\kappa} \sqrt{\frac{n-m}{m(n-1)}}. \tag{3.9}$$

Proof. Let $\alpha_1, ..., \alpha_\kappa$ be the κ distinct nonnegative values assumed by the inner product norms $|\langle f_i, f_j \rangle|, i \neq j$. Since each α_i arises the same number $\left(\frac{n(n-1)}{\kappa} \right)$ of times as such a norm by hypothesis, we have that the mean of the squared inner product norms is equal to that of the α_i^2:

$$\frac{1}{n(n-1)} \sum_{i \neq j} |\langle f_i, f_j \rangle|^2 = \frac{1}{\kappa} \sum_{i=1}^\kappa \alpha_i^2. \tag{3.10}$$

As a result, Lemma 5 gives us

$$\frac{1}{\kappa} \sum_{i=1}^\kappa \alpha_i^2 = \frac{n-m}{m(n-1)}. \tag{3.11}$$

Now we can bound the coherence as

$$\mu^2 = \max_i \{\alpha_i^2\} \leq \sum_i \alpha_i^2 = \kappa \frac{n-m}{m(n-1)}, \tag{3.12}$$

from which the (3.9) follows. □

Notice that when $\kappa = 1$ in Lemma 6, the frame becomes both tight and equiangular, and fittingly the coherence in (3.9) achieves the Welch bound. In what follows, we will discuss constructions of unit-norm tight frames in which we could control the number of distinct inner product values and ensure that each arises with the same multiplicity. These constructions appear in [96–99]

3.2   Frames from Unitary Group Representations: Slepian Group Codes

In [90], Slepian proposed a method to construct low-coherence matrices by reasoning that the key to controlling the inner products between the columns was to reduce the number of distinct inner product values which arise. His construction, which has come to be known as a *group frame*, also called a "group code," has since been generalized (see, for example [102] and [17]). On this note, let $\mathcal{U} = \{\mathbf{U}_1, \mathbf{U}_2, ..., \mathbf{U}_n\}$ be a (multiplicative) group of unitary matrices. We can equivalently view $\mathcal{U}$ as the image of a faithful, unitary representation of a group $\mathcal{G}$. In some works, e.g. [46], $\mathcal{U}$ is taken to be a group-like unitary operator system—the image of a projective representation—but normal representations will suffice for our purposes. Such representations exist for any finite group.

Suppose that for each i, we have $\mathbf{U}_i \in \mathbb{C}^{m \times m}$ (or equivalently, $\mathcal{U}$ is the image of an m-dimensional representation). Let $\mathbf{v} = [v_1, ..., v_m]^T \in \mathbb{C}^{m \times 1}$ be any vector, and let $\mathbf{M}$ be the matrix whose i^{th} column is $\mathbf{U}_i \mathbf{v}$:

$$\mathbf{M} = [\mathbf{U}_1 \mathbf{v}, ..., \mathbf{U}_n \mathbf{v}].$$

The inner product between the i^{th} and j^{th} columns of $\mathbf{M}$ is $\langle \mathbf{U}_i \mathbf{v}, \mathbf{U}_j \mathbf{v} \rangle = \mathbf{v}^* \mathbf{U}_i^* \mathbf{U}_j \mathbf{v}$. Since $\mathcal{U}$ is a unitary group, we have

$$\mathbf{U}_i^* \mathbf{U}_j = \mathbf{U}_i^{-1} \mathbf{U}_j = \mathbf{U}_k,$$

for some $k \in \{1, ..., n\}$, so we can write

$$\langle \mathbf{U}_i \mathbf{v}, \mathbf{U}_j \mathbf{v} \rangle = \mathbf{v}^* \mathbf{U}_k \mathbf{v}.$$

In particular, each column has the same norm $||\mathbf{U}_i \mathbf{v}||_2 = ||\mathbf{v}||_2$, so if we assume $\mathbf{v}$ is normalized then the columns of $\mathbf{M}$ form a unit-norm set. In this manner, we have reduced the total number of pairwise inner products between the columns of $\mathbf{M}$ from a possible $\binom{n}{2}$ to only $n - 1$, the inner products parametrized by the non-identity elements of $\mathcal{U}$. Furthermore, we have the following:

Lemma 7. *Let $\{\mathbf{U}_1, ..., \mathbf{U}_n\} \subset \mathbb{C}^{m \times m}$ be a set of distinct unitary matrices which form a group under multiplication, and let $\mathbf{v} \in \mathbb{C}^{m \times 1}$ be a nonzero vector. Each of the values $\mathbf{v}^* \mathbf{U}_k \mathbf{v}$ occurs as the inner*

product between two columns of $\mathbf{M} = [\mathbf{U}_1\mathbf{v}, ..., \mathbf{U}_n\mathbf{v}]$ *the same number of times.*

Proof. For every choice of $\mathbf{U}_k$ and $\mathbf{U}_i$, there is a unique $\mathbf{U}_j$ such that $\mathbf{U}_i^{-1}\mathbf{U}_j = \mathbf{U}_k$. Thus, for each $\mathbf{U}_k$, there are n pairs $(\mathbf{U}_i, \mathbf{U}_j)$ such that $\mathbf{v}^*\mathbf{U}_i^*\mathbf{U}_j\mathbf{v} = \mathbf{v}^*\mathbf{U}_k\mathbf{v}$. $\square$

This result suggests that group codes lend themselves to analysis via Lemmas 5 and 6. One should be wary, however, about applying the bound in Lemma 6, since in this case the number of inner product magnitudes could be as high as $\kappa = n - 1$, in which case the bound becomes $\mu \leq \sqrt{\frac{n-m}{m}}$. Since this upper bound is greater than 1 when $n > 2m$, it provides no useful information in this regime.

Remark: Note that the inner products corresponding to $\mathbf{U}_k$ and $\mathbf{U}_k^{-1}$ actually have the same norms, since

$$|\langle \mathbf{v}, \mathbf{U}_k\mathbf{v}\rangle| = |\langle \mathbf{U}_k\mathbf{v}, \mathbf{v}\rangle| = |\langle \mathbf{v}, \mathbf{U}_k^*\mathbf{v}\rangle| = |\langle \mathbf{v}, \mathbf{U}_k^{-1}\mathbf{v}\rangle|.$$

Thus in practice, depending on our group $\mathcal{U}$, Slepian's construction can in fact give us as few as $\frac{n-1}{2}$ distinct nontrivial inner product values, though it is important to bare in mind that they may not arise with the same multiplicity when grouped together in this fashion.

3.3 Abelian Groups and Harmonic Frames

When $\mathcal{U}$ is chosen to be abelian, so that all the $\mathbf{U}_i$ commute with each other, then the matrices can be simultaneously diagonalized by a unitary change of basis matrix $\mathbf{B}$ so that we may write $\mathbf{B}^*\mathbf{U}_i\mathbf{B} = \mathbf{D}_i$, where $\mathbf{D}_i$ is diagonal. In this case the inner product corresponding to $\mathbf{U}_i$ will take the form

$$\mathbf{v}^*\mathbf{U}_i\mathbf{v} = \mathbf{v}^*\mathbf{B}^*\mathbf{D}_i\mathbf{B}\mathbf{v},$$

so by replacing $\mathbf{v}$ with $\mathbf{B}^*\mathbf{v}$ without loss of generality, we may assume that the $\mathbf{U}_i$ are already diagonal. Furthermore, since each $\mathbf{U}_i$ must have a multiplicative order dividing the size of $\mathcal{U}$, we may take the diagonal entries of $\mathbf{U}_i$ to be powers of the n^{th}-root of unity $\omega := e^{\frac{2\pi i}{n}}$. The matrices will then take the form

$$\mathbf{U}_j = \text{diag}(\omega^{a_{1,j}}, ..., \omega^{a_{m,j}}) \in \mathbb{C}^{m \times m},$$

where the $a_{i,j}$ are integers between 0 and $n - 1$. In the language of representation theory we have decomposed $\mathcal{U}$ into its irreducible representations, all of which are degree-1 since $\mathcal{U}$ is abelian.

If we write the coordinates of our rotated vector as $\mathbf{v} = (v_1, ..., v_m)^T \in \mathbb{C}^{m \times 1}$, then our inner

products will now take the form

$$\mathbf{v}^* \mathbf{U}_j \mathbf{v} = \sum_{i=1}^{m} \omega^{a_{i,j}} |v_i|^2, \tag{3.13}$$

so we see that the inner products depend only on the magnitudes of the v_i, which weight the diagonal entries of the $\mathbf{U}_j$. In particular, for the sake of minimizing coherence, we may take the entries of $\mathbf{v}$ to be real. Furthermore, it turns out that in order for our abelian group frame to be tight, all the entries v_i must be of equal norm. This follows from Theorem 5.4 in [17], and we will touch on this in Section 4.1. On this note, we will consider the case where $\mathbf{v}$ is a scaled vector of all 1's,

$$\mathbf{v} = \frac{1}{\sqrt{m}} \mathbf{1}_m = \frac{1}{\sqrt{m}} [1, ..., 1]^T \in \mathbb{C}^{m \times 1}, \tag{3.14}$$

where we have again chosen $\mathbf{v}$, and hence all the vectors $\mathbf{U}_i \mathbf{v}$, to be unit-norm. Now the inner product norm corresponding to the element $\mathbf{U}_j$ becomes simply

$$\frac{|\mathbf{v}^* \mathbf{U}_j \mathbf{v}|}{||\mathbf{v}||_2^2} = \frac{1}{m} \left| \sum_{i=1}^{m} \omega^{a_{i,j}} \right|. \tag{3.15}$$

Notice that from Equation (3.15), we can see that the coherence of our final matrix would remain unchanged if we chose ω to be any other primitive n^{th} root of unity. Indeed, if we replace ω with another primitive root of unity, which we may write as ω^b where b is relatively prime to n, then the inner product associated with $\mathbf{U}_j$ will become $\frac{1}{m} \left| \sum_{i=1}^{m} \omega^{b \cdot a_{i,j}} \right|$. But this is just the original inner product associated with $\mathbf{U}_j^b$, which in turn generates the entire cyclic group $\langle \mathbf{U}_j \rangle$. Hence, the inner products which arise using any two primitive n^{th} roots of unity are the same.

When we form a frame from an abelian group in this manner, and in addition require the sets of diagonal components $\{\omega^{a_{i,j}}\}_{i=1}^{m}$ to form distinct representations of $\mathcal{U}$, we obtain what is called a *harmonic frame*, which we will define concretely as follows:

Definition 2. *Let m and n be integers, $\omega = e^{\frac{2\pi i}{n}}$, and $\mathbf{U}_j = diag(\omega^{a_{1,j}}, ..., \omega^{a_{m,j}}) \in \mathbb{C}^{m \times m}$ for $j = 1, ..., n$, where the $a_{i,j}$ are integers between 0 and $n-1$. If we set $\mathbf{v} = \frac{1}{\sqrt{m}} [1, ..., 1]^T \in \mathbb{C}^{m \times 1}$, and $\mathbf{M} = [\mathbf{U}_1 \mathbf{v}, ..., \mathbf{U}_n \mathbf{v}]$, then if the rows of $\mathbf{M}$ are distinct, we call the set of columns $\{\mathbf{U}_j \mathbf{v}\}_{j=1}^{n}$ a harmonic frame.*

Remark: Note that we have sidestepped the discussion of whether a "harmonic frame" is actually a frame in the sense of Definition 1. But indeed it is, as we will discuss in the proof of Lemma 8.

Harmonic frames are one of the most thoroughly-studied types of structured frames [25, 55]. An important example of a harmonic frame arises when we choose the group $\mathcal{U}$ to be cyclic, meaning that

each $\mathbf{U}_j$ is a power of a single matrix $\mathbf{U}$, which we will explicitly write as

$$\mathbf{U} = \mathrm{diag}(\omega^{a_1}, ..., \omega^{a_m}), \tag{3.16}$$

so we may write $\mathbf{U}_j := \mathbf{U}^j$. For cyclic groups, the inner product between the columns $\mathbf{U}^{\ell_1}\mathbf{v}$ and $\mathbf{U}^{\ell_2}\mathbf{v}$, after normalizing the columns, will take the form $\frac{|\mathbf{v}^*\mathbf{U}^{\ell_2-\ell_1}\mathbf{v}|}{\|\mathbf{v}\|_2^2}$, which is the value of the inner product determined by $\mathbf{U}^{\ell_2-\ell_1}$ in (3.15).

In this case, if we again take $\mathbf{v}$ to be the normalized vector of all 1s, our frame matrix takes the form

$$\mathbf{M} = \begin{bmatrix} \mathbf{v} & \mathbf{U}\mathbf{v} & \dots & \mathbf{U}^{n-1}\mathbf{v} \end{bmatrix} \tag{3.17}$$

$$= \frac{1}{\sqrt{m}} \begin{bmatrix} 1 & \omega^{a_1} & \omega^{a_1 \cdot 2} & \dots & \omega^{a_1 \cdot (n-1)} \\ 1 & \omega^{a_2} & \omega^{a_2 \cdot 2} & \dots & \omega^{a_2 \cdot (n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{a_m} & \omega^{a_m \cdot 2} & \dots & \omega^{a_m \cdot (n-1)} \end{bmatrix}, \tag{3.18}$$

where the columns form a harmonic frame precisely when the a_i are distinct. In this form, we see that $\mathbf{M}$ is a subset of rows of the $n \times n$ discrete Fourier matrix, so it becomes clear that the columns of $\mathbf{M}$ form a tight frame since $\mathbf{M}\mathbf{M}^* = \frac{n}{m}\mathbf{I} \in \mathbb{C}^{m \times m}$. In fact, this is true of all harmonic frames:

Lemma 8. *A harmonic frame is a tight, unit-norm frame.*

Proof. The fact that harmonic frames are unit-norm follows straight from the definition. We note that the rest of this lemma is proven in [17], and we will explain the tightness of harmonic frames in Section 4.1 when we discuss tight group frames in greater generality. For now, however, we will provide a simple, self-contained proof.

A general abelian group $\mathcal{U}$ can be represented as follows: first express G as a direct product of, say, L cyclic groups of orders $n_1, ..., n_L$, so that $\mathcal{U} \cong \frac{\mathbb{Z}}{n_1\mathbb{Z}} \times ... \times \frac{\mathbb{Z}}{n_L\mathbb{Z}}$. Then let $\omega_1, ..., \omega_L$ be the corresponding primitive roots of unity: $\omega_j = e^{2\pi i/n_j}$. Then we set $\mathbf{U}_j = \mathrm{diag}(\omega_j^{a_{1j}}, ..., \omega_j^{a_{mj}})$, where we will assume that the a_{ij} are distinct integers modulo n_j. The abelian group generated by the diagonal matrices $\{\mathbf{U}_1, ..., \mathbf{U}_L\}$ is isomorphic to $\mathcal{U}$, and an arbitrary element will take the form $\mathbf{U}_1^{b_1}\mathbf{U}_2^{b_2}\dots\mathbf{U}_L^{b_L}$, where $b_j \in \{0, ..., n_j - 1\}$. Our frame matrix $\mathbf{M}$ will then take the form $\mathbf{M} = [\dots (\mathbf{U}_1^{b_1}\mathbf{U}_2^{b_2}\dots\mathbf{U}_L^{b_L}\mathbf{v})\dots]_{0 \le b_j \le n_j - 1}$.

In this form, our previous cyclic frames clearly arise as subsets of the columns of $\mathbf{M}$. It is not too difficult to see that the frame matrix $\mathbf{M}$ is a subset of rows of the Kronecker product $\mathbf{A}_{\mathrm{Kron}} :=$

$\mathbf{A}_1 \otimes ... \otimes \mathbf{A}_L$, where $\mathbf{A}_j = \left[\mathbf{v}, \mathbf{U}_j\mathbf{v}, ..., \mathbf{U}_j^{n_j-1}\mathbf{v}\right]$. By the properties of the Kronecker product,

$$\mathbf{A}_{\mathrm{Kron}}\mathbf{A}_{\mathrm{Kron}}^* = (\mathbf{A}_1 \otimes ... \otimes \mathbf{A}_L)(\mathbf{A}_1 \otimes ... \otimes \mathbf{A}_L)^* \tag{3.19}$$

$$= (\mathbf{A}_1 \otimes ... \otimes \mathbf{A}_L)(\mathbf{A}_1^* \otimes ... \otimes \mathbf{A}_L^*) \tag{3.20}$$

$$= \mathbf{A}_1\mathbf{A}_1^* \otimes ... \otimes \mathbf{A}_L\mathbf{A}_L^*, \tag{3.21}$$

and since each $\mathbf{A}_j\mathbf{A}_j^*$ is a multiple of the identity matrix, so is their Kronecker product. It follows that the columns of $\mathbf{M}$ are indeed a tight frame. $\qquad\square$

As we will see in the next few sections, there is a lot we can do in optimizing frame coherence even if we restrict our attention to harmonic frames.

3.4 Equiangular Frames from Cyclic Group Representations

Let us examine the "cyclic" harmonic frame formed by the columns of $\mathbf{M}$ in (3.18). [110] classified the conditions on the a_i under which this frame is equiangular. Since we know these frames are tight, this determines precisely when their coherence achieves the Welch lower bound of Theorem 4.

Definition 3. *Let G be a group. A difference set $A = \{a_1, ..., a_m\} \subset G$ is a set of elements such that every nonidentity element $g \in G$ occurs as a difference $a_i - a_j$ the same number of times. That is, the sets $A_g := \{(a_i, a_j) \in A \times A \mid a_i - a_j = g\}$ have the same size for $g \neq 0$.*

Theorem 5 ([110] Equiangular Cyclic Harmonic Frames)**.** *The harmonic frame formed by the columns of $\mathbf{M}$ in (3.18) is equiangular if and only if the integers a_i form a difference set in $\mathbb{Z}/n\mathbb{Z}$.*

Proof. The proof follows from a simple but insightful Fourier connection. Let us define $\mathcal{A}_t := \{(a_i, a_j) \in A \times A \mid a_i - a_j \equiv t \mod n\}$ for any $t \in \mathbb{Z}/n\mathbb{Z}$, and set $N_t := |\mathcal{A}_t|$. Furthermore, if we index the columns as $\ell = 0, 1, ..., n-1$ then the inner product associated to the ℓ^{th} column takes the form

$$c_\ell := \frac{\mathbf{v}^*\mathbf{U}^\ell\mathbf{v}}{||\mathbf{v}||_2^2} = \frac{1}{m}\sum_{a \in A}\omega^{\ell a}.$$

Since we are concerned only with the magnitude of c_ℓ, we may consider the quantity

$$\alpha_\ell := |c_\ell|^2 = \frac{1}{m^2}\left(\sum_{a \in A}\omega^{\ell a}\right)^*\left(\sum_{a \in A}\omega^{\ell a}\right)$$

$$= \frac{1}{m^2}\sum_{a_i, a_j \in A}\omega^{\ell(a_i - a_j)}.$$

We can then write

$$\alpha_\ell = \frac{1}{m^2} \sum_{t=0}^{n-1} N_t \omega^{\ell t}, \tag{3.22}$$

which gives us a Fourier pairing between the α_ℓ and the N_t with inverse transform given by

$$N_t = \frac{m^2}{n} \sum_{\ell=0}^{n-1} \alpha_\ell \omega^{-t\ell}. \tag{3.23}$$

M will be an equiangular tight frame precisely when all of the α_ℓ are equal for $\ell \neq 0$, and from the Fourier pairing this will occur precisely when the N_t are equal for $t \neq 0$, i.e., when the a_i form a difference set. $\qquad \square$

This concept of tight equiangular frames arising from difference sets has since been generalized and elaborated [17, 35, 106]. [34] showed how slightly relaxed forms of difference sets can produce frames which have coherence almost reaching the Welch Bound. Many of our results in the following sections can also be viewed as relaxing difference sets even further to produce low-coherence frames. Difference sets have been long studied and classified [3, 8]. They have found application in other fields as well, such as designing codes for DS-CDMA systems [36], LDPC codes [105], sonar and synchronization [47], and other forms of frame design [61].

While Theorem 5 completely characterizes the optimal-coherence frames arising from representations of cyclic groups, it reveals that equiangular frames of the form (3.18) are rather scarce, since the number of known difference sets is relatively small. In the following section, we will present a new strategy for selecting the integers a_i which, while not always producing an equiangular frame, does yield frames with few distinct inner product values and provably low coherence.

3.5 Cyclic Groups of Prime Order

We have already managed to cut down the number of distinct inner products between columns from $\binom{n}{2}$ to $n-1$, simply by using a unitary group to generate our columns. For cyclic groups, however, we can reduce this number even more. We first consider the case where n is prime. Let $H = (\mathbb{Z}/n\mathbb{Z})^\times$, the multiplicative group of the integers modulo n. As usual, we identify the elements of $\mathbb{Z}/n\mathbb{Z}$ with the integers $0, 1, ..., n-1$. Since n is assumed to be a prime, H is itself a cyclic group, and consists of the $n-1$ nonzero elements of H. Now let us choose m to be any divisor of $n-1$, and set $\kappa := \frac{n-1}{m}$. Since H is cyclic, it has a unique subgroup A of order m consisting of the distinct κ^{th} powers of the elements of H. In fact, if g is any generator for H, then A will be generated by $a := g^{\frac{n-1}{m}}$. Now,

if we write out the elements of A as $\{a_1, ..., a_m\}$ (or equivalently in terms of a single generator a as $\{1, a, a^2, ..., a^{m-1}\}$), we can form our generator matrix $\mathbf{U}$ as in (3.16), choosing $\omega^{a_i} = \omega^{a^{i-1}}$ to be the i^{th} diagonal term. Note that since A consists of elements relatively prime to n, then for each i, ω^{a_i} has multiplicative order n. It follows that $\mathbf{U}$ also has order n and generates the cyclic group $\mathcal{U} = \{\mathbf{U}^\ell\}_{\ell=0}^{n-1} \cong \mathbb{Z}/n\mathbb{Z}$.

It turns out that this construction both 1) reduces the number of distinct inner product values between our columns and 2) maintains the property that each such value occurs with the same multiplicity:

Theorem 6. *Let n be a prime and m any divisor of $n-1$. Take $A = \{a_1, ..., a_m\}$ to be the unique (cyclic) subgroup of $H = (\mathbb{Z}/n\mathbb{Z})^\times$ of size m. Set $\omega = e^{\frac{2\pi i}{n}}$, $\mathbf{v} = \frac{1}{\sqrt{m}}[1, ..., 1]^T \in \mathbb{C}^m$, and $\mathbf{U} = \text{diag}(\omega^{a_1}, ..., \omega^{a_m})$. Then the columns of $\mathbf{M} = [\mathbf{v}, \mathbf{Uv}, ..., \mathbf{U}^{n-1}\mathbf{v}] \in \mathbb{C}^{m \times n}$ form a unit-norm tight frame with at most $\frac{n-1}{m}$ distinct inner product values between its columns, each occurring with the same multiplicity.*

Proof. For any integer ℓ in the set $\{1, ..., n-1\}$, the inner product corresponding to $\mathbf{U}^\ell$ (as in Equation (3.15)) will take the following form:

$$\frac{|\mathbf{v}^*\mathbf{U}^\ell\mathbf{v}|}{||\mathbf{v}||_2^2} = \frac{1}{m}\left|\sum_{i=1}^m \omega^{\ell \cdot a_i}\right|. \tag{3.24}$$

Notice the exponents of ω appearing in the above summation can be taken modulo n, since ω is an n^{th} root of unity, and are then simply the elements of the ℓ^{th} coset of K in H, $\ell A = \{\ell \cdot a_1, ..., \ell \cdot a_m\}$. The set of all cosets of A in H is denoted H/A. From elementary group theory, we know that the distinct cosets of A form a disjoint partition of H, so the number of distinct cosets of K in H is the quotient of their sizes: $|H/A| = \frac{|H|}{|A|} = \frac{n-1}{m}$. Thus, the total number of distinct pairwise inner products that we now must control is $\frac{n-1}{m}$.

It only remains to show that each of the $\frac{n-1}{m}$ inner products occurs the same number of times. Let $\{\ell_1, ..., \ell_r\}$ be a complete set of coset representatives for K in $(\mathbb{Z}/n\mathbb{Z})^\times$. Here, r is simply $\frac{n-1}{m}$. Then every element in $\{1, ..., n-1\}$ can be written uniquely as a product $\ell_i a_j$, and from Lemma 7 the $n-1$ inner products $\mathbf{v}^*\mathbf{U}^{\ell_i a_j}\mathbf{v}$ all arise the same number of times. As described above, the $\frac{n-1}{m}$ distinct inner product values correspond to the cosets of K, i.e., for a fixed ℓ_i the m inner products $\mathbf{v}^*\mathbf{U}^{\ell_i a_1}\mathbf{v}, ..., \mathbf{v}^*\mathbf{U}^{\ell_i a_m}\mathbf{v}$ will give rise to one of the distinct inner product values. Thus, since each distinct value corresponds to m inner products, each arising the same number of times, our result is proved. $\square$

We emphasize the power of this construction in reducing the number of inner products that we must control in order to maintain low matrix coherence. Since we are free to choose m to be *any*

divisor of $n-1$, then for properly chosen matrix dimensions, we can reasonably create matrices with just two or three distinct values of inner products between columns. In practice, this often creates matrices with remarkably low coherence, far outmatching that of any known randomly-generated matrices. In Table A.1, we compare the coherences of the "Group Matrices" from our construction with those of randomly-generated complex Gaussian matrices and matrices designed by randomly selecting m rows from the $n \times n$ Fourier matrix. (This latter construction is equivalent to randomly selecting the exponents k_i in our cyclic generator matrix $\mathbf{U}$ in (3.16).) For convenience, we also list the lower bound on coherence from Theorem 4, and we underline the coherences which achieve this bound. Figure 3.1 illustrates explicitly the inner products for a random Fourier matrix vs. a Group matrix.

Table 3.1: Coherences for Random and Group Matrices (for n a Prime)

(n, m)	Complex Gaussian	Random Fourier	Group Matrix	$\sqrt{\frac{n-m}{m(n-1)}}$
$(251, 125)$	.2677	.1996	<u>.0635</u>	.0635
$(499, 166)$	.3559	.1786	.0888	.0635
$(499, 249)$	.2226	.1736	<u>.0449</u>	.0449
$(503, 251)$	.2137	.1533	<u>.0447</u>	.0447
$(521, 260)$	.2208	.1504	.0458	.0439
$(521, 130)$	.3065	.2376	.1175	.0761
$(643, 321)$	.2034	.1627	<u>.0395</u>	.0395
$(643, 214)$	.2274	.1978	.0755	.0559
$(701, 175)$	.2653	.2316	.0687	.0655
$(701, 350)$	.1788	.1326	.0393	.0379
$(1009, 504)$	.1565	.1147	.0325	.0315
$(1009, 336)$	.2086	.1384	.0597	.0446
$(1009, 252)$	.2287	.1631	.0846	.0546

3.6 Sharper Bounds on Coherence for Frames from Cyclic Groups of Prime Order

In the special case where we construct our frame as in Theorem 6 (using Slepian's approach with a group $\mathcal{U} \cong \mathbb{Z}/n\mathbb{Z}$ and n prime), we have a great deal of underlying algebraic structure in our frame. So it should come as no surprise that we can derive sharper bounds on our coherence and even compute it exactly in some cases.

As before, let m be a divisor of $n-1$, and take $A = \{a_1, ..., a_m\}$ to be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m. Define $\kappa := \frac{n-1}{m}$, which is the number of distinct inner product values. If κ is small, it becomes relatively simple to analyze these values. For example:

Figure 3.1: The norms of the inner products associated to each group element for (a) randomly-chosen A, and (b) A selected to be a subgroup of $(\mathbb{Z}/n\mathbb{Z})^{\times}$ of index 3. Here, $n = 499$ (a prime) and $m = 166$. In (b), as expected, there are only three distinct values of the inner products between distinct, normalized columns.

Theorem 7 ($\kappa = 2$). *Let n be a prime, m a divisor of $n-1$, and $\omega = e^{\frac{2\pi i}{n}}$. Let $A = \{a_1, ..., a_m\}$ be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^{\times}$ of size m, and set $\mathbf{U} = \mathrm{diag}(\omega^{a_1}, ..., \omega^{a_m}) \in \mathbb{C}^{m \times m}$, $\mathbf{v} = \frac{1}{\sqrt{m}}[1, ..., 1]^T \in \mathbb{C}^{m \times 1}$, and $\mathbf{M} = [\mathbf{v}, \mathbf{Uv}, ..., \mathbf{U}^{n-1}\mathbf{v}]$.*

If $\kappa := \frac{n-1}{m} = 2$, there are two distinct inner product values between the columns of $\mathbf{M}$, both of which are real. If $n-1$ is divisible by 4, these inner products are $\frac{-1 \pm \sqrt{1+2m}}{2m}$. In this case, $\mathbf{M}$ has coherence $\sqrt{\frac{n-m-\frac{1}{2}}{m(n-1)} + \frac{1}{2m}}$.

If $n-1$ is not divisible by 4, then the columns of $\mathbf{M}$ form an equiangular frame. The two inner products are $\pm\sqrt{\frac{1}{m}\left(\frac{1}{2} + \frac{1}{2m}\right)}$, and the coherence is $\sqrt{\frac{n-m}{m(n-1)}}$.

Proof. We will hold off on the details of the proof until Appendix A.2 aside from mentioning that it is related to the connection made by Xia et al [110] between tight equiangular harmonic frames and difference sets. In fact, in the case where $n-1$ is not divisible by 4, A forms a known difference set in $\mathbb{Z}/n\mathbb{Z}$. If we view $\mathbb{Z}/n\mathbb{Z}$ as the additive group of $\mathbb{F}_n$, this particular case also overlaps with the tight equiangular frames classified in Theorem 3 of [35]. $\square$

As the number κ of inner products increases, it becomes more complicated to explicitly compute their values or even just the coherence of the resulting frame. While there were only two cases to consider when $\kappa = 2$, there are many more even for κ as low as 3. We can, however, exploit the

algebraic structure of our frames to yield bounds on their coherence which in practice prove to be nearly tight.

Theorem 8 ($\kappa = 3$). *Let n be a prime, m a divisor of $n-1$, and $\omega = e^{\frac{2\pi i}{n}}$. Let $A = \{a_1, ..., a_m\}$ be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^{\times}$ of size m, and set $\mathbf{U} = \mathrm{diag}(\omega^{a_1}, ..., \omega^{a_m}) \in \mathbb{C}^{m \times m}$, $\mathbf{v} = \frac{1}{\sqrt{m}}[1, ..., 1]^T \in \mathbb{C}^{m \times 1}$, and $\mathbf{M} = [\mathbf{v}, \mathbf{U}\mathbf{v}, ..., \mathbf{U}^{n-1}\mathbf{v}]$.*

If $\kappa := \frac{n-1}{m} = 3$, then the coherence of $\mathbf{M}$ will satisfy

$$\mu \leq \frac{1}{3}\left(2\sqrt{\frac{1}{m}\left(3 + \frac{1}{m}\right)} + \frac{1}{m}\right) \approx \sqrt{\frac{4}{3m}}, \tag{3.25}$$

and for large enough m, we will asymptotically have the following lower bound on coherence:

$$\mu \geq \frac{1}{\sqrt{m}} \ (asymptotically), \tag{3.26}$$

which is strictly greater than the Welch bound.

Proof. We present the proof in Appendix A.3. $\qquad\square$

From Theorem 8 we see that unlike when $\kappa = 2$, we can never hope to achieve the Welch bound with these frames when $\kappa = 3$. But this is not a trend, for our frames will again be able to achieve the Welch bound for certain higher values of κ, including $\kappa = 4$ and $\kappa = 8$. This again relates to the connection with difference sets from [110]. As a result, the lower bound on coherence in Theorem 8 does not generalize to all values of κ. Fortunately, the upper bound does:

Theorem 9 (General κ). *Let n be a prime, m a divisor of $n - 1$, and $\omega = e^{\frac{2\pi i}{n}}$. Let $A = \{a_1, ..., a_m\}$ be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^{\times}$ of size m, and set $\mathbf{U} = \mathrm{diag}(\omega^{a_1}, ..., \omega^{a_m}) \in \mathbb{C}^{m \times m}$, $\mathbf{v} = \frac{1}{\sqrt{m}}[1, ..., 1]^T \in \mathbb{C}^{m \times 1}$, and $\mathbf{M} = [\mathbf{v}, \mathbf{U}\mathbf{v}, ..., \mathbf{U}^{n-1}\mathbf{v}]$.*

If $\kappa := \frac{n-1}{m}$, then the coherence μ of $\mathbf{M}$ satisfies the following upper bound:

$$\mu \leq \frac{1}{\kappa}\left((\kappa - 1)\sqrt{\frac{1}{m}\left(\kappa + \frac{1}{m}\right)} + \frac{1}{m}\right). \tag{3.27}$$

Proof. This theorem will be proved in Appendix B.1. $\qquad\square$

This bound is strictly lower than the one from Lemma 6, which applies to all tight frames. In fact, when $n > 2$, we can find an even lower bound on the coherence of our frames constructed in Theorem 6, which surprisingly depends only on whether m is odd:

Theorem 10 (*m* odd)**.** *Let n be an odd prime, m a divisor of $n-1$, and $\omega = e^{\frac{2\pi i}{n}}$. Let $A = \{a_1, ..., a_m\}$ be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m, and set $\mathbf{U} = \mathrm{diag}(\omega^{a_1}, ..., \omega^{a_m}) \in \mathbb{C}^{m \times m}$, $\mathbf{v} = \frac{1}{\sqrt{m}}[1, ..., 1]^T \in \mathbb{C}^{m \times 1}$, and $\mathbf{M} = [\mathbf{v}, \mathbf{U}\mathbf{v}, ..., \mathbf{U}^{n-1}\mathbf{v}]$. Set $\kappa := \frac{n-1}{m}$.*

If m is odd, then the coherence of $\mathbf{M}$ is upper-bounded by

$$\mu \leq \frac{1}{\kappa} \sqrt{\left(\frac{1}{m} + \left(\frac{\kappa}{2} - 1\right)\beta\right)^2 + \left(\frac{\kappa}{2}\right)^2 \beta^2}, \tag{3.28}$$

where $\beta = \sqrt{\frac{1}{m}\left(\kappa + \frac{1}{m}\right)}$.

Proof. We delay the proof of this theorem until Appendix B.1. □

It is worth noting that this latter bound has no analog in the $\kappa = 3$ situation because m must always be even in that case. Indeed, if n is any prime greater than 2 it is necessarily odd, and $n - 1$ is even. Thus $m = \frac{n-1}{3}$ is also even. In the Appendix B.1 we will give an alternate classification for exactly when this latter coherence bound applies. This will allow us to apply our bound to a more general class of frames, which we will discuss in the next chapter. We illustrate the upper and lower bounds for $\kappa = 3$ in Figure 3.2 and the two upper upper bounds from Theorem 10 for when $\kappa = 4$. When $\kappa = 4$, we can also derive different lower bounds on the coherence for when m is even or odd, and together with the two upper bounds from the theorems they form two non-overlapping regions in which the coherences can fall in the graph. While these regions will exist for every κ, they will sometimes overlap (that is, the lower bound on coherence for m even could be less than the upper bound for m odd).

3.7 Optimizing Coherence Over Cosets

While the preceding results give us a deterministic way to construct very low coherence matrices, we can hope to generalize our construction to yield an entire class of group-theoretically based matrices over which we can optimize to find even lower coherences.

As before, let us take n to be a prime, and m a divisor of $n - 1$, and let $H = (\mathbb{Z}/n\mathbb{Z})^\times$. As we remarked, there is a unique subgroup A of H of any order m dividing $n - 1$, and it is cyclic. But suppose m' is a divisor of m, and let $A' = \{a'_1, ..., a'_{m'}\}$ be the unique subgroup of H of order m'. For convenience, let $d = \frac{m}{m'}$. If A is the unique subgroup of H of order m, then A' is a subgroup of A. Taking g to be a generator for H, then $g^{\frac{n-1}{m}}$ is a generator for A and $g^{d\frac{n-1}{m}}$ is a generator for A'.

Now, the set of cosets of A' in H form the group H/A'. This is a cyclic group of size $\frac{n-1}{m'} = d\frac{n-1}{m}$, generated by the coset gA'. We will construct our unitary group $\mathcal{U}$ as follows: take a set of d cosets of A' in H, $\{\ell_1 A', ..., \ell_d A'\}$. Then, for ω a primitive n^{th} root of unity, we let $\mathcal{U}$ be the cyclic group

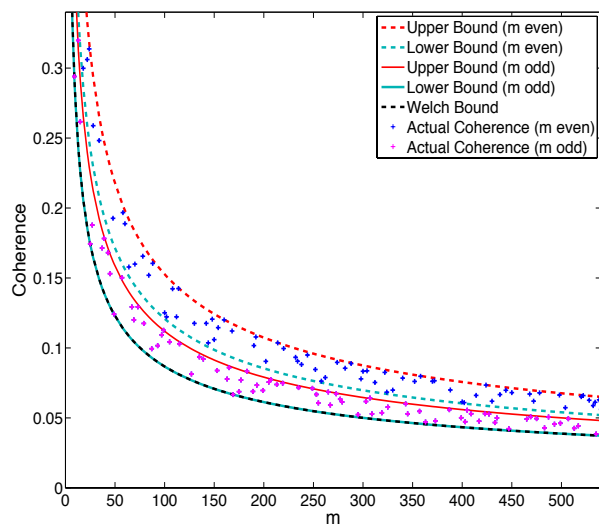Figure 3.2: The upper and lower bounds on coherence for $\kappa = 3$.

Figure 3.3: The upper and lower bounds on coherence for $\kappa = 4$.

generated by the matrix $\mathbf{U}$, which we define as follows: For any $\ell \in H$, let $\mathbf{D}_\ell = \operatorname{diag}\left(\omega^{\ell a_1'}, ..., \omega^{\ell a_{m'}'}\right)$, an $m' \times m'$ diagonal matrix with the elements of $\omega^{\ell a'}, a' \in A'$ along the diagonal. Then define $\mathbf{U}$ to be the (block) diagonal matrix

$$\mathbf{U} := \operatorname{diag}(\mathbf{D}_{\ell_1}, \mathbf{D}_{\ell_2}, ..., \mathbf{D}_{\ell_d}).$$

Now, since each coset of A' in H consists only of elements relatively prime to n, then we see that this matrix will indeed maintain the property of having multiplicative order n, as in our original framework. In fact, if we choose $\ell_i A' = g^{i\frac{n-1}{m}} A'$, for each $i = 1, ..., d$, then since $g^{\frac{n-1}{m}}$ is a generator for A, we find that the cosets $\{\ell_1 A', ..., \ell_d A'\}$ are precisely the cosets of A' as a subgroup of A. These cosets partition the elements of A, so we retrieve the matrix obtained from our original construction, with $\mathbf{U} = \operatorname{diag}\left(\omega^{a_1}, ..., \omega^{a_m}\right)$ up to a permutation of the elements of A (which will not affect the values of the inner products in (3.15)). Thus, this new construction is a direct generalization of our original work. Another special case is when $A' = 1$, the trivial subgroup. In this case, selecting cosets for A' is nothing more than selecting individual rows of the $n \times n$ Fourier matrix for $\mathbf{M}$, with the exception of the row of all 1's.

As we cycle through the powers of $\mathbf{U}$, each $\mathbf{D}_{\ell_i}$ cycles through the different cosets of A' in some order. Since some powers of $\mathbf{U}$ may give rise to permutations of the same cosets, and hence lead to the same corresponding inner product from Equation 3.24, it can take some care to determine precisely how many distinct inner products we have in our constructed matrix. We know that it can be as few as $\frac{n-1}{m}$, as is the case when the chosen cosets of A' partition A. In general, we have the following theorem:

Theorem 11. *Let n be a prime, m a divisor of $n-1$, and m' a divisor of m, with $m = dm'$. Let A be the unique subgroup of $H = (\mathbb{Z}/n\mathbb{Z})^\times$ of size m, and $A' = \{a_1', ..., a_{m'}'\}$ the unique subgroup of size m'. Let $\{\ell_1 A', ..., \ell_d A'\}$ be a set of d cosets of A' in H. Set $\omega = e^{2\pi i/n}$ and $\mathbf{v} = \frac{1}{\sqrt{m}}[1, ..., 1]^T \in \mathbb{C}^{m \times 1}$, and form the matrices $\mathbf{D}_{\ell_i} = \operatorname{diag}\left(\omega^{\ell_i a_1'}, ..., \omega^{\ell_i a_{m'}'}\right) \in \mathbb{C}^{m' \times m'}$, $\mathbf{U} := \operatorname{diag}(\mathbf{D}_{\ell_1}, \mathbf{D}_{\ell_2}, ..., \mathbf{D}_{\ell_d}) \in \mathbb{C}^{m \times m}$, and $\mathbf{M} = [\mathbf{v}, \mathbf{U}\mathbf{v}, ..., \mathbf{U}^{n-1}\mathbf{v}] \in \mathbb{C}^{m \times n}$. Then $\mathbf{M}$ has at most $\frac{d \cdot (n-1)}{m}$ distinct values of the inner products between its columns.*

Proof. We know that the distinct inner products between the normalized columns of $\mathbf{M}$ will correspond to the powers of $\mathbf{U}$. The b^{th} power of $\mathbf{U}$ can be written as $\mathbf{U}^b = \operatorname{diag}\left(\mathbf{D}_{\ell_1}^b, ..., \mathbf{D}_{\ell_d}^b\right)$. Thus, the inner product corresponding to this power of $\mathbf{U}$ is

$$\frac{1}{m}\left| \sum_{a' \in A'} \omega^{\ell_1(ba')} + \sum_{a' \in A'} \omega^{\ell_2(ba')} + ... + \sum_{a' \in A'} \omega^{\ell_d(ba')} \right|. \tag{3.29}$$

Thus, there can only be as many such sums as there are cosets bA'. Since there are $\frac{n-1}{m'} = \frac{d \cdot (n-1)}{m}$

cosets of A' in H, we have our result. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

This coset construction offers us a tradeoff. By using the smaller group A' (of size $\frac{m}{d}$) to construct our matrix as opposed to A (of size m), we gain the possibility of having nice cancelation properties among the sums $\sum_{a' \in A'} \omega^{\ell_i(ba')}$ in (3.29) at the cost of having more inner products to control. But since the number of distinct inner products can increase only by a factor of d at most, this can turn out to be a worthwhile tradeoff, and indeed we have examples where we can strictly decrease the coherence of $\mathbf{M}$ by using this construction. (See Fig. 3.4).

We can now can formulate the problem of constructing low-coherence matrices as an optimization problem, where we can optimize over both the choice of A' and the set of cosets $\{\ell_1 A', ..., \ell_d A'\}$. For fixed m and n, where n is a prime and m a divisor of $n-1$, we must solve the following:

$$\min_{m'|m,\ \underline{\ell} \in G^{\times \frac{m}{m'}}} \left(\max_{b \in H} \frac{1}{m} \left| \sum_{i=1}^{m/m'} \left(\sum_{a \in A_{m'}} \omega^{\ell_i(ba)} \right) \right| \right), \tag{3.30}$$

where $H = (\mathbb{Z}/n\mathbb{Z})^\times$, $\underline{\ell} = (\ell_1, ..., \ell_d)$, $H^{\times \frac{m}{m'}}$ denotes the Cartesian product of H with itself $\frac{m}{m'}$ times, and $A_{m'}$ denotes the unique subgroup of H of size m'.

In practice, it is typically not feasible to perform this exact optimization since it requires a search over the lattice $H^{\times \frac{m}{m'}}$ for every m' dividing m. One simple way to deal with this problem is to fix m' and randomly sample $\underline{\ell} \in H^{\times \frac{m}{m'}}$ to search for the smallest value of the objective function. Note that if $m'_2 | m'_1$ (or equivalently, $A_{m'_2} \leq A_{m'_1}$), then searching over cosets of $A_{m'_2}$ *encompasses* the search over cosets of $A_{m'_1}$ since we can express $A_{m'_1}$ as a union of cosets of $A_{m'_2}$. One might therefore be tempted to argue that it is unnecessary to search over cosets of $A_{m'_1}$ at all. There is, however, value in searching over these cosets, since this search will converge to its optimal value much faster than the search over cosets of the smaller group. See Figure 3.4.

Of course, we can still bound the coherence of the frames that can arise from this construction using that of our previous frames.

Theorem 12. *Let n be a prime, m a divisor of $n-1$, m' a divisor of m, and $d = \frac{m}{m'}$. Let $A' = \{a'_1, ..., a'_{m'}\}$ be the unique subgroup of $H := (\mathbb{Z}/n\mathbb{Z})^\times$ of order m'. Set $\omega = e^{2\pi i/n}$ and $\mathbf{v} = \frac{1}{\sqrt{m}}[1, ..., 1]^T \in \mathbb{C}^{m \times 1}$, and as before choose a set of cosets $\{\ell_1 A', ..., \ell_d A'\}$ of A' in H. Form the matrices $\mathbf{D}_\ell = \text{diag}\left(\omega^{\ell a'_1}, ..., \omega^{\ell a'_{m'}}\right) \in \mathbb{C}^{m' \times m'}$ for any $\ell \in H$, and $\mathbf{U} := \text{diag}(\mathbf{D}_{\ell_1}, \mathbf{D}_{\ell_2}, ..., \mathbf{D}_{\ell_d}) \in \mathbb{C}^{m \times m}$. If $\mathbf{M}_1 = [\mathbf{v}, \mathbf{Uv}, ..., \mathbf{U}^{n-1}\mathbf{v}] \in \mathbb{C}^{m \times n}$ has coherence μ_1 and $\mathbf{M}_2 = [\mathbf{v}, \mathbf{D}_1\mathbf{v}, ..., \mathbf{D}_1^{n-1}\mathbf{v}] \in \mathbb{C}^{m' \times n}$ has coherence μ_2, then we have $\mu_1 \leq \mu_2$.*

Proof. The result comes from a simple application of the triangle inequality: from Equation (3.29),

Figure 3.4: Randomly sampling ℓ to search for the optimal coherence over cosets of subgroups of size m' for various values of m' ($= 1, 14, 70$). (Plot shows the lowest coherence found up to a given iteration). Here, $n = 491$, and $m = 70$. The figure shows that $m' = 14$ quickly achieves the lowest values of coherence.

we see that

$$\mu_1 = \max_b \frac{1}{m} \left| \sum_{a' \in A'} \omega^{\ell_1(ba')} + ... + \sum_{a' \in A'} \omega^{\ell_d(ba')} \right| \tag{3.31}$$

$$\leq \frac{1}{dm'} \left(\max_{b_1} \left| \sum_{a' \in A'} \omega^{(\ell_1 b_1)a'} \right| + ... + \max_{b_d} \left| \sum_{a' \in A'} \omega^{(\ell_d b_d)a'} \right| \right) \tag{3.32}$$

$$= \frac{1}{d} \left(d \cdot \max_s \left| \frac{1}{m'} \sum_{a' \in A'} \omega^{sa'} \right| \right) \tag{3.33}$$

$$= \mu_2. \tag{3.34}$$

$\square$

Theorem 12 naturally allows us to use the bounds from Theorems 9 and 10 to explicitly bound the coherence from our coset optimization in terms of r, m and d, though it is worth noting that in practice we achieve coherence significantly lower than these bounds.

3.8   Generalized Dihedral Groups

Let us now investigate what changes when $\mathcal{U}$ is nonabelian. In this case the irreducible representations at our disposal will no longer all be one-dimensional, so we will no longer have all the matrices $\mathbf{U}_i$ be simultaneously diagonal. Consequently, it may not be possible to write all of our inner products

in the simple form of Equation (3.15), so it is no longer clear that we can restrict our vector $\mathbf{v}$ to be real-valued.

One simple class of nonabelian groups is that of semidirect products of cyclic groups. On this note, consider the following group presentation (which arises in [88]):

$$G_{n,r} = \langle \sigma, \tau \mid \sigma^n = 1, \tau^D = 1, \tau\sigma\tau^{-1} = \sigma^r \rangle. \tag{3.35}$$

Here, n and r are relatively prime integers, and D is the multiplicative order of r modulo n. $G_{n,r}$ is precisely a semidirect product in the form $\frac{\mathbb{Z}}{n\mathbb{Z}} \rtimes \frac{\mathbb{Z}}{D\mathbb{Z}}$, and if we take $D = 2$ and $r = n - 1$, we see that we obtain the familiar dihedral group D_{2n}.

There are $n \cdot D$ group elements in $G_{n,r}$, each of which can be written in the form $\sigma^{\ell_1}\tau^{\ell_2}$ for some integers $0 \leq \ell_1 < n$ and $0 \leq \ell_2 < D$. $G_{n,r}$ has an irreducible representation in the form

$$\sigma \mapsto \mathbf{S} := \begin{bmatrix} \omega & & & \\ & \omega^r & & \\ & & \ddots & \\ & & & \omega^{r^{D-1}} \end{bmatrix} \in \mathbb{C}^{D \times D}, \tag{3.36}$$

$$\tau \mapsto \mathbf{T} := \begin{bmatrix} & 1 & & & \\ & & 1 & & \\ & & & \ddots & \\ & & & & 1 \\ 1 & & & & \end{bmatrix} \in \mathbb{C}^{D \times D}, \tag{3.37}$$

where $\omega = e^{\frac{2\pi i}{n}}$ (see again [88]). The informed reader might note that this representation is quite similar to that of Heisenberg groups, which have been extensively applied to the construction of frames [9, 63, 84, 86]. Our following methods can be conceivably adjusted for use with Heisenberg frames as well.

In order to construct our frames, we would like to follow the example of our previous construction in Theorem 6 by selecting a representation for $G_{n,r}$ of the form

$$\sigma \mapsto [\sigma] := \begin{bmatrix} \mathbf{S}^{a_1} & & \\ & \ddots & \\ & & \mathbf{S}^{a_m} \end{bmatrix}, \ \tau \mapsto [\tau] := \begin{bmatrix} \mathbf{T} & & \\ & \ddots & \\ & & \mathbf{T} \end{bmatrix}, \tag{3.38}$$

where m and the a_i are cleverly chosen integers. Then we will select a vector $\mathbf{v} \in \mathbb{C}^{Dm \times 1}$ and take our

frame to be the columns of the matrix $\mathbf{M} := [\dots [\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v}\dots]_{0\leq \ell_1 < n,\ 0\leq\ell_2 < D}$. We need the greatest common divisor between the a_i to be relatively prime to n in order for the columns to be distinct, and again we satisfy this by taking n to be prime. Note that in our above notation, this will be a Dm-dimensional representation, so our resulting frame matrices will have dimensions $Dm \times Dn$.

At this point, we can see that in order to minimize coherence we must deviate from our original construction, for if we were to set $\mathbf{v}$ to the vector $\mathbf{1}_{Dm}$ of all ones it would be fixed by $[\tau]^{\ell_2}$ for any ℓ_2, and the inner product corresponding to $[\tau]^{\ell_1}$ would be 1. We must therefore be more clever in how we construct $\mathbf{v}$. A natural form for $\mathbf{v}$ would be to find some D-dimensional vector $\mathbf{w} = [w_1, ..., w_D]^T \in \mathbb{C}^{D\times 1}$ and set $\mathbf{v}$ equal to the periodic vector $\mathbf{v} = \begin{bmatrix} \mathbf{w}^T & \mathbf{w}^T & ... & \mathbf{w}^T \end{bmatrix}^T \in \mathbb{C}^{Dm\times 1}$. The question now becomes how to choose $\mathbf{w}$?

In order to preserve as much of the structure from our previous construction as possible, we would like each entry of $\mathbf{w}$ to have the same norm. This will ensure that the inner products corresponding to the elements $[\sigma]^{\ell_1}$ will have the same values as those in our previous construction from Theorem 6 corresponding to when $\mathcal{U}$ was the cyclic group $\mathbb{Z}/n\mathbb{Z}$ generated by $[\sigma]$. Let us require that w_d be unit norm for each d, and consider attempting to force $\mathbf{w}$ to satisfy the constraint that

$$\mathbf{w}^*\mathbf{T}^{\ell_2}\mathbf{w} = \sum_d w_d^* w_{d+\ell_2} = 0, \ \forall \ell_2, \tag{3.39}$$

where the indices are taken modulo D. It turns out that we can satisfy all our restrictions on $\mathbf{w}$ by selecting its indices to form a *Zadoff-Chu* (ZC) sequence [26, 60]:

$$w_d = \begin{cases} e^{\frac{i\pi d^2}{D}}, & \text{if } D \text{ is even} \\ e^{\frac{i\pi d(d+1)}{D}}, & \text{if } D \text{ is odd} \end{cases}. \tag{3.40}$$

This is a well-known constant amplitude zero autocorrelation (CAZAC) sequence. Our frame elements will now take the form

$$[\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v} = \begin{bmatrix} \mathbf{S}^{\ell_1 a_1}\mathbf{w}_{d+\ell_2} \\ \vdots \\ \mathbf{S}^{\ell_1 a_m}\mathbf{w}_{d+\ell_2} \end{bmatrix}, \tag{3.41}$$

where $\mathbf{w}_{d+\ell_2} := \mathbf{T}^{\ell_2}\mathbf{w}$ denotes the vector obtained by cyclically shifting the entries of $\mathbf{w}$ by ℓ_2 positions. Thus, as the notation would suggest, the d^{th} entry of $\mathbf{w}_{d+\ell_2}$ is $w_{d+\ell_2}$. (Note that by this notation,

$\mathbf{w}_d$ is simply $\mathbf{w}$ itself). Our inner products will take the form

$$\frac{\mathbf{v}^*[\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v}}{||\mathbf{v}||_2^2} = \frac{1}{m \cdot D} \sum_{j=1}^{m} \mathbf{w}_d^* \mathbf{S}^{\ell_1 a_j} \mathbf{w}_{d+\ell_2}. \qquad (3.42)$$

Our new frames remain tight:

Theorem 13. *Let n and r be relatively prime integers, and D the order of r modulo n. Let $[\sigma] \in \mathbb{C}^{Dm \times Dm}$ and $[\tau] \in \mathbb{C}^{Dm \times Dm}$ be the generating matrices for $G_{n,r}$ defined in (3.37) and (3.38). If $\mathbf{w} = [w_1, ..., w_D]^T \in \mathbb{C}^{D \times 1}$ is a ZC-sequence (3.40), and $\mathbf{v} = \begin{bmatrix} \mathbf{w}^T & ... & \mathbf{w}^T \end{bmatrix}^T \in \mathbb{C}^{Dm \times 1}$, then the columns of $\mathbf{M} = \begin{bmatrix} ... & [\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v} & ... \end{bmatrix} \in \mathbb{C}^{Dm \times Dn}$ form a tight frame.*

Proof. This be deduced from Theorem 5.4 of [17] since all the representations are of the same dimension and the corresponding components $\mathbf{w}_d$ of $\mathbf{v}$ all have the same norm. But we will give a self-contained proof here for completeness. It is not too difficult to see that $\mathbf{M}$ will have $D \cdot m$ rows which can be indexed by a pair of numbers (d, j), where $1 \le d \le D$ and $1 \le j \le m$. Row (d, j) will be given by $\begin{bmatrix} \mathbf{z}_1^{(d,j)} & \mathbf{z}_2^{(d,j)} & ... & \mathbf{z}_D^{(d,j)} \end{bmatrix}$, where $\mathbf{z}_{\ell_2+1}^{(d,j)} = \begin{bmatrix} ... & \omega^{r^{d-1}\ell_1 a_j} w_{d+\ell_2} & ... \end{bmatrix}_{0 \le \ell_1 < n}$.

Now we can see that the inner product between row (d, j) and row (d', j') will be

$$\sum_{\ell_2=0}^{D-1} \sum_{\ell_1=0}^{n-1} \omega^{(-r^{d-1}a_j + r^{d'-1}a_{j'})\ell_1} w_{d+\ell_2}^* w_{d'+\ell_2}$$

$$= \left[\sum_{\ell_1=0}^{n-1} \omega^{(-r^{d-1}a_j + r^{d'-1}a_{j'})\ell_1} \right] \cdot \left[\sum_{\ell_2=0}^{D-1} w_{d+\ell_2}^* w_{d'+\ell_2} \right]. \qquad (3.43)$$

Since the entries of $\mathbf{w}$ form a ZC-sequence, the sum $\sum_{\ell_2=0}^{D-1} w_{d+\ell_2}^* w_{d'+\ell_2}$ is equal to zero unless $d = d'$, in which case it is equal to D. In this latter case we have,

$$\sum_{\ell_1=0}^{n-1} \omega^{(-r^{d-1}a_j + r^{d'-1}a_{j'})\ell_1} = \sum_{\ell_1=0}^{n-1} \omega^{r^{d-1}(a_{j'} - a_j)\ell_1}, \qquad (3.44)$$

which is zero unless $j = j'$. Thus, the rows of $\mathbf{M}$ are indeed orthogonal, and of equal norm, so the frame is tight. $\qquad \square$

Exploiting the properties of our construction, we can bound the coherence of our new frames by that of our original frames arising from representations of cyclic groups.

Theorem 14. *Let n be an integer, and $a_1, ..., a_m$ distinct integers modulo n whose greatest common divisor is relatively prime to n. Take r an integer relatively prime to n, and D the multiplicative order*

of r modulo n. Set $\omega = e^{\frac{2\pi i}{n}}$. Consider the two frames:

1. *The columns of the "cyclic frame" $\mathbf{M}_1 = [\mathbf{v}_1, \mathbf{U}\mathbf{v}_1, \ldots, \mathbf{U}^{n-1}\mathbf{v}_1] \in \mathbb{C}^{m \times n}$ where $\mathbf{U} = \mathrm{diag}(\omega^{a_1}, \ldots, \omega^{a_m}) \in \mathbb{C}^{m \times m}$ and $\mathbf{v}_1 = [1, \ldots, 1]^T \in \mathbb{C}^{m \times 1}$.*

2. *The columns of the "generalized dihedral frame" $\mathbf{M}_2 = \begin{bmatrix} \ldots & [\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v}_2 & \ldots \end{bmatrix} \in \mathbb{C}^{Dm \times Dn}$ where $\mathbf{v}_2 = \begin{bmatrix} \mathbf{w}^T & \ldots & \mathbf{w}^T \end{bmatrix}^T \in \mathbb{C}^{Dm \times 1}$ and $\mathbf{w} = [w_1, \ldots, w_D]^T \in \mathbb{C}^{D \times 1}$ is a ZC-sequence.*

If μ_A^{cyc} is the coherence of the cyclic frame $\mathbf{M}_1$ and μ_A^D the coherence of the generalized dihedral frame $\mathbf{M}_2$, then $\mu_A^D \leq \mu_A^{cyc}$.

Proof. From (3.42), we see that the inner products for the generalized dihedral representation will take the form

$$\frac{\mathbf{v}^*[\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v}}{||\mathbf{v}||_2^2} = \frac{1}{m \cdot D} \sum_{a \in A} \sum_d w_d^* w_{d+\ell_2} \omega^{a\ell_1 r^{d-1}} \tag{3.45}$$

$$= \frac{1}{m \cdot D} \sum_d w_d^* w_{d+\ell_2} \sum_{a \in A} \omega^{a\ell_1 r^{d-1}} \tag{3.46}$$

$$= \frac{1}{m \cdot D} \sum_d w_d^* w_{d+\ell_2} \sum_{a \in A} \omega^{a\ell'}, \tag{3.47}$$

where $\ell' = \ell_1 r^{d-1}$. Furthermore,

$$\frac{|\mathbf{v}^*[\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v}|}{||\mathbf{v}||_2^2} \leq \frac{1}{m \cdot D} \sum_d \left| w_d^* w_{d+\ell_2} \sum_{a \in A} \omega^{a\ell'} \right| \tag{3.48}$$

$$= \frac{1}{m \cdot D} \sum_d \left| \sum_{a \in A} \omega^{a\ell'} \right| \tag{3.49}$$

$$\leq \frac{1}{m \cdot D} \sum_d m \mu_A^{cyc} = \mu_A^{cyc}, \tag{3.50}$$

so $\mu_A^D \leq \mu_A^{cyc}$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Theorem 14 allows us to bound the coherence of our generalized dihedral frames using the same bounds from Theorems 9 and 10:

Corollary 2. *Let n be a prime and m a divisor of $n - 1$, and let $A = \{a_1, \ldots, a_m\}$ be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m. Set $\kappa = \frac{n-1}{m}$. Take r an integer relatively prime to n, and D the multiplicative order of r modulo n.*

Let $[\sigma] \in \mathbb{C}^{Dm \times Dm}$ and $[\tau] \in \mathbb{C}^{Dm \times Dm}$ be the generating matrices for $G_{n,r}$ defined in (3.37) and (3.38). If $\mathbf{w} = [w_1, \ldots, w_D]^T \in \mathbb{C}^{D \times 1}$ is a ZC-sequence (3.40), and $\mathbf{v} = \begin{bmatrix} \mathbf{w}^T & \ldots & \mathbf{w}^T \end{bmatrix}^T \in \mathbb{C}^{Dm \times 1}$,

then the the columns of $\mathbf{M} = \begin{bmatrix} \ldots & [\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v} & \ldots \end{bmatrix} \in \mathbb{C}^{Dm \times Dn}$ *have at most* $D \cdot \frac{n-1}{m}$ *distinct inner product values between them, and the coherence* μ *of* $\mathbf{M}$ *is bounded by*

$$\mu \leq \frac{1}{\kappa}\left((\kappa-1)\sqrt{\frac{1}{m}\left(\kappa+\frac{1}{m}\right)}+\frac{1}{m}\right). \tag{3.51}$$

If m *is odd, then the coherence of* $\mathbf{M}$ *is upper-bounded by*

$$\mu \leq \frac{1}{\kappa}\sqrt{\left(\frac{1}{m}+\left(\frac{\kappa}{2}-1\right)\beta\right)^2+\left(\frac{\kappa}{2}\right)^2\beta^2}, \tag{3.52}$$

where $\beta = \sqrt{\frac{1}{m}\left(\kappa+\frac{1}{m}\right)}$.

Proof. From (3.47), we can write out the inner product corresponding to the group element $\sigma^{\ell_1}\tau^{\ell_2}$ in the form

$$\frac{\mathbf{v}^*[\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v}}{||\mathbf{v}||_2^2} = \frac{1}{m \cdot D}\sum_d w_d^* w_{d+\ell_2}\sum_{a \in A}\omega^{a\ell'}, \tag{3.53}$$

where $\ell' = \ell_1 r^{d-1}$. In this form, we see that for each value of d in the summation, there are $\frac{n-1}{m}$ possible distinct inner product values associated to the different cosets $\ell'A$, so there are at most $D\frac{n-1}{m}$ possible values. The last two bounds (3.51) and (3.52) follow from Theorem 14 and the bounds given in Theorems 9 and 10. $\square$

In the case of regular dihedral groups ($D = 2$), our $\mathbf{w}$ becomes $[1, i]^T$, and we can readily calculate our inner products to be

$$\frac{\mathbf{v}^*[\sigma]^{\ell}\mathbf{v}}{||\mathbf{v}||_2^2} = \text{Re}\left(\frac{1}{m}\sum_{j=1}^m \omega^{\ell a_j}\right),$$

$$\frac{\mathbf{v}^*[\sigma]^{\ell}[\tau]\mathbf{v}}{||\mathbf{v}||_2^2} = \text{Im}\left(-\frac{1}{m}\sum_{j=1}^m \omega^{\ell a_j}\right).$$

As we can clearly see, each of these has magnitude bounded by that of the corresponding inner product in the cyclic counterpart, $\left|\frac{1}{m}\sum_{j=1}^m \omega^{\ell a_j}\right|$. In general, the dihedral coherence could be substantially smaller than the corresponding cyclic coherence. Most importantly, by extending to generalized dihedral groups, we allow for frame matrices $\mathbf{M}$ with a greater variety of dimensions. In particular, the number of columns (nD) no longer need be prime. In Figure 3.5, we plot the coherences arising from these frames along with the upper bounds predicted by Corollary 2 for the case $\kappa = 4$. From this figure, it becomes apparent that in practice the coherence of our frames significantly outperforms

these bounds.

Figure 3.5: Coherences arising from dihedral representations for $\kappa = 4$, which we show can also be realized by abelian representations. We also plot the upper bounds from Theorems 9 and 10.

3.8.1 Simulating Generalized Dihedral Frames with Harmonic Frames

It turns out that in this case, we could have created frames with the same dimensions whose inner products have exactly the same magnitudes as the those of the above generalized dihedral frames had we replaced $\mathbf{T}$ and $[\tau]$ with

$$\mathbf{T}' := \text{diag}(1, \gamma, \gamma^2, ..., \gamma^{D-1}) \in \mathbb{C}^{D \times D} \tag{3.54}$$

$$[\tau'] := \text{diag}(\mathbf{T}', ..., \mathbf{T}') \in \mathbb{C}^{Dm \times Dm}, \tag{3.55}$$

where $\gamma = e^{\frac{2\pi i}{D}}$, and replaced $\mathbf{v}$ with $\mathbf{v}' = \frac{1}{\sqrt{Dm}} \begin{bmatrix} 1 & ... & 1 \end{bmatrix}^T \in \mathbb{C}^{Dm \times 1}$, the vector of all ones. Here we have altered $\mathbf{T}$ to be a diagonal matrix, $\mathbf{T}'$, but maintain the property that $[\tau']$ is a block diagonal matrix with m copies of $\mathbf{T}'$ on the diagonal. Together with $[\sigma]$, this is no longer a representation of a generalized dihedral group, but rather a representation of the abelian group $\frac{\mathbb{Z}}{n\mathbb{Z}} \times \frac{\mathbb{Z}}{D\mathbb{Z}}$. The group elements again take the form $[\sigma]^{\ell_1} [\tau']^{\ell_2}$, where $0 \le \ell_1 \le n - 1$ and $0 \le \ell_2 \le D - 1$.

The resulting frame

$$\mathbf{M} = \begin{bmatrix} \ldots & [\sigma]^{\ell_1} [\tau']^{\ell_2} \mathbf{v}' & \ldots \end{bmatrix}$$

is harmonic and therefore tight by Lemma 8, and it is not too difficult to see that the inner product

corresponding to $[\sigma]^{\ell_1}[\tau']^{\ell_2}$ will be

$$\frac{\mathbf{v}'^*[\sigma]^{\ell_1}[\tau']^{\ell_2}\mathbf{v}'}{||\mathbf{v}'||_2^2} = \frac{1}{m \cdot D} \sum_d \gamma^{d\ell_2} \sum_{a \in A} \omega^{a\ell_1 r^{d-1}}. \tag{3.56}$$

We would like to compare (3.56) to (3.47). A quick calculation shows that when the terms w_d in (3.47) are chosen to be a ZC sequence as in (3.40), we have

$$w_d^* w_{d+\ell} = \begin{cases} e^{\frac{i\pi}{D}(2d\ell+\ell^2)} = \gamma^{d\ell} e^{\frac{i\pi\ell^2}{D}} & \text{if } D \text{ is even,} \\ e^{\frac{i\pi}{D}(2d\ell+\ell^2+\ell)} = \gamma^{d\ell} e^{\frac{i\pi(\ell^2+\ell)}{D}} & \text{if } D \text{ is odd.} \end{cases} \tag{3.57}$$

We can now see from (3.56) and (3.47) that when choosing $\mathbf{v}$ and $[\tau]$ as in Theorem 13, we have

$$\frac{\mathbf{v}^*[\sigma]^{\ell_1}[\tau]^{\ell_2}\mathbf{v}}{||\mathbf{v}||_2^2} = \begin{cases} e^{\frac{i\pi\ell_2^2}{D}}\left(\frac{\mathbf{v}'^*[\sigma]^{\ell_1}[\tau']^{\ell_2}\mathbf{v}'}{||\mathbf{v}'||_2^2}\right) & \text{if } D \text{ is even,} \\ e^{\frac{i\pi(\ell_2^2+\ell_2)}{D}}\left(\frac{\mathbf{v}'^*[\sigma]^{\ell_1}[\tau']^{\ell_2}\mathbf{v}'}{||\mathbf{v}'||_2^2}\right) & \text{if } D \text{ is odd,} \end{cases} \tag{3.58}$$

so indeed the inner products from our two frames have the same norm.

3.9   Summary

In this chapter, we have presented a method to select a set of representations of a finite cyclic group to construct tight, unit-norm group frames such that the frame elements take on very few distinct pairwise inner product values. Our construction ensures that each such inner product value arises the same number of times, allowing us to derive upper bounds on the coherence of the frames which approach the Welch lower bound. In certain cases, our construction has yielded instances of previously known tight, equiangular frames which achieve the Welch bound. We have then demonstrated how our method can be applied to constructing tight group frames from abelian and generalized dihedral groups to obtain a richer set of frames of different sizes and dimensions. We have derived similar bounds on coherence in these situations. In the Chapter 4, we will realize our method in a more general context, showing how to choose representations of a general group to construct group frames. We will develop a general framework which will tie all of our previous constructions together, and it will become apparent why our cyclic group construction extends so naturally to generalized dihedral groups. Furthermore, we will identify other groups for which our method produces frames with particularly low coherence, including certain other tight, equiangular frames.

Chapter 4

Frames from Generalized Fourier Matrices

4.1 Tight Group Frames and the Group Fourier Matrix

In light of Lemmas 5 and 6, we will first establish the tools we need to ensure that our group frames are tight. It turns out that the tight group frames have been completely classified [17, 102]. On this note, we review some basics on representation theory, which the interested reader can read about in greater depth in the first few chapters of [87].

Let G be a group of size n, and recall that a complex representation of G is formally defined as a complex vector space V together with a function $\rho : G \to GL(V)$ such that $\rho(gg') = \rho(g)\rho(g')$, $\forall g, g' \in G$. If V has dimension d, then $\rho(g)$ is simply a $d \times d$ invertible complex matrix—a *degree* d representation. Two representations ρ_1 and ρ_2 with corresponding vector spaces V_1 and V_2 are *equivalent* if there is an invertible transformation $T : V_1 \to V_2$ such that $T\rho_1(g)T^{-1} = \rho_2(g)$ for all $g \in G$. A basic result in representation theory says that every representation of a finite group is equivalent to a unitary representation, in which all the $\rho(g_i)$ are unitary matrices, which is why we have used the notation $\rho(g_i) = \mathbf{U}_i$ in our previous discussion. We will typically assume our representations are unitary without loss of generality.

A representation ρ is *reducible* if there is a nontrivial subspace V' of V which is mapped to itself by $\rho(g)$ for every $g \in G$. Otherwise, it is called *irreducible*. As matrices, the representation is reducible if the $\rho(g)$ can be simultaneously block-diagonalized by a similarity transformation. For any finite group G of size n, there are only a finite number of inequivalent, irreducible unitary representations. If we call them $\rho_1, ..., \rho_{n_r}$ with corresponding degrees $d_1, ..., d_{n_r}$, then it can be shown [87] that these

degrees satisfy the relation

$$\sum_{i=1}^{n_r} d_i^2 = |G|. \tag{4.1}$$

Every complex representation of G is equivalent to an orthogonal direct sum of irreducible representations. Formally, this means that there is an invertible linear transformation $T : V \to V_1 \oplus ... \oplus V_m$ such that the V_i are mutually orthogonal vector spaces and $T\rho(g)T^{-1} = \rho_1(g) \oplus ... \oplus \rho_m(g)$, where for each i, ρ_i and V_i give an irreducible representation of G. These irreducible representations can again be taken to be unitary. As matrices, this means that the $\rho(g)$ can be simultaneously block-diagonalized in the form $\rho(g) = \text{diag}(\rho_1(g), ..., \rho_m(g))$. A basic result of representation theory is that this decomposition into irreducible components is unique up to isomorphism. We are now ready to give a classification of all the tight G-frames:

Theorem 15 ([17]). *Let $G = \{g_i\}_{i=1}^n$ be a finite group, and $\rho : G \to GL(V)$ a complex representation of G which has the decomposition into orthogonal unitary irreducible representations:*

$$V = V_1 \oplus ... \oplus V_m,$$

$$\rho(g) = \rho_1(g) \oplus ... \oplus \rho_m(g).$$

Let $\mathbf{v} = \mathbf{v}_1 + ... + \mathbf{v}_m$, $\mathbf{v}_k \in V_k$, and set $\mathbf{f}_i = \rho(g_i)\mathbf{v}$. Then $\{\mathbf{f}_i\}_{i=1}^n$ is a tight G-frame if and only if

- $\frac{||\mathbf{v}_i||_2^2}{||\mathbf{v}_j||_2^2} = \frac{dim(V_i)}{dim(V_j)}$, *and*

- *if the i^{th} and j^{th} irreducible components are equivalent via $T : V_i \to V_j$, then $T\mathbf{v}_i$ and $\mathbf{v}_j$ are orthogonal.*

Proof. This is Theorem 5.4 in [17]. It follows from considering the frame matrix $\mathbf{M} := \begin{bmatrix} ... & \rho(g_i)\mathbf{v} & ... \end{bmatrix}_{i=1}^n$ and applying Schur's Lemma (Section 2.2, [87]) to the product $\mathbf{M}\mathbf{M}^*$ to see when it is a scalar matrix, which is equivalent to the columns of $\mathbf{M}$ forming a tight frame. $\square$

We will now establish a tool that will allow us to easily use this theorem to construct tight frames. On this note, consider the following well-studied generalization of the classical discrete Fourier transform [94]:

Definition 4. *We define the group Fourier transform of a complex-valued function on G, $f : G \to \mathbb{C}$, to be the function that maps a degree d representation ρ to the $d \times d$ complex matrix*

$$\hat{f}(\rho) = \sum_{g \in G} f(g)\rho(g). \tag{4.2}$$

There is an inverse transformation given by

$$f(g) = \frac{1}{|G|} \sum_{i=1}^{n_r} d_i \text{Tr}(\rho_i(g^{-1})\hat{f}(\rho_i)), \tag{4.3}$$

where the sum is taken over all the inequivalent irreducible representations of G.

Much like the traditional discrete Fourier transform, this transformation has a matrix representation in the form

$$\mathcal{F} = \begin{bmatrix} \sqrt{d_1}\text{vec}(\rho_1(g_1)) & \cdots & \sqrt{d_1}\text{vec}(\rho_1(g_n)) \\ \sqrt{d_2}\text{vec}(\rho_2(g_1)) & \cdots & \sqrt{d_2}\text{vec}(\rho_2(g_n)) \\ \vdots & \ddots & \vdots \\ \sqrt{d_{n_r}}\text{vec}(\rho_{n_r}(g_1)) & \cdots & \sqrt{d_{n_r}}\text{vec}(\rho_{n_r}(g_n)) \end{bmatrix}, \tag{4.4}$$

where for a $d \times d$ matrix A, $\text{vex}(A)$ is the *vectorization* of A, i.e., the vector formed by stacking the columns of A into a single $d^2 \times 1$ column. From equation (4.1), we see that $\mathcal{F}$ is a square matrix.

Notice that when G is a cyclic group of size n, then the group elements are $\{0, 1, ..., n-1\}$ (with the group operation being addition modulo n). Since this group is abelian, there are exactly n irreducible representations, $\{\rho_\ell\}_{\ell=1}^n$, each degree-1. ρ_ℓ is simply the function that maps $k \mapsto \omega^{k\ell}$, $k \in \{0, ..., n-1\}$, where $\omega = e^{\frac{2\pi i}{n}}$. In this case our group Fourier transform and matrix become the familiar discrete time Fourier transform and DFT matrix.

Theorem 16. *Let $G = \{g_i\}_{i=1}^n$ be a finite group with inequivalent, irreducible representations $\{\rho_i\}_{i=1}^{n_r}$, and $\mathcal{F}$ the group Fourier matrix of G as in (4.4). Then the columns of $\mathcal{F}$ form a tight G-frame, so $\mathcal{F}$ is a unitary matrix. In fact, if $\tilde{\rho}: G \to \mathbb{C}^{d \times d}$ is a representation of G and $\tilde{\mathbf{v}} \in \mathbb{C}^{d \times 1}$ such that the columns of $\mathbf{M} := [\ldots \ \tilde{\rho}(g_i)\tilde{\mathbf{v}} \ \ldots]_{i=1}^n$ form a tight frame, then the rows of $\mathbf{M}$ are a subset of the rows of $\mathcal{F}$ up to an equivalence of $\tilde{\rho}$ or a change of basis of $\mathbb{C}^{d \times 1}$.*

Proof. The group Fourier matrix $\mathcal{F}$ can be realized as a G-frame as follows: For each $i = 1, ..., n_r$, define the representation

$$\tilde{\rho}_i(g) = \text{diag}(\rho_i(g), ..., \rho_i(g)) \in \mathbb{C}^{d_i^2 \times d_i^2}, \tag{4.5}$$

a direct sum of d_i copies of the irreducible representation ρ_i. Also define the vector

$$\mathbf{v}_i = \text{vec}(\mathbf{I}_{d_i}) = [\mathbf{e}_i^{(1)T}, ..., \mathbf{e}_i^{(d_i)T}]^T \in \mathbb{C}^{d_i^2}, \tag{4.6}$$

where $\mathbf{I}_{d_i}$ is the $d_i \times d_i$ identity matrix and $\mathbf{e}_i^{(j)} \in \mathbb{C}^{d_i \times 1}$ is the j^{th} column of $\mathbf{I}_{d_i}$—a vector of all zeros

except for a 1 in the j^{th} position.

Now choose the representation

$$\rho(g) = \text{diag}(\tilde{\rho}_1(g), \tilde{\rho}_2(g), ..., \tilde{\rho}_{n_r}(g)) \in \mathbb{C}^{n \times n}, \tag{4.7}$$

and the vector

$$\mathbf{v} = [\sqrt{d_1}\mathbf{v}_1^T, \sqrt{d_2}\mathbf{v}_2^T, \ldots, \sqrt{d_{n_r}}\mathbf{v}_{n_r}^T]^T \in \mathbb{C}^n. \tag{4.8}$$

Then $\mathcal{F}$ is the G-frame with columns $\rho(g_i)\mathbf{v}$. For any i, $\{\mathbf{e}_i^{(j)}\}_{j=1}^{d_i}$ is a complete orthonormal set in $\mathbb{C}^{d_i}$, and $\frac{||e_{i_1}^{(j_1)}||_2^2}{||e_{i_2}^{(j_2)}||_2^2} = \frac{d_{i_1}}{d_{i_2}}$. From Theorem 15, we see that not only do the columns of $\mathcal{F}$ form a tight G-frame, but in fact up to a change of basis of the $\mathbf{e}_i^{(j)}$ or a similarity transformation of the ρ_i, every tight G-frame can be realized as a subset of the rows of $\mathcal{F}$ by forming each $\mathbf{v}_i$ from a corresponding *subset* of the columns $\{\mathbf{e}_i^{(j)}\}_{j=1}^{d_i}$. $\qquad\qquad\square$

Theorem 16 reduces the task of constructing tight G-frames to selecting blocks of rows of the corresponding group Fourier matrix $\mathcal{F}$. Our job will now be to find good choices of the group G, and to identify which rows of $\mathcal{F}$ to choose to create a tight group frame with low coherence. We should mention that this problem was explored for abelian groups G in [35], with a focus on finding frames with coherence equal to the Welch Bound. We will find, however, that by not placing any restrictions on our group G, and by allowing our coherence to be slightly above the Welch lower bound, we can produce a vastly larger and richer collection of frames.

4.2 Reducing the Number of Distinct Inner Products in Tight Group Frames

In our original construction from Theorem 6, we designed harmonic frames in the form of $\mathbf{M}$ from (3.18) which arose from representations of the cyclic group $G = \mathbb{Z}/n\mathbb{Z}$, where n is a prime. Indeed, the j^{th} row of $\mathbf{M}$ is $[1, \omega^{k_j}, \omega^{2k_j}, ..., \omega^{(n-1)k_j}$, where $\omega = e^{\frac{2\pi i}{n}}$, and we can now see that this is simply the row of the group Fourier matrix of G corresponding to the n-dimensional representation $\rho_{k_j}(\ell) = \omega^{\ell k_j}$, for $\ell \in \{0, ..., n-1\}$. We wish to generalize our original method from Theorem 6 of constructing frames with few distinct inner product values.

On this note, we will consider constructing frames by choosing the blocks of rows corresponding to m of the representations, which we may assume are $\rho_1, ..., \rho_m$ up to a reordering, so that our frame

matrix takes the form

$$\mathbf{M} = \begin{bmatrix} \sqrt{d_1}\text{vec}(\rho_1(g_1)) & \cdots & \sqrt{d_1}\text{vec}(\rho_1(g_n)) \\ \vdots & \ddots & \vdots \\ \sqrt{d_m}\text{vec}(\rho_m(g_1)) & \cdots & \sqrt{d_m}\text{vec}(\rho_m(g_n)) \end{bmatrix}. \tag{4.9}$$

As an analog to Equations (4.7) and (4.8) from the proof of Theorem 16, this corresponds to the tight group frame whose elements are the images of the vector $\mathbf{v} = [\sqrt{d_1}\mathbf{v}_1^T, \sqrt{d_2}\mathbf{v}_2^T, \ldots, \sqrt{d_m}\mathbf{v}_m^T]^T$ under the representation $\rho(g) = \text{diag}(\tilde{\rho}_1(g), \tilde{\rho}_2(g), ..., \tilde{\rho}_m(g))$, where $\tilde{\rho}_i$ and $\mathbf{v}_i$ are defined as in Equations (4.5) and (4.6) respectively. The dimension of this representation is easily seen to be $\sum_{i=1}^m d_i^2$. Note that in the setting of Theorem 6, the representations ρ_i are all 1-dimensional, so the block $\begin{bmatrix} \sqrt{d_i}\text{vec}(\rho_i(g_1)) & \cdots & \sqrt{d_i}\text{vec}(\rho_i(g_n)) \end{bmatrix}$ is just a single row.

The inner product between the i^{th} and j^{th} columns of $\mathbf{M}$ in (4.9) takes the form

$$\sum_{t=1}^m d_t \text{vec}(\rho_t(g_i))^* \text{vec}(\rho_t(g_j)) = \sum_{t=1}^m d_t \text{Tr}(\rho_t(g_i)^* \rho_t(g_j)) \tag{4.10}$$

$$= \sum_{t=1}^m d_t \text{Tr}(\rho_t(g_i^{-1} g_j)) \tag{4.11}$$

$$= \sum_{t=1}^m d_t \chi_t(g_i^{-1} g_j). \tag{4.12}$$

Here, $\chi_i(g) := \text{Tr}(\rho_i(g))$ is the *character function* associated to the representation ρ_i. Equation (4.12) actually arises in [35], though only 1-dimensional representations are considered, in which case each representation is essentially just its own character. Note that in this form the frame is unnormalized, but all of the columns have the same norm, which is given by the square root of the inner product associated to the identity element:

$$||\rho(g)\mathbf{v}||_2 = \sqrt{\sum_{t=1}^m d_t \chi_t(1)} = \sqrt{\sum_{t=1}^m d_t^2}, \tag{4.13}$$

where we have used the fact the character evaluated at 1 is simply the dimension of the representation. Alternatively, we could have simply seen this to be the norm of $\mathbf{v}$ by speculation.

Basic representation theory tells us that a character χ completely determines its representation up to isomorphism, and as such the characters of many groups are well-studied. In light of this fact, we can often compute the coherence of frames in the form of (4.9) for different choices of representations $\{\rho_i\}_{i=1}^m$ without explicitly building the frame matrix $\mathbf{M}$, which can often be a tedious computation. From (4.11) and (4.12) we see that the inner product depends only on the group element $g_k := g_i^{-1} g_j$,

so a priori there are only $n-1$ possible nontrivial distinct inner product values, and each of these values arises the same number of times as the inner product between two columns. This was to be expected, since the columns of $\mathbf{M}$ form a group frame in light of Theorem 16. If we could generalize our method for choosing rows of the classical Fourier matrix, however, we could hope to reduce this number even further.

Toward this end, we consider the group of *automorphisms* of G. An automorphism of G is a bijective function $\sigma : G \to G$ which respects the group multiplication, i.e., $\sigma(gg') = \sigma(g)\sigma(g')$ for any $g, g' \in G$. The automorphisms of G form a group under composition, denoted $Aut(G)$. An important subgroup of $Aut(G)$ is that of the *inner automorphisms*, denoted $Inn(G)$. These are the automorphisms which arise from *conjugation* by an element $h \in G$, which is the function $\sigma_h(g) = hgh^{-1}$. Two elements g and g' are said to be *conjugate* if there is some $h \in G$ such that $g' = hgh^{-1}$, and the set of all elements conjugate to g is called the *conjugacy class* $\mathcal{C}_g$. We see that the relation $\{g \sim g' \iff g$ is conjugate to $g'\}$ is an equivalence relation on G, so G can be partitioned into a disjoint union of its conjugacy classes. $Inn(G)$ is easily verified to be a normal subgroup of $Aut(G)$, and the quotient group $Aut(G)/Inn(G)$ is called the group of *outer automorphisms*, denoted $Out(G)$.

Any conjugation $\sigma_h \in Inn(G)$ fixes a representation's character function. Indeed, if ρ is a representation of G with associated character χ, then

$$\chi(\sigma_h(g)) = \chi(hgh^{-1}) = \mathrm{Tr}(\rho(h)\rho(g)\rho(h)^{-1}) = \mathrm{Tr}(\rho(g)) = \chi(g). \tag{4.14}$$

Thus, since the inner products between the columns of $\mathbf{M}$ in (4.9) can be expressed as in (4.12) in terms of the characters of the irreducible representations of G (i.e., a so-called *character function* on the group elements), we see that there is really only one inner product value for each conjugacy class of G. Note that while this observation has the advantage of reducing the number of distinct inner product values to consider, we unfortunately cannot readily apply Lemma 6 to obtain a tighter coherence bound since these values no longer occur with the same multiplicity. Indeed, for each $g \in G$, the corresponding inner product value $\sum_{t=1}^{m} d_t \chi_t(g)$ will arise once for each element in the conjugacy class $\mathcal{C}_g$, and the conjugacy classes need not have the same size.

Since an automorphism essentially preserves the structure of the group G, it is no surprise that it also preserves the structure of its representations:

Lemma 9. $\rho(g)$ *is an irreducible representation of the finite group* G *if and only if* $\rho(\sigma(g))$ *is also an irreducible representation for any* $\sigma \in Aut(G)$. *Furthermore,* $\rho(g)$ *and* $\rho(\sigma(g))$ *have the same degrees.*

Proof. If $\rho : G \to GL(V)$ is a representation, then composing with the automorphism $\sigma : G \to G$ yields

a function $\rho \circ \sigma : G \to GL(V)$ which respects the group multiplication: $\rho(\sigma(gg')) = \rho(\sigma(g)\sigma(g')) = \rho(\sigma(g))\rho(\sigma(g'))$. Thus, $\rho(\sigma(g))$ is a well-defined representation which clearly has the same dimension as $\rho(g)$. Furthermore, since σ is a bijection of G, the matrices $\{\rho(\sigma(g)) : g \in G\}$ are simply a permutation of the matrices $\{\rho(g) : g \in G\}$, so the first representation is irreducible if and only if the second is. $\qquad\square$

If ρ is a representation with character χ, and $\sigma \in Aut(G)$, we will use the notation ρ_σ to indicate the representation

$$\rho_\sigma(g) := \rho(\sigma(g)), \tag{4.15}$$

which is irreducible if ρ is. ρ_σ has corresponding character

$$\chi_\sigma(g) := \chi(\sigma(g)). \tag{4.16}$$

Under this notation, if $\mathbf{1} \in Aut(G)$ denotes the identity automorphism $\mathbf{1}(g) = g$, then $\rho_{\mathbf{1}}$ and $\chi_{\mathbf{1}}$ are simply ρ and χ, respectively. From Lemma 9, we see that $Aut(G)$ has a group action on the irreducible representations and characters of G given by

$$\sigma' \cdot \rho_\sigma := \rho_{\sigma\sigma'}, \tag{4.17}$$

$$\sigma' \cdot \chi_\sigma := \chi_{\sigma\sigma'}. \tag{4.18}$$

Let us consider case in our original construction from Theorem 6 where G was the (additive) cyclic group $\mathbb{Z}/n\mathbb{Z} = \{0, ..., n-1\}$. In this case, $Aut(G)$ is isomorphic to the (multiplicative) group of elements relatively prime to n, $(\mathbb{Z}/n\mathbb{Z})^\times$. For each $\ell \in (\mathbb{Z}/n\mathbb{Z})^\times$, the corresponding automorphism $\sigma_\ell \in Aut(G)$ is given by $\sigma_\ell(g) = \ell g$. When we required that n be prime in Theorem 6, we ensured that every nonzero element had a multiplicative inverse modulo n, so in this case $(\mathbb{Z}/n\mathbb{Z})^\times$ is the set $\{1, ..., n-1\}$.

Refer back to the structure of our harmonic frame from (3.18):

$$\mathbf{M} = \frac{1}{\sqrt{m}} \begin{bmatrix} 1 & \omega^{a_1} & \omega^{a_1 \cdot 2} & \ldots & \omega^{a_1 \cdot (n-1)} \\ 1 & \omega^{a_2} & \omega^{a_2 \cdot 2} & \ldots & \omega^{a_2 \cdot (n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{a_m} & \omega^{a_m \cdot 2} & \ldots & \omega^{a_m \cdot (n-1)} \end{bmatrix}, \tag{4.19}$$

where $\omega = e^{2\pi i/n}$. As we have discussed, selecting the frequencies $\{a_1, ..., a_m\}$ is equivalent to choosing

rows of the group Fourier matrix corresponding to $\mathbb{Z}/n\mathbb{Z}$, each of which corresponds to a degree-1 representation. By choosing the frequencies $\{a_1, ..., a_m\}$ in (4.19) to be a subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ as in Theorem 6, we can now see that we are actually choosing a subgroup of $Aut(G)$. Without loss of generality, let $a_1 = 1$ so that the first row of (4.19) corresponds to the representation $\rho(g) = \omega^g$. Then the i^{th} row corresponds to the representation $\rho_i(g) := \rho(\sigma_{a_i}(g)) = \omega^{a_i g}$. Thus, we have formed $\mathbf{M}$ by choosing the rows of the group Fourier matrix corresponding to a subset of representations of the form $\{\sigma_i \cdot \rho\}$, where the $\{\sigma_i\}$ form a subgroup of automorphisms.

We wish to generalize this process to groups G other than $\mathbb{Z}/n\mathbb{Z}$ by choosing an irreducible representation ρ of G and taking its image under a subgroup of automorphisms $\{\sigma_i\} \leq Aut(G)$. Note that from Lemma 9, the representations $\{\sigma_i \cdot \rho\}$ will all be irreducible, and hence correspond to easily-identified blocks of rows from the group Fourier matrix $\mathcal{F}$ in (4.4). It is not clear, however, whether these representations will be distinct. The question now becomes how to choose the subgroup of automorphisms?

4.3 Choosing the Automorphism Subgroup

Let $H \leq Aut(G)$ be a group of automorphisms of G, and fix an irreducible representation ρ with character χ. Define K to be the subgroup of H which fixes χ:

$$K = \{\sigma \in H \ : \ \chi(\sigma(g)) = \chi(g), \ \forall g \in G\}. \tag{4.20}$$

Immediately we see that K contains every inner automorphism in H. Thus, it is effectively the group of *outer* automorphisms which acts nontrivially on the representations. Now choose a subgroup $A \leq H$ such that the group product $KA := \{ka \ : \ k \in K, \ a \in A\}$ is a subgroup of H. This is equivalent to the group products KA and AK being equal as sets. We consider choosing the rows of the generalized Fourier matrix corresponding to the representations $\{\rho_a \ : \ a \in A\}$, with notation as in (4.15). From Lemma 9, all of these representations have the same degree d. Thus, if $A = \{a_1, ..., a_m\} \leq Aut(G)$, then $\mathbf{M}$ takes the form

$$\mathbf{M} = \sqrt{d} \begin{bmatrix} \text{vec}(\rho_{a_1}(g_1)) & \dots & \text{vec}(\rho_{a_1}(g_n)) \\ \vdots & \ddots & \vdots \\ \text{vec}(\rho_{a_m}(g_1)) & \dots & \text{vec}(\rho_{a_m}(g_n)) \end{bmatrix}. \tag{4.21}$$

Notice that if A and K have nontrivial intersection, then some of the blocks of rows of $\mathbf{M}$ above may correspond to repeated or isomorphic representations. If this is the case our frame will no longer be tight. We can avoid this by assuming that $|K \cap A| = 1$, though we will typically not make use of this

assumption in our following proofs.

Now let us examine the inner products between our frame elements. From (4.12), the inner product corresponding to the group element g is

$$d \sum_{a \in A} \chi_a(g) = d \sum_{a \in A} \chi(a(g)). \tag{4.22}$$

Our aim is to generalize the concept from Theorem 6 of having one inner product per coset of a subgroup of $Aut(G)$. We first establish the following preliminary lemma:

Lemma 10. *Let A and K be subgroups of a finite group H such that the set product KA is a group, and let $\{a_i\}_{i=1}^{|A|/|A \cap K|}$ be a set of right coset representatives for $(A \cap K) \backslash A$. Then for each fixed a_i and $k \in K$, there is a unique $a_{i'}$ and $k' \in K$ such that $a_{i'}k = k'a_i$, and a unique $a_{i''}$ and $k'' \in K$ such that $a_i k = k''a_{i''}$.*

Proof. Since KA is a group (by assumption) which obviously contains both K and A, we can write $a_i k = \tilde{k}\tilde{a}$ for some $\tilde{k} \in K$ and $\tilde{a} \in A$. Then $\tilde{a}$ can further be written uniquely in the form $\tilde{k}_2 a_{i''}$ for some $\tilde{k}_2 \in A \cap K$ and $a_{i''}$ one of the right coset representatives of $A \cap K$ in A. Setting $k'' = \tilde{k}\tilde{k}_2$ gives us the second part of this theorem.

Now suppose there are two pairs (a_j, k_j') and (a_t, k_t') such that

$$a_j k = k_j' a_i, \tag{4.23}$$

$$a_t k = k_t' a_i. \tag{4.24}$$

Then from (4.24) we have $a_t(a_j)^{-1}a_j k = k_t'(k_j')^{-1}k_j' a_i$, and we can use (4.23) to cancel out $a_j k$ and $k_j' a_i$ from this expression to arrive at

$$a_t(a_j)^{-1} = k_t'(k_j')^{-1} \in A \cap K. \tag{4.25}$$

But since a_t and a_j are representatives of distinct right cosets of $A \cap K$ in A, they must be equal, hence $a_t = a_j$ and $k_t' = k_j'$. This shows that there can only be *at most* one pair $(a_{i'}, k')$ such that $a_{i'}k = k'a_i$. But since we have already shown that every $a_j k$ can be written uniquely in the form $k''a_{j''}$ for some $a_{j''}$, then since our groups are finite there must be some j for which $a_{j''} = a_i$, so there is *exactly* one such pair $(a_{i'}, k') = (a_j, k'')$ which satisfies the hypotheses of the lemma. $\qquad \square$

The next lemma now extends the coset idea of Theorem 6 to drastically reduce the number of distinct inner product values we need consider.

Lemma 11. *Let G be a finite group, $H \leq Aut(G)$, ρ an irreducible representation of G with character*

χ, and K the subgroup of H which fixes χ as in (4.20). Let A be a subgroup of H such that KA is a group. Then for any $\sigma_1, \sigma_2 \in H$ which are in the same right coset of KA, the inner products associated to $\sigma_1(g)$ and $\sigma_2(g)$ respectively are equal for any $g \in G$. That is,

$$d \sum_{a \in A} \chi_a(\sigma_1(g)) = d \sum_{a \in A} \chi_a(\sigma_2(g)). \tag{4.26}$$

Proof. Since σ_1 and σ_2 are in the same right coset of KA (which is equal to AK), there is some $h \in H$ such that $\sigma_1 = a_1 k_1 h$ and $\sigma_2 = a_2 k_2 h$ for some $a_1, a_2 \in A$ and some $k_1, k_2 \in K$. Thus, (4.22) becomes

$$d \sum_{a \in A} \chi(a\sigma_1(g)) = d \sum_{a \in A} \chi(aa_1 k_1 h(g)) \tag{4.27}$$

$$= d \sum_{a \in A} \chi(ak_1 h(g)), \tag{4.28}$$

which follows from the fact that multiplication by a_1 permutes the elements of A.

Now let $\{a_i\}$ be a set of right coset representatives for $(A \cap K) \backslash A$. Our sum now becomes

$$d \sum_{a \in A} \chi(ak_1 h(g)) = d \sum_{a_i} \sum_{\gamma \in A \cap K} \chi(\gamma a_i k_1 h(g)) \tag{4.29}$$

$$= d \sum_{a_i} |A \cap K| \chi(a_i k_1 h(g)), \tag{4.30}$$

which follows from fact that elements of K fix χ. Now for each a_i, we know from Lemma 10 that $a_i k_1$ is uniquely expressible in the form $k_j' a_j$ for some right coset representative a_j and some $k_j' \in K$. Thus, since the $\{a_i\}$ and $\{a_j\}$ are in one to one correspondence by Lemma 10, we can further rewrite our sum as

$$d \sum_{a_i} |A \cap K| \chi(a_i k_1 h(g)) = d \sum_{a_j} |A \cap K| \chi(k_j' a_j h(g)) \tag{4.31}$$

$$= d \sum_{a_j} |A \cap K| \chi(a_j h(g)) \tag{4.32}$$

$$= d \sum_{a_j} \sum_{\gamma \in A \cap K} \chi(\gamma a_j h(g)) \tag{4.33}$$

$$= d \sum_{a \in A} \chi(ah(g)). \tag{4.34}$$

Since the inner product depends only on h, we are done. $\square$

We can now express each inner product in terms of a right coset of KA and an *orbit* of G under

the automorphism group H. Two elements $g, g' \in G$ are said to be in the same orbit if there is an automorphism $h \in H$ such that $h(g) = g'$. Note that since $g = h^{-1}(g')$, this is an equivalence relation, so the orbits partition G. We may write this orbit as $Hg := \{h(g) \mid h \in H\}$, and we say that g is a *representative* of this orbit. It should be clear that the identity element $1 \in G$ is in its own orbit.

We are now equipped to bound both the number of distinct inner product values, as well as the coherence of our new frames. The following theorem contains the analogs of Lemma 6 and Theorem 6 to the broader class of frames we have just constructed.

Theorem 17. *Let G be a finite group of size n and ρ a degree-d irreducible representation of G with character χ. Define*

- *$H \leq Aut(G)$ a group of automorphisms of G,*

- *$K := \{\sigma \in H \ : \ \chi(\sigma(g)) = \chi(g), \ \forall g \in G\}$, the subgroup of H consisting of automorphisms which fix χ,*

- *$A = \{a_i\}_{i=1}^{m} \leq H$, any subgroup of H such that the set product KA is also subgroup of H with $A \cap K = 1$,*

- *$\{h_i\}_{i=1}^{n_c}$ representatives of the distinct cosets of KA in H*

- *$\{g_j\}_{j=1}^{n_o}$ representatives of the distinct orbits of G under H*

Finally, let $\mathbf{M}$ be the frame with elements $\{\sqrt{d}[\text{vec}(\rho_{a_1}(g))^T, ..., \text{vec}(\rho_{a_m}(g))^T]^T\}_{g \in G}$ as in (4.21). Then $\mathbf{M}$ is a tight frame with at most $n_c(n_o - 1)$ distinct inner product values between its vectors. If μ_W is the lower bound on coherence given by the Welch bound (explicitly $\mu_W = \sqrt{\frac{n-dm}{dm(n-1)}}$), then the coherence μ of our frame is bounded by

$$\mu \leq \sqrt{\frac{|G| - 1}{\min_{\{(i,j):g_j \neq 1\}} |KAh_i(g_j)|}} \mu_W. \tag{4.35}$$

Proof. By hypothesis, G is partitioned into distinct orbits $Hg_1, ..., Hg_{n_o}$ with representatives $g_1, ..., g_{n_o}$. Let $g \in G$ be in the j^{th} orbit so that for some $h \in H$ we have $h(g_j) = g$. Suppose that $h \in KAh_i$. Then from Lemma 11, the inner product associated to g is

$$d \sum_{a \in A} \chi(a(g)) = d \sum_{a \in A} \chi(ah(g_j)) = d \sum_{a \in A} \chi(ah_i(g_j)). \tag{4.36}$$

Thus, excluding the orbit corresponding to the identity element (which corresponds to taking the inner product of a column of $\mathbf{M}$ with itself), the number of nontrivial inner products that we must

consider is $n_c(n_o - 1)$, and the number of times the inner product corresponding to the pair (h_i, g_j) arises is

$$|KAh_i(g_j)| = \#\{kah_i(g_j) \ : \ k \in K, \ a \in A\}. \tag{4.37}$$

Now since our frame $\mathbf{M}$ is tight by Theorem 16, then from Lemma 6, the mean squared inner product between the frame vectors is equal to μ_W^2, and this mean can be written as

$$\mu_W^2 = \frac{1}{\sum_{h_i} \sum_{g_j \neq 1} |KAh_i(g_j)|} \cdot \sum_{h_i} \sum_{g_j \neq 1} |KAh_i(g_j)||\alpha_{i,j}|^2 \tag{4.38}$$

$$= \frac{1}{|G| - 1} \cdot \sum_{h_i} \sum_{g_j \neq 1} |KAh_i(g_j)||\alpha_{i,j}|^2, \tag{4.39}$$

where $\alpha_{i,j}$ is the inner product associated to the pair (h_i, g_j). From this, it follows that

$$(|G| - 1)\mu_W^2 \geq \left(\min_{\{(i,j):g_j \neq 1\}} |KAh_i(g_j)| \right) \left(\max_{\{(i,j):g_j \neq 1\}} |\alpha_{i,j}|^2 \right),$$

from which our result follows.

$\square$

We can see from Theorem 17 that in general our coherence will be closer to the Welch bound if we have fewer orbits, and the sets $KAh_i(g_j)$ are close to each other in size. We articulate this in the following corollary.

Corollary 3. *In Theorem 17, if the sets $KAh_i(g_j)$ are the same size for all h_i and all nonidentity g_j, we achieve our optimal upper bound in (4.35):*

$$\mu \leq \sqrt{n_c(n_o - 1)}\mu_W. \tag{4.40}$$

Proof. If there are n_c cosets of KA in H, and n_o orbits of G under the action of H, then since $\sum_{h_i} \sum_{g_j \neq 1} |KAh_i(g_j)| = |G| - 1$, we have

$$\min_{\{(i,j):g_j \neq 1\}} |KAh_i(g_j)| \leq \frac{|G| - 1}{n_c(n_o - 1)}, \tag{4.41}$$

with equality if and only if the sets $KAh_i(g_j)$ are all the same size. The result follows immediately. $\square$

For clarity, let us reiterate how our frames from Theorem 6 fall into the more general framework of Theorem 17. In this case,

- G is the cyclic additive group $\mathbb{Z}/n\mathbb{Z} = \{0, 1, ..., n-1\} \mod n$, where n is a prime.

- ρ is the representation $\rho(x) = e^{\frac{2\pi i x}{n}}$ for any $x \in G$.

- $\chi(x)$ is equal to $\rho(x)$ for any $x \in G$, since ρ is a degree-1 representation.

- H is the multiplicative group $(\mathbb{Z}/n\mathbb{Z})^{\times} = \{1, 2, ..., n-1\} \mod n$, where each element $h \in (\mathbb{Z}/n\mathbb{Z})^{\times}$ is viewed as an automorphism $h(x) = h \cdot x$.

- K is the subgroup of H such that $e^{\frac{2\pi i k x}{n}} = e^{\frac{2\pi i x}{n}}$, $\forall x \in G$. In this case, we can see that $K = \{1\}$.

- A is the size m subgroup of H, where $m | (n-1)$. Since K is trivial, KA is automatically a subgroup of H, and $A \cap K = 1$.

- n_c is the number of cosets of A in H, which is $\frac{n-1}{m}$. $\{h_i\}_{i=0}^{n_c}$ are the representatives of these cosets. If x is a cyclic generator for H, then the h_i can be taken to be the powers of x: $h_i = x^i$, $i = 1, ..., n_c$.

- $n_o = 2$, because there are only two orbits of G under H. One of these is the trivial orbit, $\{0\}$, and indeed $h \cdot 0 = 0$, $\forall h \in H$. All the nonzero elements $\{1, ..., n-1\} \subset G$ are in the same orbit, since any two of these elements differ only by a multiplicative factor in H. Thus we may take our two orbit generators to be $g_1 = 1$ (the generator of the nontrivial orbit) and $g_2 = 0$ (the generator of the tribal orbit).

In light of this last point, we see that these frames trivially satisfy the hypothesis of Corollary 3 since the sets $KAh_i(g_j)$ are simply the cosets Ah_i, which all have the same size as desired. (Note that since we write G additively in this situation, the identity element is 0 instead of 1, so the hypothesis of Corollary 3 effectively becomes that the sets $KAh_i(g_j)$ are the same size for $g_j \neq 0$). Thus the frames from Theorem 6 give us our optimal bound in Theorem 17, and the bound in (4.40) becomes $\mu \leq \sqrt{n_c}\mu_W$, which is the same bound we saw in Corollary 3. We will explore this connection more in the next section.

4.4   Subgroups and Quotients of General Linear Groups

We will now identify a class of groups that yield frames with remarkably low coherence using this framework, a subclass of which consists of the groups used in Theorem 6. Recall that in our original construction of Theorem 6, we chose G to be the additive group $\mathbb{Z}/n\mathbb{Z}$, where n was a prime p, and H was isomorphic to the multiplicative group $(\mathbb{Z}/n\mathbb{Z})^{\times}$, which contains all the nonzero elements of $\mathbb{Z}/n\mathbb{Z}$ when n is prime. This is equivalent to choosing G and H respectively to be the additive and

multiplicative groups of the finite field with p elements, $\mathbb{F}_p$. In this case, H is the simplest example of a general linear group. Indeed, H can be interpreted as the 1-dimensional invertible matrices with entries in $\mathbb{F}_p$. As we will now see, subgroups and quotients of matrix groups over finite fields lend themselves naturally to our construction.

4.4.1 Frames from Vector Spaces Over Finite Fields

Recall from our discussion following Theorem 17 that in general our coherence will be closer to the Welch bound if we have fewer orbits, and the sets $KAh_i(g_j)$ are close to each other in size. The optimal case is when their sizes are all equal, in which case we obtain the bound in Corollary 3. Equation (4.40) in this corollary closely resembles the result from Lemma 6. This is no coincidence, since the condition that the sets $KAh_i(g_j)$ have the same size is equivalent to requiring that each corresponding inner product value arises the same number of times as the inner product between two frame elements. (Recall that we exploited this latter property in deriving Lemma 6.) In a sense, the best case is when we have exactly one nontrivial orbit, so that $n_o = 2$. And if in addition the sets $KAh_i(g_j)$ have the same size for all h_i and $g_j \neq 1$, Corollary 3 shows that the coherence is bounded by a factor of $\sqrt{n_c}$ of the Welch bound.

We saw at the end of Section 4.3 that this happens in our original frames constructed in Theorem 6, when G was the additive group of a prime-sized finite field $\mathbb{F}_p \cong \mathbb{Z}/p\mathbb{Z}$ and H the set of automorphisms given by multiplication by elements of $\mathbb{F}_p^\times \cong (\mathbb{Z}/p\mathbb{Z})^\times$. As we remarked at the beginning of this section, H is the simplest example of a general linear group $GL(r, \mathbb{F}_p)$—the multiplicative group of $r \times r$ invertible matrices with entries in $\mathbb{F}_n$ (in this case $r = 1$). It turns out that even higher-dimensional general linear groups fit the framework of Corollary 3. If we set $H := GL(r, \mathbb{F}_p)$ then it is the automorphism group of $G := (\mathbb{F}_p)^r$, the r-dimensional vector space over $\mathbb{F}_p$ (viewed only as an additive abelian group). For any two nonzero vectors $\mathbf{v}_1$ and $\mathbf{v}_2$ in $(\mathbb{F}_p)^r$, there is an invertible matrix $\mathbf{W} \in GL(r, \mathbb{F}_p)$ such that $\mathbf{W}\mathbf{v}_1 = \mathbf{v}_2$, so all nontrivial elements of $(\mathbb{F}_p)^r$ lie in the same orbit under H.

Alternatively, we may view $(\mathbb{F}_p)^r$ as the additive group of the finite field with p^r elements, $\mathbb{F}_{p^r}$, which is a vector space over its subfield $\mathbb{F}_p$. An irreducible representation ρ of $\mathbb{F}_{p^r}$ (and hence of $(\mathbb{F}_p)^r$) is the function

$$\rho(x) = e^{\frac{2\pi i \operatorname{Tr}(x)}{p}}, \tag{4.42}$$

where $\text{Tr}(x)$ is the *trace* of the field element x, defined as

$$\text{Tr} : \mathbb{F}_{p^r} \to \mathbb{F}_p,$$

$$\text{Tr}(x) = x + x^p + x^{p^2} + \ldots + x^{p^{r-1}}. \tag{4.43}$$

The trace function in our context is the sum of the automorphisms of $\mathbb{F}_{p^r}$ fixing the subfield $\mathbb{F}_p$, and is so named because $\text{Tr}(x)$ is the trace of the matrix associated with the linear transformation of multiplication by x. This transformation acts on the additive group of $\mathbb{F}_{p^r}$ viewed as a vector space over $\mathbb{F}_p$. As such, the trace is an additive function: $\text{Tr}(x + y) = \text{Tr}(x) + \text{Tr}(y)$, and consequently $\text{Tr}(-x) = -\text{Tr}(x)$. In the case where $r = 1$, the trace becomes the identity function, and we see that as expected we recover a familiar representation of $\mathbb{F}_p$ similar to the ones used in Theorem 6.

We should point out that the *general* form of an irreducible representation of the additive group of $\mathbb{F}_{p^r}$ is $\rho_a(x) := \omega^{\text{Tr}(ax)}$, where $\omega = e^{\frac{2\pi i}{p}}$ and $a \in \mathbb{F}_{p^r}$. This is the image of the function ρ in (4.42) under the action of k viewed as a matrix in $GL(r, \mathbb{F}_p)$ as just described. As such, it is fitting that the notation "ρ_k" bears resemblance to that of equation (4.15). Note also that since each of these is a degree-1 representation, each is equal to its own character: $\chi_k(x) = \rho_k(x)$. Each of the p^r representations ρ_k, $k \in \mathbb{F}_{p^r}$, is unique, and from equation (4.1) we see that they indeed comprise *all* of the inequivalent irreducible representations of G.

Now, we can concisely describe the group K of character-preserving automorphisms from Theorem 17 as follows: K is simply the set of automorphisms in H which preserve the field trace, $K = \{k \in H \mid \text{Tr}(kx) = \text{Tr}(x), \ \forall x \in G\}$. It can be easily shown that the size of K is $|K| = |H|/|\mathbb{F}_{p^r}^\times| = (p^r - p) \ldots (p^r - p^{r-1})$. What is not clear, however, is the form that each element of K will take as a matrix in $H = GL(r, \mathbb{F}_p)$. The same issue arises when we attempt to compute the group A from Theorem 17.

To rectify this issue, we will shift our focus to the interpretation of G as the additive group of the field $\mathbb{F}_{p^r}$. And instead of choosing H to be the entire automorphism group $GL(r, \mathbb{F}_p)$, we will let H be the size-$(p^r - 1)$ subgroup of matrices corresponding to the nonzero field elements $\mathbb{F}_{p^r}^\times$. (Recall, each element of $\mathbb{F}_{p^r}^\times$ acts linearly on $\mathbb{F}_{p^r}$ by multiplication, and as such has a matrix representation when viewed as a linear transformation of $(\mathbb{F}_p)^r$.) In this new setting, the only element of H which fixes the field trace is 1, so K is now the trivial group.

It is reasonable to ask if we lose anything by choosing H to be only a proper subgroup of $GL(r, \mathbb{F}_p)$. But in fact, we can see from Lemma 11 and Theorem 17 that the coherence of our frames depends only on the right cosets of K in H. The following lemma shows that we do not lose anything by choosing H to be $\mathbb{F}_{p^r}^\times$ instead of $GL(r, \mathbb{F}_p)$:

Lemma 12. *Let $G = (\mathbb{F}_p)^r$ (which is the additive group of $\mathbb{F}_{p^r}$), and χ a character of G. Let $H_1 = GL(r, \mathbb{F}_p)$ with $K_1 \leq H_1$ the subgroup that fixes χ, and $H_2 = \mathbb{F}_{p^r}^\times \leq H_1$ with corresponding subgroup $K_2 = H_2 \cap K_1$. For every subgroup A_1 of H_1 with $A_1 \cap K_1 = 1$, there is a subgroup A_2 of H_2 with $A_2 \cap K_2 = 1$ such that the groups A_1 and A_2 give rise to the same inner products described by Lemma 11.*

Proof. As we touched on above, since our character is a function of the form $\chi(x) = e^{\frac{2\pi i \operatorname{Tr}(ax)}{p}}$, we observe that no nontrivial element of H_2 fixes χ. Thus, $K_2 = 1$. Since the right cosets $K_1 \, H_1 = \{K_1 h_1 \; : \; h_1 \in H_1\}$ partition H_1, each element $h_2 \in H_2$ must lie in some such coset. We claim that no two elements of H_2 are in the same right coset of K_1. To see this, assume we have h_2 and h_2' in H_2 which lie in the same right coset of K_1. This means that $h_2' h_2^{-1} \in K_1 \cap H_2 = K_2$, hence h_2 and h_2' must be equal. Furthermore, we know that there is one element of H_2 in *each* right coset of K_1 in H_1, since *every* character of G can be written in the form $\chi(h_2(x))$ for some multiplicative field element $h_2 \in H_2$. (This is a well-known fact that can be found, for example, in [83].)

Now, if A_1 is a subgroup of H_1 which intersects K_1 trivially, each element of A_1 must lie in a distinct right coset of K_1. For each element $a_1 \in A_1$, let a_2 be the unique element of H_2 which lies in the same such coset, and let A_2 be the set of all these elements. Clearly A_2 has trivial intersection with K_2, since it is a subset of H_2. The fact that A_2 is itself a group is easy to verify. For example, for elements a_2 and a_2' in A_2, with corresponding elements a_1 and a_1' in A_1, we see that the product $a_2' a_2^{-1}$ is also an element of A_2 since it is the field element lying in the same right coset of K_1 as $a_1' a_1^{-1} \in A_1$. Since elements a in the same right coset of K_1 give rise to the same character $\chi(a(x))$, we see also that the groups A_1 and A_2 will give rise to the same frame inner products as described in Lemma 11. $\qquad \square$

Let us explicitly match this example with the framework of Theorem 17. We note that

- G is the additive group of the vector space $(\mathbb{F}_p)^r$, or equivalently the additive group of the field $\mathbb{F}_{p^r}$.

- $\rho(x) = e^{\frac{2\pi i \operatorname{Tr}(x)}{p}}$.

- $\chi(x) = \rho(x)$, since ρ is a 1-dimensional representation, hence is equal to its own character.

- $H = \mathbb{F}_{p^r}^\times = \mathbb{F}_{p^r} \setminus \{0\}$, where G is viewed as the additive group of $\mathbb{F}_{p^r}$. Basic field theory tells us that H is isomorphic to the cyclic group of size $p^r - 1$.

- $K = 1$, since the only field element $h \in H$ such that $\chi(h(x)) = \chi(x)$ is the identity.

- $A = \{a_1, ..., a_m\}$ is any subgroup of H, which will necessarily be a cyclic group of size m, where m is a divisor of $p^r - 1$. Since H is cyclic, there is a unique subgroup for each such m, and

it consists of the $\left(\frac{p^r-1}{m}\right)^{th}$ powers in H. Thus, if x is a cyclic generator for H, we may set $y = x^{\frac{p^r-1}{m}}$ and $a_i = y^i$ for each $i = 1, ..., m$.

- $n_c = \frac{p^r-1}{m}$, the number of cosets of A in H. If x is a generator for the cyclic group H, these cosets are $h_i = x^i$, $i = 1, ..., n_c$.

- $n_o = 2$, since again $0 \in \mathbb{F}_{p^r}$ is in its own orbit, and all the nontrivial elements are in their own orbit under H (generated by $1 \in \mathbb{F}_{p^r}$).

Our new frame matrix $\mathbf{M}$ from (4.21) becomes

$$\mathbf{M} = \begin{bmatrix} \omega^{\text{Tr}(a_1 x_1)} & \omega^{\text{Tr}(a_1 x_2)} & \cdots & \omega^{\text{Tr}(a_1 x_n)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \omega^{\text{Tr}(a_m x_1)} & \omega^{\text{Tr}(a_m x_2)} & \cdots & \omega^{\text{Tr}(a_m x_n)} \end{bmatrix}, \tag{4.44}$$

where we have expressed the elements of our field as $\{x_i\}_{i=1}^n$. If $x_j - x_i = x_\ell$, the inner product between the i^{th} and j^{th} columns now becomes

$$\sum_{a_t} \left(\omega^{\text{Tr}(a_t x_i)}\right)^* \left(\omega^{\text{Tr}(a_t x_j)}\right) = \sum_{a_t} \omega^{\text{Tr}(a_t(x_j - x_i))} \tag{4.45}$$

$$= \sum_{a_t} \omega^{\text{Tr}(a_t x_\ell)}. \tag{4.46}$$

We can see from (4.46) that as in our original frames from Theorem 6, we have exactly $\frac{n-1}{m}$ nontrivial inner product values: one for each element of $\mathbb{F}_{p^r}^\times$ (each of which represents a right coset of K in H). Again, each of these values arises as an inner product the same number of times.

Since these new frames are a generalization our original frames constructed in Theorem 6, it should come as no surprise that the bounds in Theorems 9 and 10 generalizes as well:

Theorem 18. *If n is prime power p^r, m a divisor of $n-1$, and $\{a_i\}$ the unique subgroup of $\mathbb{F}_{p^r}^\times$ of size m, then setting $\omega = e^{\frac{2\pi i}{p}}$, and $\kappa := \frac{n-1}{m}$, the coherence μ of our frame $\mathbf{M}$ in (4.44) satisfies*

$$\mu \le \frac{1}{\kappa} \left((\kappa - 1)\sqrt{\frac{1}{m}\left(\kappa + \frac{1}{m}\right)} + \frac{1}{m} \right). \tag{4.47}$$

If both p and m are odd, μ satisfies the tighter bound

$$\mu \le \frac{1}{\kappa}\sqrt{\left(\frac{1}{m} + \left(\frac{\kappa}{2} - 1\right)\beta\right)^2 + \left(\frac{\kappa}{2}\right)^2 \beta^2}, \tag{4.48}$$

where $\beta = \sqrt{\frac{1}{m}\left(\kappa + \frac{1}{m}\right)}$.

Proof. We present this proof in Appendix B.1. □

4.4.2 Smaller Alphabets and Frames from Hadamard Matrices

We emphasize that these generalized frames have several advantages over the original frames constructed in Theorem 6. First, the number n of frame vectors is no longer limited to being a prime, but is instead a power of a prime, $n = p^r$. Furthermore, the entries of our frame matrix $\mathbf{M}$ in (4.44) are no longer n^{th} roots of unity, but rather p^{th} roots of unity. This allows for more practical implementations of our frames. Indeed, while our original frames did achieve low coherence, the entries of the frame vectors came from an alphabet size as large as the frame itself. Thus even for small examples our frames could require an alphabet size of at least several hundred. In our new frames, we could fix p to be a small prime and take a number of frame elements that is substantially larger, yet our frame vectors will only have entries from an alphabet of size p.

For instance, if $p = 2$, then even though our frame can have $n = 2^r$ elements for any r, the matrix $\mathbf{M}$ will always have ± 1 entries. In this case, we have the following:

Theorem 19. *When $p = 2$ in our above framework, our frame matrix $\mathbf{M}$ in (4.44) is a subset of rows of an $n \times n$ Hadamard matrix.*

Proof. We already commented above that when $p = 2$, $\mathbf{M}$ will have ± 1 entries. The theorem then follows from the fact that the frame is tight (i.e. the rows of $\mathbf{M}$ are orthogonal with equal norm) by Theorem 16. □

This is not the first time that frames with ± 1 entries have been explored. For example, [75] designed such frames using codes constructed by [6] and [113], and analyzed the frames' geometry. Figure 4.1 illustrates the benefit of using our frames to control coherence. Depicting histograms of the inner products resulting from selecting two sets of 341 rows of from a 1024×1024 Hadamard matrix using our method (red) versus randomly (blue), we can see that our construction actually yields just two distinct inner product values in this case, both much closer to zero than the largest magnitude inner products from the random case. In Table 4.1, we compute the coherences of several random vs. group Hadamard frames, and compare to the Welch bound for reference. The group Hadamard frames have consistently lower coherence than the random Hadamard frames, particularly when the frame dimensions m and n are large but the quotient $\kappa = \frac{n-1}{m}$ is small.

4.4.3 Difference Sets

On one final note, we point out that in certain cases the group A forms a *difference set* in $\mathbb{F}_{p^r}$, that is, each nonzero element of $\mathbb{F}_{p^r}$ occurs as a difference $a_i - a_j$ the same number of times. In this case,
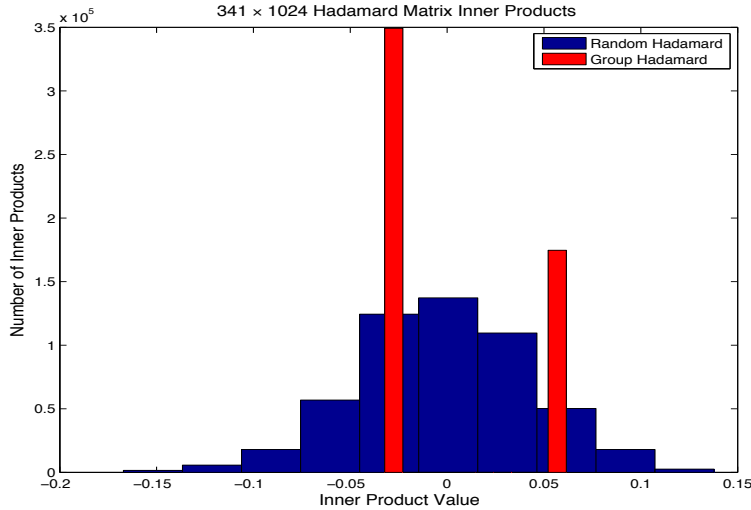
Figure 4.1: Superimposed histograms of the inner product values between elements of a 341×1024 frame formed from our method of selecting rows of the Hadamard matrix (red) versus selecting the rows randomly (blue). The red values are concentrated to only two points in this case, resulting in coherence closer to zero.

Table 4.1: Coherences for Random vs. Group Hadamards

(n, m)	Random Hadamard	Group Hadamard	$\sqrt{\frac{n-m}{m(n-1)}}$
$(256, 51)$	.3725	.2549	.1256
$(256, 85)$	.2941	.1294	.0888
$(512, 73)$	.3425	.2329	.1085
$(1024, 341)$	.2023	.0616	.0442
$(4096, 455)$	.1868	.1253	.0442

our frames yield examples of those constructed [110] and [34]:

Theorem 20. *The columns of* $\mathbf{M}$ *in (4.44) form a tight equiangular frame if and only if the elements in* $A = \{a_i\}_{i=1}^m$ *form a difference set in* $\mathbb{F}_{p^r}$. *In this case, the coherence of* $\mathbf{M}$ *achieves the Welch bound. In particular, our construction yields a difference set when* $\frac{p^r-1}{m} = 2$ *and* m *is odd.*

Proof. Again, this follows from the arguments in [110] and [34] (see Theorem 3 of the latter). When $\frac{p^r-1}{m} = 2$, A is the group of squared elements in $\mathbb{F}_{p^r}^{\times}$, which is a well-known difference set when $p^r \equiv 3$ mod 4 (an example of what is called a "Paley difference set"). [107] This is precisely the case when m is odd. $\qquad \square$

Unfortunately, the Hadamard frames we constructed in the previous section cannot satisfy the condition $\frac{p^r-1}{m} = 2$, since they require that $p = 2$. We can, however, use our construction to produce

tight, equiangular frames whose entries are from an alphabet only of size three—the third roots of unity:

Corollary 4. *Let $p \equiv 3 \mod 4$ be a prime, r an odd integer, and set $m := \frac{p^r - 1}{2}$. Choose the set $A = \{a_i\}_{i=1}^m$ to be the unique subgroup of $\mathbb{F}_{p^r}$ of size m. Then the columns of $\mathbf{M}$ in (4.44) form a tight equiangular frame whose entries are each one of the distinct p^{th} roots of unity. In particular, when $p = 3$, the entries of $\mathbf{M}$ come from an alphabet of size three.*

Proof. Since r is odd, we have $p^r \equiv 3 \mod 4$, so the set A forms a Paley difference set as mentioned in the proof of Theorem 20. Thus the columns of M form a tight, equiangular frame whose elements are integer powers of $\omega = e^{2\pi i/p}$, i.e., the p^{th} roots of unity. $\square$

In Table 4.2, we list the coherences of several of the tight, equiangular frames arising from Corollary 4, and compare the coherence to when the matrix $\mathbf{M}$ in (4.44) is formed by randomly choosing the elements $\{a_i\}_{i=1}^m$. As expected, our frames consistently have lower coherence, in this case meeting the Welch bound.

Table 4.2: Coherences for Random vs. Group Matrices with Small Alphabets, $m = \frac{n-1}{2}$

n	Random	Group	$\sqrt{\frac{n-m}{m(n-1)}}$
3^3	.3353	.2035	.2035
3^5	.1577	.0645	.0645
3^7	.0509	.0214	.0214
7^3	.1110	.0542	.0542
11^3	.0674	.0274	.0274

Coherences of $m \times n$ frame matrices formed from rows of the group Fourier matrices for the finite fields $\mathbb{F}_q$, $q = n$. We compare choosing the rows randomly with using the group method from Section 4.4.1, which produces tight, equiangular frames by Corollary 4. When $n = p^r$, the matrix entries are p^{th} roots of unity.

4.5   Frames from Special Linear Groups

To show how our framework can be applied to more complicated groups, we will demonstrate how to obtain frames with low coherence in the case where G is the special linear group $SL_2(\mathbb{F}_q)$. Frames of this type were discussed in [100]. This matrix group is easy to describe, but it is nonabelian and has irreducible representations of degree greater than 1, hence will be interesting for our purposes.

Let $\mathbb{F}_q$ be the finite field containing q elements, where q is some integral power of a prime number. Then $SL_2(\mathbb{F}_q)$ is the set of 2×2 determinant-1 matrices with entries in $\mathbb{F}_q$,

$$SL_2(\mathbb{F}_q) := \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \mid a, b, c, d \in \mathbb{F}_q, \ ad - bc = 1 \right\}.$$

Table 4.3: Character Table of $SL_2(\mathbb{F}_q)$, q even

Class Representative:	$\begin{bmatrix}1&0\\0&1\end{bmatrix}$	$\begin{bmatrix}1&1\\0&1\end{bmatrix}$	$\begin{bmatrix}c&0\\0&c^{-1}\end{bmatrix}$	$B\begin{bmatrix}s&0\\0&s^{-1}\end{bmatrix}B^{-1}$
No. of such classes:	1	1	$\frac{1}{2}(q-2)$	$\frac{1}{2}q$
Size of class:	1	q^2-1	$q(q+1)$	$q(q-1)$
1_G	1	1	1	1
St_G	q	0	1	-1
ρ_χ	$q+1$	1	$\chi(c)+\chi(c^{-1})$	0
π_η	$q-1$	-1	0	$-\eta(s)-\eta(s^{-1})$

Here, $c \in \mathbb{F}_q$ and $s \in \mathbb{F}_{q^2}$, where s is an element of norm 1. B is an invertible matrix with entries in $\mathbb{F}_{q^2}$.

It is not difficult to check that the size of this group is $|SL_2(\mathbb{F}_q)| = q(q+1)(q-1)$.

Table 4.3 is the character table of $SL_2(\mathbb{F}_q)$ for when q is even (a power of 2). As we can see, in this case the matrices fall into four types of conjugacy classes based on how they diagonalize. The first is simply the identity matrix, $\begin{bmatrix}1&0\\0&1\end{bmatrix}$. The second consists of the matrices that are not diagonalizable, and have the Jordan canonical form $\begin{bmatrix}1&1\\0&1\end{bmatrix}$. These first two conjugacy classes contain all the matrices in $SL_2(\mathbb{F}_q)$ with repeated eigenvalues of 1.

Each conjugacy class of the third type has a representative which is a diagonal matrix: $\begin{bmatrix}c&0\\0&c^{-1}\end{bmatrix}$, where $c \in \mathbb{F}_q \setminus \{0,1\}$. Since the diagonal matrices $\mathrm{diag}(c, c^{-1})$ and $\mathrm{diag}(c^{-1}, c)$ are conjugate to each other, there are $\frac{1}{2}(q-2)$ such classes.

The fourth type of conjugacy class consists of matrices whose eigenvalues do not lie in $\mathbb{F}_q$. These are the matrices that take the form $B\begin{bmatrix}s&0\\0&s^{-1}\end{bmatrix}B^{-1}$, where $B \in SL_2(\mathbb{F}_{q^2})$ and $s \in \mathbb{F}_{q^2}$ is one of the norm-1 elements of $\mathbb{F}_{q^2} \setminus \mathbb{F}_q$, that is, $s^{q+1} = 1$. Note that here $\mathbb{F}_{q^2}$ is the finite field of q^2 elements, which contains $\mathbb{F}_q$ as a subfield. There are $q+1$ elements of $\mathbb{F}_{q^2}$ which satisfy the equation $s^{q+1} = 1$. Of these, the only element lying in $\mathbb{F}_q$ is 1, and the remaining q lie in $\mathbb{F}_{q^2} \setminus \mathbb{F}_q$. As in the previous case, these q elements pair up to represent a total of $q/2$ distinct conjugacy classes of the fourth type.

There are four types of characters of $SL_2(\mathbb{F}_q)$ for q even, arising as a consequence of the four types of conjugacy classes. The interested reader can refer to [79, 83, 93] to learn in depth how these characters come about, but for now we will give brief descriptions. The first two characters both correspond to degree-1 representations. They include the character of the identity representation 1_G, which maps every element to 1, and that of the Steinberg representation St_G, which maps elements of the various conjugacy classes to the values shown in Table 4.3. For our purposes, the third and fourth types of characters in the last two rows of the table are of greater interest. The third corresponds to

what is called an *induced representation*, denoted here as ρ_χ. It is a degree-$(q+1)$ representation built from an underlying nontrivial degree-1 representation χ of the multiplicative group $\mathbb{F}_q^\times$, a cyclic group of size $q - 1$. If $\tilde{c}$ is a cyclic generator for $\mathbb{F}_q^\times$ (so that every element can be written as a power of $\tilde{c}$), and we set $\omega_- = e^{\frac{2\pi i}{q-1}}$, then χ is a function of the form $\chi(\tilde{c}^\ell) = \omega_-^{a\ell}$, for some fixed $a \in \{1, 2, ..., q-2\}$. (It is required that a be nonzero modulo $q - 1$ in order for ρ_χ to be irreducible.)

The final type of character, denoted π_η, corresponds to a degree-$(q-1)$ *cuspidal representation*. A cuspidal representation is constructed from a degree-1 representation η of the set of norm-1 elements of $\mathbb{F}_{q^2}$, which is a cyclic multiplicative group of size $q + 1$. Given a cyclic generator $\tilde{s}$ for this group, and setting $\omega_+ = e^{\frac{2\pi i}{q+1}}$, then η will take the form $\eta(\tilde{s}^\ell) = \omega_+^{h\ell}$, where h is some fixed integer in the set $\{1, 2, ..., q\}$. (Again we require $h \not\equiv 0 \mod q + 1$ for irreducibility of π_η.

4.5.1 Frames from Induced and Cuspidal Representations

We can now use our previous results to design low-coherence frames in the form of $\mathcal{F}$ in (4.4) using the characters of $SL_2(\mathbb{F}_q)$ for q even. We emphasize that while explicitly writing out our frame vectors can be cumbersome and requires a certain amount of work in its own right, we will find that *identifying* which representations to use will be quick, as will computing the coherence of the resulting frame.

We will first focus our attention on only the induced representations. For convenience, we will write χ_a and ρ_a, respectively, for the representations χ and ρ_χ where $\chi(\tilde{c}) = \omega_-^a$. It remains to identify a suitable group A of automorphisms of $SL_2(\mathbb{F}_q)$ under which we can take the image of an induced representation to construct our frames, as prescribed by Theorem 17. In the last section, when our group was just the additive group of a finite field $\mathbb{F}_q$, our automorphisms corresponded to the nonzero field elements which formed the cyclic multiplicative group $\mathbb{F}_q^\times$. These automorphisms were well-described and easy to work with. It turns out that each automorphism φ of $\mathbb{F}_q$ induces an automorphism of $SL_2(\mathbb{F}_q)$ by simply applying φ to the entries of the 2×2 matrices in the special linear group:

$$\varphi\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) := \left(\begin{bmatrix} \varphi(a) & \varphi(b) \\ \varphi(c) & \varphi(d) \end{bmatrix}\right). \tag{4.49}$$

This observation enables us to continue working with the automorphisms of $\mathbb{F}_q$, so we can again choose A to be a subgroup of $\mathbb{F}_q^\times$. If $a' \in A \leq \mathbb{F}_q^\times$, then as an automorphism a' acts on ρ_a as

$$a' \cdot \rho_a = \rho_{a' \cdot a}. \tag{4.50}$$

Thus, it would be natural to choose for A to act on the representation ρ_1, so that the images under

A will be the representations $\{\rho_a \mid a \in A\}$. For the sake of simplicity, we will set $K = 1$ and $H = A$ in our Theorem 17 notation.

One caveat that we now face by choosing this set of automorphisms is the following: notice that each element of A fixes the element $u = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \in SL_2(\mathbb{F}_q)$, which means that there is a size-1 orbit $KA(u)$. This means that the bound we gave in Equation (4.35) of Theorem 17 will be somewhat ineffective. We can get around this problem by noticing that from Equations (4.38) and (4.39), the magnitude of the largest inner product will still be small as long as the inner product corresponding to u is small in magnitude. We quickly see this to be the case based on the equation for the inner product given in (4.22) and the fact that, from Table 4.3, the character values $\rho_a(u)$ are all equal to 1, a relatively small constant. We will give an explicit formula for the inner product corresponding to u in Equation (4.52), and after normalizing our frame elements (dividing the inner product by the squared norm of a frame element) this inner product becomes very small as q grows.

Since we are working with such a familiar set of automorphisms A, we would like to exploit some of the tools we developed for our frames constructed from finite fields. Consider choosing q such that $q - 1$ is some prime p. In this case, χ_a is simply a representation of the cyclic group $\mathbb{Z}/p\mathbb{Z}$, which is isomorphic to the additive group of the field $\mathbb{F}_p$. From the preceding sections, we already have powerful tools at our disposal for bounding certain sums of these characters. Since the character χ_a appears in the main part of the character ρ_a (as shown in Table 4.3), we would like to apply these tools to bound sums of the ρ_a as well. This will allow us to use our bounds from Theorem 18 to obtain even tighter bounds on coherence than those we could obtain from Theorem 17.

Intuitively, if we take m to be a divisor of $p - 1$, and let $A = \{a_1, ..., a_m\}$ be the unique size-m subgroup of $(\mathbb{Z}/p\mathbb{Z})^\times$ (explicitly the set $\{1, ..., p - 1\}$ with multiplication modulo p), then we should achieve frames with low coherence by using A to choose the representations ρ_{a_i} to use in our frame matrix $\mathcal{F}$ from (4.4). Note that from our previous notation,

Now, notice that based on Table 4.3, the characters corresponding to ρ_a and ρ_{-a} are the same (where $-a$ is taken modulo p). This indicates that ρ_a and ρ_{-a} are in fact equivalent representations. If -1 is contained in A and is not equivalent to 1 in $\mathbb{Z}/p\mathbb{Z}$ (which is always the case when q is even, since $p \neq 2$), then for each $a_i \in A$ we also have $-a_i \in A$, and $-a_i \not\equiv a_i$ in $\mathbb{Z}/p\mathbb{Z}$. In this case, the set of chosen representations $\{\rho_a \mid a \in A\}$ has repetition, and using these representations as the rows of $\mathcal{F}$ would yield repeated rows of the Group Fourier Matrix of $SL_2(\mathbb{F}_q)$, and hence would not produce a tight frame (based on Theorem 16). More importantly for our purposes, the resulting frame would not fit our criteria from Theorem 17, which means we could not use the tools we have built to bound its coherence. Therefore, if -1 lies in the unique subgroup of $(\mathbb{Z}/p\mathbb{Z})^\times$ of size m, we must choose A slightly differently. First, let us explicitly describe how the size-m subgroup decomposes into pairs

$\{a, -a\}$:

Lemma 13. *Let $q = 2^d$ for some positive integer d, such that $p = q - 1$ is a prime. Take a divisor m of $p - 1$, and let A_m be the unique size-m subgroup of $(\mathbb{Z}/p\mathbb{Z})^{\times}$. Then A_m contains -1 if and only if m is even. In this case, $m/2$ is odd, and $A_m = A_{m/2} \cup -A_{m/2}$ where $A_{m/2}$ is the unique size-$\frac{m}{2}$ subgroup and $-A_{m/2} = \{-a \mid a \in A_{m/2}\}$.*

Proof. Since p is necessarily odd, -1 generates the unique size-2 subgroup of $(\mathbb{Z}/p\mathbb{Z})^{\times}$, and A_m contains this subgroup if and only if its size m is even.

Since $p = q - 1$ is prime, then writing q in the form 2^d for some integer d, we must have $d > 1$. In this case, m is a divisor of $q - 2 = 2(2^{d-1} - 1)$. In this form, it is clear that $q - 2$ can never be divisible by 4 (since the factor $(2^{d-1} - 1)$ is odd), so neither can its divisor m. Thus, if m is even, $m/2$ must be odd, so $-1 \notin A_{m/2}$. As a result, for any $a \in A_{m/2}$, we must have $-a \in A_m \setminus A_{m/2}$ (since $A_{m/2}$ is a subgroup of A_m). By comparing sizes, we see that A_m must be equal to the union $A_{m/2} \cup -A_{m/2}$. $\qquad \square$

From Lemma 13, we see that when m is an even divisor of $p - 1$, the obvious candidate for the group A is the unique size-$\frac{m}{2}$ subgroup of $(\mathbb{Z}/p\mathbb{Z})^{\times}$, which will ensure that -1 is not in A and that our resulting frame is tight. With this in mind, we will simply assume that we choose m to be odd. The following theorem uses our previous results on frames constructed from finite fields to give a bound on the coherence of the frames we can construct from the induced representations of $SL_2(\mathbb{F}_q)$, for q even.

Theorem 21. *Take q a power of 2 such that $q - 1$ is a prime p, and let m be an odd divisor of $p - 1$ and $\kappa = \frac{p-1}{2m}$. Let $A = \{a_1, ..., a_m\}$ be the unique subgroup of $(\mathbb{Z}/p\mathbb{Z})^{\times}$ of size m, and form $\mathcal{F}$ (as in (4.4)) from the induced representations ρ_{a_i}. Then the coherence $\mu_{\mathcal{F}}$ of $\mathcal{F}$ is bounded by*

$$\mu_{\mathcal{F}} \leq \frac{1}{q+1} \max \left(1, \frac{1}{\kappa} \left((\kappa - 1) \sqrt{\frac{1}{2m} \left(\kappa + \frac{1}{2m} \right)} + \frac{1}{2m} \right) \right). \tag{4.51}$$

Proof. From Equation (4.12) and Table 4.3, we see that the only nontrivial inner products between the columns of $\mathcal{F}$ are those corresponding to the conjugacy classes represented by $u := \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \in SL_2(\mathbb{F}_q)$

and $w_\ell := \begin{bmatrix} \tilde{c}^\ell & 0 \\ 0 & \tilde{c}^{-\ell} \end{bmatrix} \in SL_2(\mathbb{F}_2)$ for $\ell \in \{1, ..., q-2\}$. These inner products are

$$u: \quad \sum_{i=1}^{m} d_i \chi_{\rho_{a_i}}(u) = m(q+1) \tag{4.52}$$

$$\tag{4.53}$$

$$w_\ell: \quad \sum_{i=1}^{m} d_i \chi_{\rho_{a_i}}(w_\ell) = \sum_{i=1}^{m} (q+1) \cdot (\chi_{a_i}(\tilde{c}^\ell) + \chi_{a_i}(\tilde{c}^{-\ell})) \tag{4.54}$$

$$= (q+1) \sum_{i=1}^{m} (\omega_-^{\ell a_i} + \omega_-^{-\ell a_i}). \tag{4.55}$$

From Lemma 13 and the fact that m is odd by assumption, we can see that the union $A \cup -A = \{\pm a_1, ..., \pm a_m\}$ is actually the unique subgroup of $(\mathbb{Z}/p\mathbb{Z})^\times$ of size $2m$. If we denote this subgroup as A_{2m}, then we can write the sum in (4.55) in the form

$$\sum_{i=1}^{m} (\omega_-^{\ell a_i} + \omega_-^{-\ell a_i}) = \sum_{a \in A_{2m}} \omega_-^{\ell a}. \tag{4.56}$$

But this is just a scaled version of one of our original inner products between the elements of the harmonic frames that we constructed in Theorems 6 and 9, so we can use Theorem 9 to bound its magnitude.

To complete the proof, we simply need to take the maximum of the inner product magnitudes corresponding to the elements u and w_ℓ. This maximum becomes scaled after we normalize the columns of $\mathcal{F}$ by $\sqrt{m(q+1)^2}$, where we obtain the column norm from Equation (4.13) and the fact that the induced representations are $(q+1)$-dimensional. $\qquad\square$

Table 4.4: $SL_2(\mathbb{F}_q)$ vs. Gaussian Frame Coherences

Frame Dimensions	$SL_2(\mathbb{F}_q)$	Random Gaussian	Welch Bound
25×60	.2000	.5214	.1540
81×504	.2002	.3482	.1019
243×504	.1111	.2274	.0462

Theorem 21 gives us a recipe for constructing low-coherence frames from the induced representations of $SL_2(\mathbb{F}_q)$ for q even. These frames will consist of $q(q+1)(q-1)$ vectors (one for each element of $SL_2(\mathbb{F}_q)$) which are $m(q+1)^2$-dimensional. Figure 4.2 shows how our upper bound from the theorem comes decently close to the Welch lower bound on coherence. In table 4.4, we provide some explicit values of our frames' coherence, and for comparison we have included the coherence of frames of the same dimensions and number of elements whose coordinates are chosen independently from a Gaus-

Figure 4.2: Comparison of the Welch lower bound on coherence with the upper bound given by Theorem 21 for frames constructed from the induced representations of $SL_2(\mathbb{F}_q)$, for q a power of 2 such that $q - 1$ is prime. The number of frame vectors is $|SL_2(\mathbb{F}_q)| = q(q + 1)(q - 1)$, which are $m(q + 1)^2$-dimensional. Here, we have fixed $\kappa := \frac{q-2}{2m} = 3$.

sian distribution. While the frame matrix $\mathcal{F}$ can be concretely written out using the explicit forms of the representations given in [79, 83, 93], we will omit this process since we have already described it in depth and since these particular frames tend to have rather large dimensions. We remark that we can obtain similar results using a parallel construction of $\mathcal{F}$ with only cuspidal representations π_η, which works when $q + 1$ is prime.

4.6 Satisfying the Strong Coherence Property

A closely related quantity to the coherence of a frame $\{f_i\}_{i=1}^n$ in $\mathbb{C}^m$ is the *average coherence* ν, defined as

$$\nu = \frac{1}{n-1} \max_{i \in [n]} \left| \sum_{j \neq i} \langle f_i . f_j \rangle \right|. \tag{4.57}$$

When discussing the average coherence, the usual quantity μ is sometimes referred to as the *worst-case coherence*. [1] and [75] use the average coherence to describe the following properties of certain frames:

Definition 5. *A frame $\{f_i\}_{i=1}^n$ in $\mathbb{C}^m$ with average coherence ν and worst-case coherence μ is said to satisfy the Coherence Property if*

1. $\mu \leq \frac{0.1}{\sqrt{2 \log n}}$, *and*

2. $\nu \leq \frac{\mu}{\sqrt{m}}$.

It satisfies the Strong Coherence Property if

1. $\mu \leq \frac{1}{164 \log n}$, *and*

2. $\nu \leq \frac{\mu}{\sqrt{m}}$.

These works also give theoretical guarantees on the sparse-signal-recovery abilities of frames satisfying these properties. In particular, they discuss the One-Step Thresholding (OST) algorithm (Algorithm 1) described in [1]. If $F \in \mathbb{C}^{m \times n}$ has columns which form a unit-norm frame, $x \in \mathbb{C}^{n \times 1}$ is a sparse signal, and $e \in \mathbb{C}^{n \times 1}$ is a noise vector, OST produces an estimate $\hat{x}$ for x given $y := Fx + e$. We assume the entries of e are iid complex Gaussian values with mean 0 and variance σ^2, and the OST threshold λ is chosen to be

$$\lambda := \left(\sqrt{2\sigma^2 \log n} \right) \max\{\frac{10}{t}\mu_F \sqrt{m} \cdot \mathsf{SNR}, \ \frac{\sqrt{2}}{1-t}\}, \tag{4.58}$$

where μ_F is the worst-case coherence of F, SNR is the signal to noise ratio $\frac{||x||_2^2}{\mathbf{E}[||e||_2^2]}$, and t is a parameter chosen between 0 and 1. If F satisfies the coherence property, [1] finds regimes in which the support of $\hat{x}$ is equal to that of x with high probability. If F further satisfies the *strong* coherence property, [75] further provides high-probability bounds on the error $||x - \hat{x}||_2$. In the absence of an error vector e, [1] also finds cases where $\hat{x}$ is identically equal to x with high probability, though this calls for a different threshold, $\lambda = 10\mu_F||y||_2 \sqrt{\frac{2 \log n}{1-e^{-1/2}}}$. For our purposes, however, we will mainly focus on recovering signals with complex Gaussian error.

Algorithm 1 One-Step Thresholding (OST) Algorithm [1]

1: **Input:** $F \in \mathbb{C}^{m \times n}$ whose columns form a unit-norm frame, a vector $y = Fx + e$, and a threshold $\lambda > 0$.
2: **Output:** Estimates $\hat{K}$ for supp(x) and $\hat{x} \in \mathbb{C}^{n \times 1}$ for x.
3: $\hat{x} \leftarrow 0$
4: $z \leftarrow F^*y$
5: $\hat{\mathcal{K}} \leftarrow \{i \ : \ |z_i| > \lambda\}$
6: $\hat{x}_{\hat{\mathcal{K}}} \leftarrow (F_{\hat{\mathcal{K}}})^\dagger y$

It turns out that we can explicitly compute the average coherence of our frames from Theorem 17, and indeed any group frame constructed from a set of distinct irreducible representations of the same degree:

Theorem 22. *Let G be a finite group of size n and $\rho_1, ..., \rho_m$ a set of distinct nontrivial degree-d irreducible representations of G. Then the columns of the matrix $\mathbf{M} = \sqrt{d}[\text{vec}\rho_i(g_j)] \in \mathbb{C}^{md \times n}$ from*

(4.9) form a frame with average coherence $\nu = \frac{1}{n-1}$. If μ is the worst-case coherence of $\mathbf{M}$, then $\nu \leq \frac{\mu}{\sqrt{md}}$ provided that $n \geq 2md$.

Proof. From equations (4.12) and (4.13), we have that after normalizing the columns of $\mathbf{M}$, the inner product between the i^{th} and j^{th} columns is

$$\frac{\sum_{t=1}^{m} d \cdot \text{vec}(\rho_t(g_i))^* \text{vec}(\rho_t(g_j))}{md^2} = \frac{1}{md} \sum_{t=1}^{m} \chi_t(g_i^{-1} g_j). \tag{4.59}$$

Then the average coherence of $\mathbf{M}$ (after normalizing the columns) becomes

$$\nu = \frac{1}{n-1} \max_{i \in [n]} \left| \sum_{j \neq i} \langle f_i, f_j \rangle \right| = \frac{1}{md(n-1)} \max_{i \in [n]} \left| \sum_{j \neq i} \sum_{t=1}^{m} \chi_t(g_i^{-1} g_j) \right| \tag{4.60}$$

$$= \frac{1}{md(n-1)} \left| \sum_{g \neq 1} \sum_{t=1}^{m} \chi_t(g) \right| \tag{4.61}$$

$$= \frac{1}{md(n-1)} \left| \sum_{t=1}^{m} \left(\sum_{g \neq 1} \chi_t(g) \right) \right|. \tag{4.62}$$

Now from basic character theory (see for example [87]), we know that for any character χ_t of a *nontrivial* irreducible representation, we have the relation

$$\frac{1}{|G|} \sum_{g \in G} \chi_t(g) = 0. \tag{4.63}$$

This is due to the orthogonality of irreducible characters, and the above sum is simply the inner product between χ_t and the trivial character. But this equation gives us

$$\sum_{g \neq 1} \chi_t(g) = -\chi_t(1) = -d, \tag{4.64}$$

since $chi_t(1)$ is the degree of the representation ρ_t. Thus,

$$\nu = \frac{1}{md(n-1)} \left| \sum_{t=1}^{m} (-d) \right| = \frac{md}{md(n-1)} = \frac{1}{n-1}. \tag{4.65}$$

Now from the Welch bound, $\mu \geq \sqrt{\frac{n-md}{md(n-1)}}$. Thus, to show that $\nu \leq \frac{\mu}{\sqrt{md}}$ it is sufficient to show that $\frac{1}{n-1} \leq \frac{1}{\sqrt{md}} \sqrt{\frac{n-md}{md(n-1)}}$, or equivalently that

$$md \leq \sqrt{(n-md)(n-1)}. \tag{4.66}$$

But since $n - 1 \geq n - md$, we have $\sqrt{(n - md)(n - 1)} \geq n - md$, so (4.66) is satisfied provided that $2md \leq n$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

[75] explored the geometry of several types of frames to see when they satisfied the Coherence and Strong Coherence Properties. In particular, they stated the following theorem:

Theorem 23 ([75]). *Let F be an $n \times n$ discrete Fourier matrix, $F_{k\ell} = e^{2\pi i k\ell/n}$, $k, \ell = 0, ..., n-1$. Then let M be the submatrix formed by randomly selecting a subset of rows of F, each row independently selected with probability $\frac{m}{n}$, and then normalizing the columns. If $16 \log n \leq m \leq \frac{n}{3}$, then with probability exceeding $1 - 4n^{-1} - n^{-2}$ the worst-case coherence of M satisfies $\mu_M \leq \sqrt{\frac{118(n-m)\log n}{mn}}$.*

In Figure 4.3, we compare this bound with the bound on our harmonic frames from Theorem 9 and the Welch lower bound on coherence, in the regimes where $m = \frac{n-1}{3}$ (i.e. $\kappa = 3$) and when $m = n^{4/5}$. In both cases, we can see that the frames from our group-based construction are guaranteed to satisfy the Coherence and Strong Coherence Properties for a wider range of values of n than random harmonic frames, as suggested by Theorem 23.

4.7 Summary

In this chapter, we have generalized our methods from Chapter 3 to yield a way to select rows of the group Fourier matrix of a finite group G to produce frames with low coherence. By choosing the rows corresponding to the image of a representation under a subgroup of $Aut(G)$, we can reduce the number of distinct inner product values which arise between our frame elements. By exploiting the tightness of the resulting frames, we identified cases in which the coherence comes very close to the Welch lower bound.

We have demonstrated that our method is particularly effective when G is a subgroup or quotient of a group of matrices with entries in a finite field. This is a consequence of the manner in which the field automorphisms permute the elements of G. It is certainly possible that other groups of automorphisms of G can lead to even better coherence when applying our method, though these remain to be explored.

Furthermore, we emphasize that using the character table of G to identify suitable representations to use in our frame allows us to avoid dealing with the explicit forms of the matrices involved in the representations. These matrices are often quite large in dimension and tedious to construct, particularly in the case of the special linear groups we examined in Section 4.5.1. While exploiting the character table makes coherence calculations relatively painless, however, it is ultimately necessary to use the representation matrices to construct the actual frame vectors. It is desirable to find a class
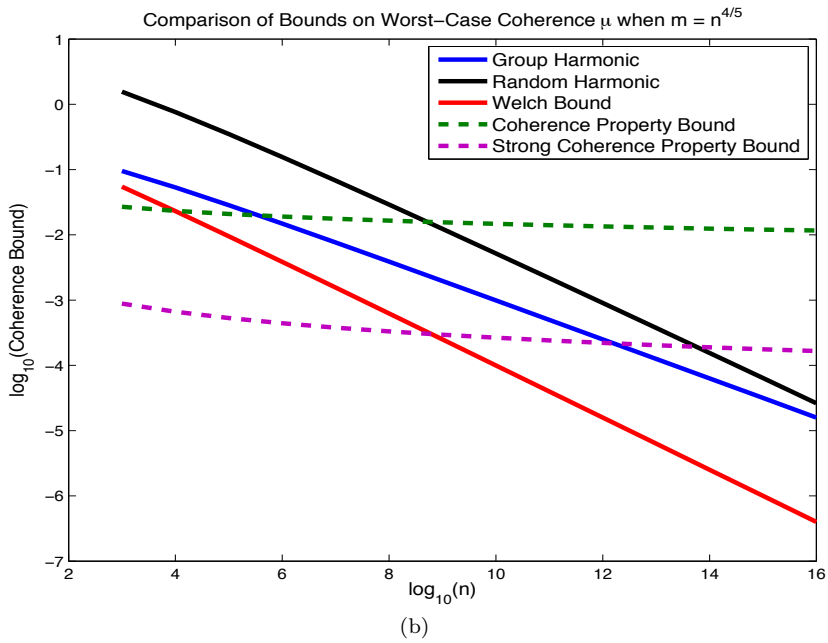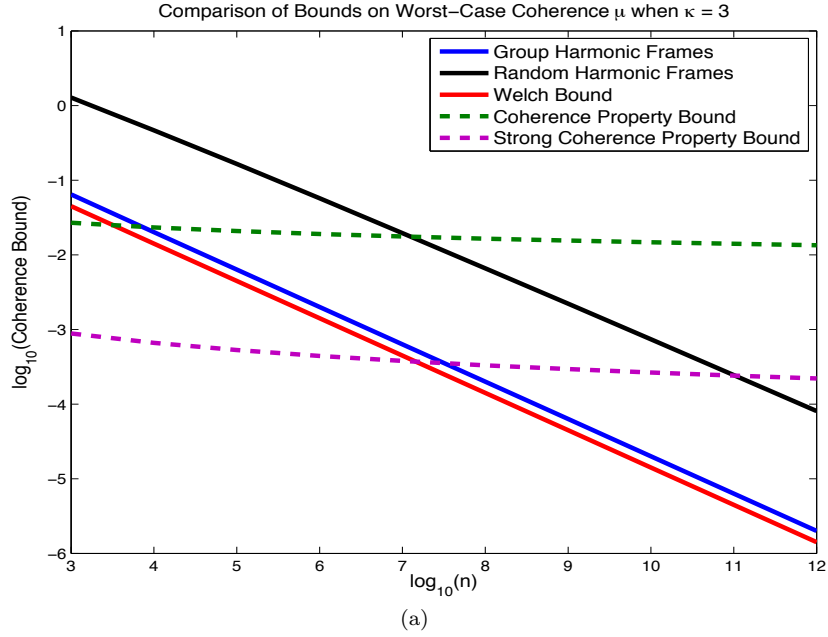
Figure 4.3: Comparison of the upper bounds on the coherence of $m \times n$ harmonic frames using our group construction (from Theorem 9) versus choosing rows randomly from a DFT matrix (Theorem 23). In 4.3(a), $\kappa = \frac{n-1}{m} = 3$, while in 4.3(b), $m = n^{4/5}$. In both these regimes, the frames from our constructions are guaranteed to satisfy both the Coherence Property and the Strong Coherence Property for smaller dimensions than randomly chosen harmonic frames.

of groups with uncomplicated representations that allow us to build low-coherence frames in a wide variety of dimensions.

Chapter 5

Coding With Constraints: Distance Bounds and Systematic Constructions

5.1 Introduction: Coding with Constraints

We consider a scenario in which we must encode s message symbols using a length n error-correcting code subject to a set of encoding constraints. Specifically, each coded symbol is a function of only a subset of the message symbols. This setup arises in various situations such as in the case of a sensor network in which each sensor can measure a certain subset of a set of parameters. The sensors would like to collectively encode the readings to allow for the possibility of measurement errors. Another scenario is one in which a client wishes to download data files from a set of servers, each of which stores information about a subset of the data files. The user should be able to recover all of the data even in the case when some of the file servers fail. Ideally, the user should also be able to download the files faster in the absence of server failures. To protect against errors, we would like the coded symbols to form an error-correcting code with reasonably high minimum distance. On the other hand, efficient download of data is permitted when the error-correcting code is of systematic form. Therefore, in this chapter, we present an upper bound on the minimum distance of an error-correcting code when subjected to encoding constraints, reminiscent of the cut-set bounds presented in [33]. In certain cases, we provide a code construction that achieves this bound. Furthermore, we refine our bound in the case that we demand a systematic linear error-correcting code, and present a construction that achieves the bound. In both cases, the codes can be decoded efficiently due to the fact that our construction utilizes Reed-Solomon codes.

5.1.1 Prior Work

The problem of constructing error-correcting codes with constrained encoding has been addressed by a variety of authors. Dau et al. [28–30] considered the problem of finding linear MDS codes with constrained generator matrices. They have shown that, under certain assumptions, such codes exist over large enough finite fields, as well as over small fields in a special case. A similar problem known as the weakly secure data exchange problem was studied in [111], [112]. The problem deals with a set of users, each with a subset of messages, who are interested in broadcasting their information securely when an eavesdropper is present. In particular, the authors of [112] conjecture the existence of secure codes based on Reed-Solomon codes and present a randomized algorithm to produce them. The problem was also considered in the context of multisource multicast network coding in [33,50,51]. In [51], the capacity region of a simple multiple access network with three sources is achieved using Reed-Solomon codes. An analogous result is derived in [50] for general multicast networks with 3 sources using Gabidulin codes.

There has been a recent line of work involving what are known as locally repairable codes (LRCs), in which every parity symbol is a function of a predetermined set of data symbols. Codes with local repair properties were described as early as 2007 in the works of [24, 54, 58]. In [48], Gopalan et al introduced bounds on code distance in terms of the locality constraints of LRCs, and since then there have been a number of new specific code constructions and extensions of these bounds [62,76,78,81,92]. Our work will also include theoretical distance bounds reminiscent of those in [48]. Another recent paper is that of Mazumdar [74] in which code symbols are represented as vertices of a partially connected graph. Each code symbol is a function of its neighbors and, if erased, can be recovered from them. Our code also utilizes a graph structure, though solely to describe the encoding procedure. In other words, there is not necessarily a notion of an individual code symbol being repairable from a local subset of the other code symbols.

5.2   Problem Setup

Let q be a prime power, and s a positive integer. Our task is to encode a set of q^s messages, represented as each of the s-dimensional vectors over the finite field $\mathbb{F}_q$ of size q. As such, we will refer to a message as such a vector $\mathbf{m} = [m_1, ..., m_s] \in \mathbb{F}_q^s$. We would like to map each of these message vectors to a codeword consisting of n symbols each coming from an alphabet of size q, again represented as a vector $\mathbf{c} = [c_1, ..., c_n] \in \mathbb{F}_q^n$. Here, $n \geq s$. In our case, each of the symbols c_i is a function of only a *subset* of the message symbols $\{m_i\}_{i=1}^s$. We will denote this subset as $\mathcal{I}_{c_i}$. For example, the m_i could represent incoming signals to a sensor array, and each c_i could represent a sensor with access to only

Figure 5.1: A bipartite graph representing the coding constraints. Here, there are 3 message symbols and 7 codeword symbols. Each c_i is a function of the message symbols to which it is connected. For example, c_1 is a function of $\{m_1, m_2\}$.

some of these signals. Alternatively, the m_i could be data files which must be stored in each member of a set of file servers. Each server might only have access to a local set of data files and seeks to store a function of these files, represented as c_i. In either case, we would like to select our encoding scheme subject to these constraints so that the original message $\mathbf{m}$ can be determined from $\mathbf{c}$ even in the case that some of the symbols c_i are erased or corrupted.

We can represent our encoding constraints in the form of a bipartite graph, $G = (\mathcal{M}, \mathcal{V}, \mathcal{E})$, with vertex sets $\mathcal{M}$ of size s and $\mathcal{V}$ of size n representing the message symbols and codeword symbols respectively. As such, we will label the vertices in $\mathcal{M}$ as $\{m_i\}_{i=1}^s$ and the vertices in $\mathcal{V}$ as $\{c_i\}_{i=1}^n$. A pair $(m_i, c_j) \in \mathcal{M} \times \mathcal{V}$ is in the edge set $\mathcal{E}$ if and only if $m_i \in \mathcal{I}_{c_j}$, that is, c_j is a function of m_i. For example, in Figure 5.1 we have $\mathcal{I}_{c_1} = \{m_1, m_2\}$ and $\mathcal{I}_{c_2} = \{m_2\}$.

Let us quickly establish some notation. For any subset $m_i \in \mathcal{M}$, we will let $\mathcal{N}(m_i)$ denote the neighborhood of m_i in $\mathcal{V}$:

$$\mathcal{N}(m_i) := \{c_j \in \mathcal{V} \ : \ (m_i, c_j) \in \mathcal{E}\}.$$

Likewise, we will consider neighborhoods of arbitrary subsets $\mathcal{M}' \subseteq \mathcal{M}$:

$$\mathcal{N}(\mathcal{M}') := \bigcup_{m_i \in \mathcal{M}'} \mathcal{N}(m_i).$$

We will denote neighborhoods of elements $c_j \in \mathcal{V}$ and subsets $\mathcal{V}' \subseteq \mathcal{V}$ similarly. With this notation, it is clear that $\mathcal{N}(c_j) = \mathcal{I}_{c_j}$.

When the m_i are assigned values from $\mathbb{F}_q$, then each c_j has an associated function $f_j : \mathbb{F}_q^{|\mathcal{N}(c_j)|} \to \mathbb{F}_q$ which maps the set of values $\{m_i \in \mathcal{N}(c_j)\}$ to a value of c_j. By abuse of notation, we will sometimes simply write $c_j = f_j(\mathbf{m})$, with the understanding that c_j depends only on the coordinates of $\mathbf{m}$ which are in $\mathcal{N}(c_j)$. If we let $[\mathbf{c}]_{\mathcal{J}}$ be the subvector of $\mathbf{c}$ with elements indexed by $\mathcal{J} \subseteq \{1, ..., n\}$, then

we will write $f_{\mathcal{J}} : \mathbb{F}_q^{|\mathcal{N}(\{c_j \ : \ j \in \mathcal{J}\})|} \to \mathbb{F}_q^{|\mathcal{J}|}$ for the function which sends $[\mathbf{m}]_{\mathcal{N}(c_j \ : \ j \in \mathcal{J})}$ to the vector $[\mathbf{c}]_{\mathcal{J}} = (f_j(\mathbf{m}), j \in \mathcal{J})$. Under this notation, we have $\mathbf{c} = f_{[n]}(\mathbf{m})$, where $[n] := \{1, ..., n\}$. If we restrict the functions $f_j(\cdot)$ to be linear, then $\mathcal{C}$ becomes a linear code.

If we define

$$\mathcal{C} := \{\mathbf{c} \in \mathbb{F}_q^n \ : \ \exists \mathbf{m} \in \mathbb{F}_q^s \text{ s.t. } \mathbf{c} = f_{[n]}(\mathbf{m})\},$$

then $\mathcal{C}$ is the set of all valid codewords, which is an error-correcting code of length n and size at most q^s. Let $d(\mathcal{C})$ be the minimum distance of this code:

$$d(\mathcal{C}) := \min_{\{\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C} \ : \ \mathbf{c}_1 \neq \mathbf{c}_2\}} d_{\mathcal{H}}(\mathbf{c}_1, \mathbf{c}_2),$$

where $d_{\mathcal{H}}(\cdot, \cdot)$ denotes the Hamming distance between two vectors. In the case that our $f_j(\cdot)$ are linear, we have the following well-known equivalent definition of the code's minimum distance:

Lemma 14. *If $\mathcal{C}$ is a linear code, then $d(\mathcal{C}) = \min_{\mathbf{c} \in \mathcal{C}} ||\mathbf{c}||_{\mathcal{H}}$, where $||\mathbf{c}||_{\mathcal{H}}$ is the Hamming weight (the number of nonzero entries) of $\mathbf{c}$.*

Proof. Since $\mathcal{C}$ is linear, the all-zero codeword $\mathbf{0}$ is in $\mathcal{C}$. Also, for any $\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C}$, we have that $\mathbf{c}_1 - \mathbf{c}_2 \in \mathcal{C}$. The result now follows from noting that $d_{\mathcal{H}}(\mathbf{c}_1, \mathbf{c}_2) = d_{\mathcal{H}}(\mathbf{c}_1 - \mathbf{c}_2, \mathbf{0}) = ||\mathbf{c}_1 - \mathbf{c}_2||_{\mathcal{H}}$, so $d(\mathcal{C}) \leq \min_{\mathbf{c} \in \mathcal{C}} ||\mathbf{c}||_{\mathcal{H}}$. On the other hand, for any $\mathbf{c} \in \mathcal{C}$ we have $||\mathbf{c}||_{\mathcal{H}} = d_{\mathcal{H}}(\mathbf{c}, \mathbf{0})$, so the reverse inequality also holds. $\square$

Let us assume our functions $f_j(\cdot)$ are linear, and $\mathcal{C}$ a linear code. This means that for each $j \in \{1, ..., n\}$, there is a column vector $\mathbf{g}_j \in \mathbb{F}_q^{s \times 1}$, such that $c_j = f_j(\mathbf{m}) = \mathbf{m} \cdot \mathbf{g}_j$. Since c_j is a function of only the $m_i \in \mathcal{N}(c_j)$, we see that the support of $\mathbf{g}_j$ must lie in the entries indexed by the elements of $\mathcal{N}(c_j)$. If we concatenate the columns $\mathbf{g}_j$, we form the *generator matrix* of $\mathcal{C}$,

$$\mathbf{G} = [\mathbf{g}_1, \ldots, \mathbf{g}_n] \in \mathbb{F}_q^{s \times n}.$$

For any message vector $\mathbf{m}$, the corresponding codeword will be given by $\mathbf{c} = \mathbf{m}\mathbf{G}$.

We can describe the support of $\mathbf{G}$ by examining the adjacency matrix $\mathbf{A} \in \{0, 1\}^{s \times n}$ of the bipartite graph $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ describing our code:

$$[\mathbf{A}]_{i,j} := \begin{cases} 1 & \text{if } (m_i, c_j) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases}. \tag{5.1}$$

Thus the j^{th} column of $\mathbf{A}$ has support precisely on $\mathcal{N}(c_j)$. Hence by our discussion above, a matrix $\mathbf{G}$ will be a "valid" generator matrix for a code $\mathcal{C}$ with constraints defined by the bipartite graph $\mathcal{G}$ if

the support of $\mathbf{G}$ is a subset of the support of $\mathbf{A}$. In the example given in Figure 5.1, our adjacency matrix is

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}. \tag{5.2}$$

The choice of support entries for a valid generator matrix $\mathbf{G}$ determines both the rank of the code (which can be between 0 and s) and its minimum distance. In general, we seek to find a valid generator matrix which produces a full-rank code (yielding the maximum number of distinct codewords, q^s) while simultaneously maximizing the minimum distance of our code (allowing us to correctly determine a codeword even in the presence of up to $\lceil \frac{d(\mathcal{C})}{2} \rceil - 1$ errors). Furthermore, we would like to ensure efficient methods to decode our codewords in the presence of errors. To this end, we will look to constructing our codes from Reed-Solomon codes, a common class of error-correcting codes with efficient decoding algorithms.

5.3 Minimum Distance Bounds for General and Constrained Codes

While it can be difficult in general to determine the optimal minimum distance of a constrained code, we have a handful of tools at our disposal to help bound it. For instance, we can always appeal to the well-known Singleton bound [89]:

Theorem 24 (Singleton Bound). *If $\mathcal{C}$ is a length-n code over an alphabet of size q, then $|\mathcal{C}| \leq q^{n-d(\mathcal{C})+1}$.*

Proof. Take any subset $\mathcal{I} \subseteq \{1, ..., n\}$ of size $d(\mathcal{C}) - 1$, and let $\mathcal{C}_{\mathcal{I}^c}$ denote the set $\{[\mathbf{c}]_{[n] \setminus \mathcal{I}} \ : \ \mathbf{c} \in \mathcal{C}\}$, i.e., the vectors of $\mathcal{C}$ with their entries in $\mathcal{I}$ removed. As such, the elements of $\mathcal{C}_{\mathcal{I}^c}$ are subvectors of length $n - d(\mathcal{C}) + 1$, so there can be no more than $q^{n-d(\mathcal{C})+1}$ of them. Since any two vectors in $\mathcal{C}$ differ in at least $d(\mathcal{C})$ entries, all of the vectors in $\mathcal{C}_{\mathcal{I}^c}$ must be distinct, and $|\mathcal{C}| = |\mathcal{C}_{\mathcal{I}^c}|$, and we are done. $\square$

In our framework, we would like to have a distinct codeword for each of our messages $\mathbf{m} \in \mathbb{F}_q^s$, hence we would like to have $\mathcal{C} = q^s$. We will accordingly rephrase the Singleton bound in the following form:

Corollary 5. *Let $\mathcal{C}$ be a length-n code over an alphabet of size q such that $|\mathcal{C}| = q^s$. Then $d(\mathcal{C}) \leq n - s + 1$.*

Proof. This follows directly from Theorem 24. □

A linear code which meets the Singleton bound with equality is called a *maximum distance separable* (MDS) code, and has the following alternative characterization:

Theorem 25. *Let $\mathcal{C} \subset \mathbb{F}_q^n$ be a linear code of dimension s and generator matrix $\mathbf{G} \in \mathbb{F}_q^{s \times n}$. Then $\mathcal{C}$ is an MDS code if and only if every s columns of $\mathbf{G}$ are full rank.*

Proof. Assume every s columns of $\mathbf{G}$ are full rank. Consider two distinct codewords $\mathbf{c}_1$ and $\mathbf{c}_2$ in $\mathcal{C}$, and let $\mathbf{m}_1$ and $\mathbf{m}_2$ be corresponding message vectors in $\mathbb{F}_q^s$ such that $\mathbf{c}_i = \mathbf{m}_i \mathbf{G}$, $i = 1, 2$. Fix a subset $\mathcal{I} \subseteq [n]$, of size $d \leq n - s$, and remove the coordinates of $\mathcal{I}$ from $\mathbf{c}_1$ and $\mathbf{c}_2$ to form $[\mathbf{c}_1]_{\mathcal{I}^c}$ and $[\mathbf{c}_2]_{\mathcal{I}^c}$ respectively. Likewise, remove the columns of $\mathbf{G}$ indexed by $\mathcal{I}$ to form $\mathbf{G}_{\mathcal{I}^c}$. Since $\mathbf{G}_{\mathcal{I}^c}$ has at least s columns, it must have full rank s. Thus, if $[\mathbf{c}_1]_{\mathcal{I}^c}$ and $[\mathbf{c}_2]_{\mathcal{I}^c}$ are identical, it implies that $\mathbf{m}_1 = \mathbf{m}_2$, and $\mathbf{c}_1 = \mathbf{c}_2$. Thus, the distance of our code is greater than $n - s$, so it must achieve the Singleton bound.

Conversely, suppose our code is MDS, and fix a subset $\mathcal{S} \subseteq [n]$ of size s. Let $\mathbf{G}_{\mathcal{S}} \in \mathbb{F}_q^{s \times s}$ be the submatrix of $\mathbf{G}$ consisting of the columns indexed by $\mathcal{S}$. Since $\mathcal{C}$ has minimum distance $n - s + 1$, then for any distinct $\mathbf{m}_1, \mathbf{m}_2 \in \mathbb{F}_q^s$, the codewords $\mathbf{c}_1 := \mathbf{m}_1 \mathbf{G}$ and $\mathbf{c}_2 := \mathbf{m}_2 \mathbf{G}$ must differ within the s coordinates indexed by $\mathcal{S}$. That is, $\mathbf{m}_1 \mathbf{G}_{\mathcal{S}} \neq \mathbf{m}_2 \mathbf{G}_{\mathcal{S}}$, so the columns indexed by $\mathcal{S}$ are full rank. □

For constrained codes, it turns out the Singleton bound is often rather loose. In this case, we can derive a tighter class of bounds reminiscent of those in [32].

Theorem 26. *Let $\mathcal{C} \subseteq \mathbb{F}_q^n$ be a code which is constrained by the bipartite graph $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$, that is, for each $\mathbf{c} = (c_1, ..., c_n) \in \mathcal{C}$ there is some $\mathbf{m} = (m_1, ..., m_s) \in \mathbb{F}_q^s$ such that c_j is a function of the set $\{m_i \in \mathcal{N}(c_j)\}$. Assume that for any subset $\mathcal{I} \subseteq \mathcal{M}$, we have $|\mathcal{N}(\mathcal{I})| \geq |\mathcal{I}|$. Then for any subset $\mathcal{I} \subseteq \mathcal{M}$, the minimum distance of $\mathcal{C}$ satisfies*

$$d(\mathcal{C}) \leq |\mathcal{N}(\mathcal{I})| - |\mathcal{I}| + 1. \tag{5.3}$$

Proof. Fix any set $\mathcal{I} \subseteq \mathcal{M}$. This proof is essentially a variation of the proof of the Singleton bound when restricted to the code induced by the subvectors $[\mathbf{c}]_{\mathcal{N}(\mathcal{I})}, \mathbf{c} \in \mathcal{C}$, consisting of the codewords in $\mathcal{C}$ with their coordinates removed outside of the set $\mathcal{N}(\mathcal{I})$. To be explicit, consider the set $S_{\mathcal{I}} \subseteq \mathbb{F}_q^s$ of vectors which are zero outside of the indices in $\mathcal{I}$:

$$S_{\mathcal{I}} = \{(m_1, ..., m_s \in \mathbb{F}_q^s \; : \; m_i = 0 \forall i \notin \mathcal{I}\}.$$

Figure 5.2: Partitions of $\mathcal{M}$ and of $\mathcal{V}$ used in the proof of Theorem 26. The set $\mathcal{N}(\mathcal{I})$ is a function of both $\mathcal{I}$ and $\mathcal{I}^{\mathsf{c}}$, while the set $\mathcal{N}(\mathcal{I})^{\mathsf{c}}$ is a function of $\mathcal{I}^{\mathsf{c}}$ only.

It is clear that $S_{\mathcal{I}}$ has $q^{|\mathcal{I}|}$ elements. By examining the bipartite graph, we can see that the symbols $c_j, j \in \mathcal{N}(\mathcal{I})^c$ are only a function of the message variables $m_i, i \notin \mathcal{I}$ (see Fig. 5.2), hence the subcode $\mathcal{C}_{\mathcal{I}} := \{\mathbf{c} \in \mathcal{C} \ : \ \mathbf{c} = f_{[n]}(\mathbf{m}), \ \mathbf{m} \in S_{\mathcal{I}}\}$ must have constant values in the indices corresponding to $\mathcal{N}(\mathcal{I})^c$. Furthermore, if $|\mathcal{N}(\mathcal{I})| \geq |\mathcal{I}|$, and we consider any $|\mathcal{I}|-1$ of the indices corresponding to $\mathcal{N}(\mathcal{I})$, then by the Pigeonhole Principal, there must be distinct $\mathbf{m}_1$ and $\mathbf{m}_2$ in $S_{\mathcal{I}}$ such that $\mathbf{c}_1 := f_{[n]}(\mathbf{m}_1)$ and $\mathbf{c}_2 := f_{[n]}(\mathbf{m}_2)$ have the same values in these $|\mathcal{I}| - 1$ indices. Thus, since $\mathbf{c}_1$ and $\mathbf{c}_2$ are both in $\mathcal{C}_{\mathcal{I}}$, they have at least $|\mathcal{N}(\mathcal{I})^c| + |\mathcal{I}| - 1$ entries in common, hence can have Hamming distance at most $n - (|\mathcal{N}(\mathcal{I})^c| + |\mathcal{I}| - 1) = |\mathcal{N}(\mathcal{I})| - |\mathcal{I}| + 1$, and we are done. $\qquad\square$

Remark: In the case where $|\mathcal{N}(\mathcal{I})| < |\mathcal{I}|$ for some subset $\mathcal{I}$, then the proof above can produce two distinct vectors $\mathbf{m}_1$ and $\mathbf{m}_2$ in $S_{\mathcal{I}}$, yielding $\mathbf{c}_1$ and $\mathbf{c}_2$ in $\mathcal{C}_{\mathcal{I}}$ which have the same entries in all coordinates of $\mathcal{N}(\mathcal{I})$, and hence $\mathbf{c}_1 = \mathbf{c}_2$ so our code has minimum distance equal to 0.

As a direct corollary, we have

Corollary 6. *Let $\mathcal{C}$ be a code constrained by the bipartite graph $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$, as in Theorem 26. Then the minimum distance $d(\mathcal{C})$ satisfies*

$$d(\mathcal{C}) \leq \min_{\mathcal{M}' \subseteq \mathcal{M}} |\mathcal{N}(\mathcal{M}')| - |\mathcal{M}'| + 1. \tag{5.4}$$

As we can see, when $\mathcal{M}'$ is taken to be the full set $\mathcal{M}$ in Corollary 6, then we obtain the original Singleton bound (assuming $\mathcal{N}(\mathcal{M})$ is the entire set $\mathcal{V}$, which we may do without loss of generality). In general, however, it remains an interesting task to provide constructions of codes $\mathcal{C}$ that meet the constraints imposed by arbitrary graphs $\mathcal{G}$ which 1) achieve the upper bound on distance in Corollary 6 with equality, and 2) have efficient decoding algorithms to recover a message $\mathbf{m} \in \mathbb{F}_q^s$ from the vector $\mathbf{c} := f_{[n]}(\mathbf{m}) \in \mathcal{C}$ even in the presence of errors. Our method will be to attempt to construct $\mathcal{C}$ as a subcode of a Reed-Solomon code, which is a well-known MDS code with known fast decoding algorithms. We will briefly review Reed-Solomon codes in the next section.

5.4 Subcodes of Reed-Solomon Codes

While there are several equivalent ways to define Reed-Solomon codes, we will use the original definition from [82], which we will see fits quite naturally into our current framework. Let n and k be integers, with $n \geq k$, and q a power of a prime. To any vector

$$\mathbf{m} = (m_1, ..., m_k) \in \mathbb{F}_q^k,$$

we can associate the polynomial

$$m(x) := \sum_{i=1}^{k} m_i x^{i-1}$$

of degree at most $k - 1$. If we fix distinct elements $\{\alpha_1, ..., \alpha_n\} \subseteq \mathbb{F}_q$, which we will call the *defining set* of our code, then a $[n, k]_q$ Reed-Solomon code is defined as

$$\mathcal{C}_{RS} := \{(m(\alpha_1), ..., m(\alpha_n)) \in \mathbb{F}_q^n \ : \ \deg(m(x)) < k\},$$

which is a k-dimensional subspace of $\mathbb{F}_q^n$.

If we define the matrix

$$\mathbf{G}_{RS} := \begin{bmatrix} 1 & \dots & 1 \\ \alpha_1 & \dots & \alpha_n \\ \vdots & \ddots & \vdots \\ \alpha_1^{k-1} & \dots & \alpha_n^{k-1} \end{bmatrix}, \tag{5.5}$$

then we can express any codeword in $\mathbf{c} \in \mathcal{C}_{RS}$ in the form $\mathbf{c} = \mathbf{m}\mathbf{G}_{RS}$ for some $\mathbf{m} \in \mathbb{F}_q^k$, so we see that $\mathbf{G}_{RS}$ is the generator matrix for $\mathcal{C}_{RS}$. Since $\mathbf{G}_{RS}$ is Vandermonde, any k of its rows are full rank, and we see that $\mathcal{C}_{RS}$ is an MDS code by Theorem 25. Thus, its minimum distance achieves the Singleton bound, and we have $d(\mathcal{C}_{RS}) = n - k + 1$.

Given a bipartite graph $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ defining a set of constraints for our code, where $|\mathcal{V}| = n$ and $|\mathcal{M}| = s \leq k$, we can find a subspace $\mathcal{C}$ of $\mathcal{C}_{RS}$ which forms a valid code for the constraints imposed by $\mathcal{G}$ as follows: For each $i \in \{1, ..., s\}$, let $\mathcal{N}(m_i)^c := \mathcal{V} \setminus \mathcal{N}(m_i)$ be the set of codeword symbols c_j to which m_i is *not* connected in $\mathcal{G}$. Identifying α_j with the node c_j for each $j = 1, ..., n$, define a polynomial $t_i(x)$ of degree at most $k - 1$ such that $t_i(x)$ is divisible by $\prod_{\alpha_j \in \mathcal{N}(m_i)^c} (x - \alpha_j)$. If we write $t_i(x) = \sum_{i'=1}^{k} t_{i,i'} x^{i'-1}$ and identify $t_i(x)$ with the vector $\mathbf{t}_i := [t_{i,1}, ..., t_{i,k}] \in \mathbb{F}_q^k$, we see that

the codeword

$$\mathbf{c}_i := \mathbf{t}_i \mathbf{G}_{RS} = [t_i(\alpha_1), ..., t_i(\alpha_n)] \tag{5.6}$$

has zeros in the entries corresponding to the elements $\alpha_j \in \mathcal{N}(m_i)^c$. Thus, if we stack the rows $\mathbf{c}_i$ to form the matrix

$$\mathbf{G} := \begin{bmatrix} \mathbf{c}_1 \\ \vdots \\ \mathbf{c}_s \end{bmatrix} = \begin{bmatrix} \mathbf{t}_1 \\ \vdots \\ \mathbf{t}_s \end{bmatrix} \cdot \begin{bmatrix} 1 & \cdots & 1 \\ \alpha_1 & \cdots & \alpha_n \\ \vdots & \ddots & \vdots \\ \alpha_1^{k-1} & \cdots & \alpha_n^{k-1} \end{bmatrix}, \tag{5.7}$$

then we see that $\mathbf{G}$ has the same zero-pattern as the adjacency matrix $\mathbf{A}$ of $\mathcal{G}$ defined in (5.1). Thus, $\mathbf{G}$ is the generator matrix for a linear code $\mathcal{C}$ which is valid for the graph $\mathcal{G}$. For convenience, we will define the matrix of the $\mathbf{t}_i$ to be

$$\mathbf{T} := \begin{bmatrix} \mathbf{t}_1 \\ \vdots \\ \mathbf{t}_s \end{bmatrix}, \tag{5.8}$$

and write the above equation in the compact form $\mathbf{G} = \mathbf{T}\mathbf{G}_{RS}$. The dimension of the code $\mathcal{C}$ is equal to the row rank of $\mathbf{T}$, which is determined by how we choose the polynomials $t_i(x)$ subject to the aforementioned requirements that $\deg(t_i(x)) < k$ and $\prod_{\alpha_j \in \mathcal{N}(m_i)^c}(x - \alpha_j)$ divides $t_i(x)$. Furthermore, since $\mathcal{C}$ is a subspace of a Reed-Solomon code, we can apply pre-existing efficient decoders to recover any message $\mathbf{m} \in \mathbb{F}_q^s$ from the codeword $\mathbf{c} := \mathbf{m}\mathbf{G}$, even in the presence of errors. Some well-known such decoders include the Peterson [77], the Berlekamp-Massey [5, 71], and the Welch-Berlekamp [7] algorithms.

5.5 Systematic Codes

In many scenarios, it is desirable to have the s symbols our original message $\mathbf{m}$ appear as a subset of the symbols of the corresponding codeword $\mathbf{c}$. This allows $\mathbf{m}$ to be retrieved immediately in the absence of errors in $\mathbf{c}$ without alluding to a lookup table, inverting the function $f_{\mathcal{M}} : \mathbf{m} \mapsto \mathbf{c}$, or performing any other method of decoding which could be costly in computation or storage. For example, if our original message symbols $\mathbf{m} = [m_1, ..., m_s]$ collectively represent a collection of data files, the codeword symbols $\mathbf{c} = [c_1, ..., c_n]$ could represent encoded files stored in n different servers,

where $n > s$ to protect the data in the case of some servers crashing. Suppose $c_i = m_i$ for $i = 1, ..., s$. Then in the case where crashes occur only in the servers corresponding to $c_{s+1}, ..., c_n$, we can still easily access our original data $\mathbf{m} = [c_1, ..., c_s]$, which can be used to quickly recompute the c_j in the servers which have crashed (for $j > s$).

Definition 6. *Let s and n be integers, $n \geq s$, and q a power of a prime. For any vector $\mathbf{c} \in \mathbb{F}_q^n$ and any subset $\mathcal{I} \subseteq [n]$, let $[\mathbf{c}]_{\mathcal{I}}$ denote the subvector of $\mathbf{c}$ in the entries indexed by $\mathcal{I}$. Let $f : \mathbb{F}_q^s \to \mathbb{F}_q^n$ be a function such that for some fixed subset $\mathcal{I}_{sys} \subseteq [n]$ of size s, we have $[f(\mathbf{m})]_{\mathcal{I}_{sys}} = \mathbf{m}$. Then the set $\mathcal{C} = \{\mathbf{c} \in \mathbb{F}_q^n : \mathbf{c} = f(\mathbf{m}), \mathbf{m} \in \mathbb{F}_q^s\}$ is called a systematic code, or a code in systematic form. For any $\mathbf{c} = [c_1, ..., c_n] \in \mathcal{C}$, the symbols c_j, $j \in \mathcal{I}_{sys}$, are called the systematic symbols of $\mathbf{c}$, and the remaining c_j, $j \notin \mathcal{I}_{sys}$, are the parity symbols.*

If $\mathcal{C}$ is a *linear* code with generator matrix $\mathbf{G} \in \mathbb{F}_q^{s \times n}$, then $\mathcal{C}$ being systematic is equivalent to the columns of the $s \times s$ identity matrix $\mathbf{I}_s$, arising as a subset of the columns of $\mathbf{G}$. Let us examine what this means in the context of codes with constraints. As before, let $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ be a bipartite graph with $|\mathcal{M}| = s$ and $|\mathcal{V}| = n$, where we identify the message symbols $m_1, ..., m_s$ with the vertices of $\mathcal{M}$ and the codewords symbols $c_1, ..., c_n$ with those of $\mathcal{V}$. For each $c_j \in \mathcal{V}$, we have an associated function $c_j = f_j(\{m_i \in \mathcal{N}(c_j)\})$. Thus if c_j is a systematic symbol in a systematic code $\mathcal{C}$ such that $c_j = m_i$, it must be that $m_i \in \mathcal{N}(c_j)$. In other words, $(m_i, c_j) \in \mathcal{E}$. On this note, we refer to the following definition from basic graph theory:

Definition 7. *Let $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ be a bipartite graph. A subset $\tilde{\mathcal{E}} \subset \mathcal{E}$ is called a matching for $\mathcal{G}$ if no two edges in $\tilde{\mathcal{E}}$ share a common vertex. $\tilde{\mathcal{E}}$ is said to be a maximal matching if it is not a proper subset of any other matching. A subset $\mathcal{S} \subseteq \mathcal{M} \cup \mathcal{V}$ is said to be covered by $\tilde{\mathcal{E}}$ if each vertex in $\mathcal{S}$ is incident to an edge in $\tilde{\mathcal{E}}$.*

Under this terminology, the following is clear:

Lemma 15. *Let $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ be a bipartite graph, with $|\mathcal{M}| = s$ and $|\mathcal{V}| = n$. Then there exists a systematic code $\mathcal{C} = \{\mathbf{c} \in \mathbb{F}_q^n : \mathbf{c} = f_{[n]}(\mathbf{m}), \mathbf{m} \in \mathbb{F}_q^s\}$ for some $f_{[n]}(\cdot)$ which fits the constraints of $\mathcal{G}$ only if there is an $\mathcal{M}$-covering matching $\tilde{\mathcal{E}} \subseteq \mathcal{E}$ for $\mathcal{G}$.*

Proof. Let $\mathcal{I}_{sys} \subseteq [n]$ be the indices of the systematic symbols of each $\mathbf{c} = [c_1, ..., c_n] \in \mathcal{C}$. For each $j \in \mathcal{I}_{sys}$, let $i_j \in [s]$ be such that $c_j = m_{i_j}$. Since $\mathcal{C}$ is constrained by $\mathcal{G}$, we necessarily have $(m_{i_j}, c_j) \in \mathcal{E}$. Note that by the nature of the systematic code, for any two distinct j_1 and j_2 in $\mathcal{I}_{sys}$, we necessarily have $i_{j_1} \neq i_{j_2}$. Furthermore, for any $m_i \in \mathcal{M}$, there must be some $j \in \mathcal{I}_{sys}$ such that $i = i_j$. Thus, the set of edges $\tilde{\mathcal{E}} := \{(m_{i_j}, c_j) : j \in \mathcal{I}_{sys}\}$ is a matching for $\mathcal{G}$ which covers $\mathcal{M}$. $\square$

A crucial tool in examining matchings is Hall's Theorem:

Theorem 27 (Hall's Theorem). *Let $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ be a bipartite graph. An $\mathcal{M}$-covering matching exists if and only if $|\mathcal{M}'| \leq |\mathcal{N}(\mathcal{M}')|$ for all subsets $\mathcal{M}' \subseteq \mathcal{M}$.*

Proof. This is a well-known result in graph theory, proven by Philip Hall in [52]. An accessible proof appears on p. 53 of [67]. $\qquad\square$

5.5.1 Systematic Code Construction Using Reed-Solomon Codes

We now present a sufficient condition on our bipartite constraint graph $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ which allows us to construct a systematic code meeting our constraints which achieves the upper bound on distance from Corollary 6. Our code will be a linear subcode of a Reed-Solomon code, and will have dimension equal to $|\mathcal{M}|$. Loosely speaking, our construction relies on a sufficient amount of connectivity in $\mathcal{G}$.

Theorem 28. *Let $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ be a bipartite graph where $\mathcal{N}(\mathcal{M}) = \mathcal{V}$, with $|\mathcal{M}| = s$ and $|\mathcal{V}| = n$. Define the set $\mathcal{A} := \{c_j \in \mathcal{V} \ : \ \mathcal{N}(c_j) = \mathcal{M}\}$, the set of code symbols which are connected to all the message symbols. Let $d_{min} := \min_{\mathcal{M}' \subseteq \mathcal{M}} |\mathcal{N}(\mathcal{M}')| - |\mathcal{M}'| + 1$ and $k_{min} := n - d_{min} + 1$. Then if q is a prime power greater than or equal to n, a linear code $\mathcal{C} \subseteq \mathbb{F}_q^n$ can be constructed with a generator matrix $\mathbf{G} \in \mathbb{F}_q^{s \times n}$ in systematic form provided that $k_{min} \geq |\mathcal{V} \setminus \mathcal{A}|$.*

Proof. By our hypotheses, we have

$$n = |\mathcal{V} \setminus \mathcal{A}| + |\mathcal{A}| \leq k_{min} + |\mathcal{A}|, \tag{5.9}$$

and by our definition of k_{min}, this gives us $|\mathcal{A}| \geq d_{min} - 1$. Let $\mathcal{B} \subseteq \mathcal{A}$ be a set of size $d_{min} - 1$, and set $\mathcal{A}^* := \mathcal{A} \setminus \mathcal{B}$, $\mathcal{V}^* := \mathcal{V} \setminus \mathcal{B}$, and $\mathcal{E}^* := \{(m_i, c_j) \in \mathcal{E} \ : \ c_j \in \mathcal{V}^*\}$. Then define the corresponding subgraph of our bipartite graph, $\mathcal{G}^* = (\mathcal{M}, \mathcal{V}^*, \mathcal{E}^*)$, in which we can see that $\mathcal{A}^*$ is precisely the set of vertices in $\mathcal{V}^*$ which are connected to all of $\mathcal{M}$. Its cardinality is

$$|\mathcal{A}^*| = |\mathcal{A}| - (d_{min} - 1). \tag{5.10}$$

To avoid confusion, for any subset $\mathcal{M}' \subseteq \mathcal{M}$, we will denote the neighborhood of $\mathcal{M}'$ in $\mathcal{V}^*$ as $\mathcal{N}^*(\mathcal{M}')$, while still using the notation $\mathcal{N}(\mathcal{M}')$ to denote the neighborhood of $\mathcal{M}'$ in the entire set $\mathcal{V}$. We can express $\mathcal{N}^*(\mathcal{M}')$ as the disjoint union $(\mathcal{N}(\mathcal{M}') \setminus \mathcal{A}) \sqcup \mathcal{A}^*$, so we have

$$|\mathcal{N}^*(\mathcal{M}')| = |\mathcal{N}(\mathcal{M}') \setminus \mathcal{A}| + |\mathcal{A}^*|. \tag{5.11}$$

On the other hand, by the definition of d_{min} we have

$$|\mathcal{M}'| \leq |\mathcal{N}(\mathcal{M}')| - (d_{min} - 1) = |\mathcal{N}(\mathcal{M}') \setminus \mathcal{A}| + |\mathcal{A}| - (d_{min} - 1). \tag{5.12}$$

Combining our relations from (5.10), (5.11), and (5.12), we obtain

$$|\mathcal{M}'| \leq |\mathcal{N}^*(\mathcal{M}')|, \ \forall \mathcal{M}' \subseteq \mathcal{M}. \tag{5.13}$$

Thus we can apply Hall's Theorem to the subgraph $\mathcal{G}^*$ to find a matching $\tilde{\mathcal{E}} \subseteq \mathcal{E}^*$ which covers $\mathcal{M}$. If we let $c_{j(i)}$ be the vertex matched to m_i, then we can write this matching as $\tilde{\mathcal{E}} = \{(m_i, c_{j(i)})\}_{i=1}^s \subseteq \mathcal{E}^*$. Let $\tilde{\mathcal{V}} = \{c_{j(i)}\}_{i=1}^s$ be the subset of $\mathcal{V}^*$ which is covered by $\tilde{\mathcal{E}}$.

The symbol $c_{j(i)}$ will correspond to the systematic coordinate of our codeword which is equal to message symbol m_i. As such, any edge $(m_{i'}, c_{j(i)})$, for $i' \neq i$, is effectively ignored. As such, define the set of ignored edges

$$\mathcal{E}_{\text{neg}} := \{(m_i, c_j) \in \mathcal{E} \ : \ j \in \tilde{\mathcal{V}}, \ j \neq j(i)\}.$$

Let $\mathbf{A}_{\tilde{\mathcal{E}}}$ be the adjacency matrix of the graph $\tilde{\mathcal{G}} := (\mathcal{M}, \mathcal{V}, \mathcal{E} \setminus \mathcal{E}_{\text{neg}})$, which is the graph $\mathcal{G}$ after removing the ignored edges. Note that any any code fitting the constraints imposed by $\tilde{\mathcal{G}}$ will automatically fit those of the original graph $\mathcal{G}$. We claim that the number of zeros in any row of $\mathbf{A}_{\tilde{\mathcal{E}}}$ is at most $n - d_{min}$. Indeed, each message symbol vertex m_i is connected to one vertex in $\mathcal{V}^*$ and all $d_{min} - 1$ vertices in $\mathcal{B}$, so the corresponding row of $\mathbf{A}_{\tilde{\mathcal{E}}}$ must have at least d_{min} ones.

Now we can construct a linear code with a generator matrix $\mathbf{G}$ having the same support set as $\mathbf{A}_{\tilde{\mathcal{E}}}$, and thus meeting the constraints imposed by the graph $\tilde{\mathcal{G}}$ (and therefore $\mathcal{G}$). We will form our code as a linear subcode of a Reed-Solomon code as described in Section 5.4. Select distinct elements $\{\alpha_1, ..., \alpha_n\} \subseteq \mathbb{F}_q$, and form the generator matrix $\mathbf{G}_{RS}$ of equation (5.5) for an $[n, k_{min}]_q$ Reed-Solomon code. To each $m_i \in \mathcal{M}$, associate the polynomial $t_i(x) := \prod_{\{j \ : \ [\mathbf{A}_{\tilde{\mathcal{E}}}]_{i,j}=0\}} (x - \alpha_j)$. By our above discussion, we have $\deg(t_i(x)) \leq n - d_{min} = k_{min} - 1$ for each i. Now for each i, expressing $t_i(x)$ in the form $\sum_{i'=1}^{k_{min}} t_{i,i'} x^{i'-1}$, we define the coefficient vector $\mathbf{t}_i := \frac{1}{t_i(\alpha_{j(i)})} [t_{i,1}, ..., t_{i,k_{min}}]$, where we have normalized the polynomial's coefficients so that its evaluation at $\alpha_{j(i)}$ is 1. Then if we stack the vectors $\mathbf{t}_i$ to form the matrix $\mathbf{T}$ as in (5.8), and set

$$\mathbf{G} = \mathbf{T}\mathbf{G}_{RS} = \left[\frac{t_i(\alpha_j)}{t_i(\alpha_{j(i)})} \right], \tag{5.14}$$

we see that $\mathbf{G}$ has zeros precisely in the locations of the zeros of $\mathbf{A}_{\tilde{\mathcal{E}}}$, so it is the generator matrix for a linear code $\mathcal{C}$ fitting our constraints. It is in systematic form, since the columns in the indices corresponding to $\{c_{j(i)}\}_{i=1}^s$ form a permutation of the columns of the $s \times s$ identity matrix. This also

immediately shows that the code is full rank s. Finally, the minimum distance $d(\mathcal{C})$ of our code must be at least that of the $[n, k_{min}]_q$ Reed-Solomon code from which it is derived, thus $d(\mathcal{C}) \geq d_{min}$. But by Corollary 6, the reverse inequality also holds, and we see that we must have $d(\mathcal{C}) = d_{min}$. $\qquad\square$

5.6 Minimum Distance for Systematic Linear Codes

In this section, we will restrict our attention to the case where a code valid for $\mathcal{G}$ is linear, so that each $c_j \in \mathcal{V}$ is a linear function of the message symbols $m_i \in \mathcal{N}(c_j)$. We seek to answer the following: what is the greatest minimum distance attainable by a *systematic* linear code valid for $\mathcal{G}$?

Any systematic code must correspond to a matching $\tilde{\mathcal{E}} \subseteq \mathcal{E}$ which identifies each message symbol $m_i \in \mathcal{M}$ with a unique codeword symbol $c_{j(i)} \in \mathcal{V}$, where $j(i) \in \{1, \ldots, n\}$. Explicitly, $\tilde{\mathcal{E}}$ consists of s edges of the form $\{(m_i, c_{j(i)})\}$ for $i = 1, \ldots, s$ such that $c_{j(i_1)} \neq c_{j(i_2)}$ for $i_1 \neq i_2$. As before, $\tilde{\mathcal{V}}$ is the subset of vertices in $\mathcal{V}$ which are involved in the matching: $\tilde{\mathcal{V}} = \{c_{j(i)}\}_{i=1}^s$. Our code becomes systematic by setting $c_{j(i)} = m_i$ for $i = 1, \ldots, s$, and choosing each remaining codeword symbol $c_j \notin \tilde{\mathcal{V}}$ to be some linear function of its neighboring message symbols $m_i \in \mathcal{N}(c_j)$.

Definition 8. *For $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$, let $\tilde{\mathcal{E}} \subseteq \mathcal{E}$ be an $\mathcal{M}$-covering matching so that $\tilde{\mathcal{E}} = \{(m_i, c_{j(i)})\}_{i=1}^s$. Let $\tilde{\mathcal{V}} = \{c_{j(i)}\}_{i=1}^s$ be the vertices in $\mathcal{V}$ which are covered by $\tilde{\mathcal{E}}$. Define the matched adjacency matrix $\mathbf{A}_{\tilde{\mathcal{E}}} \in \{0, 1\}^{s \times n}$ so that $[\mathbf{A}_{\tilde{\mathcal{E}}}]_{i,j} = 1$ if and only if either $(m_i, c_j) \in \tilde{\mathcal{E}}$, or $c_j \notin \tilde{\mathcal{V}}$ and $(m_i, c_j) \in \mathcal{E}$. In other words, $\mathbf{A}_{\tilde{\mathcal{E}}}$ is the adjacency matrix of the bipartite graph formed by starting with $\mathcal{G}$ and deleting the edges $\mathcal{E}_{neg} = \{(m_i, c_j) \in \mathcal{E} \; : \; c_j \in \tilde{\mathcal{V}} \text{ and } j \neq j(i)\}$.*

Definition 9. *Let $\tilde{\mathcal{E}} \subseteq \mathcal{E}$ be a matching for the $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ which covers $\mathcal{M}$. Let $z_{\tilde{\mathcal{E}}}$ be the maximum number of zeros in any row of the corresponding matched adjacency matrix $\mathbf{A}_{\tilde{\mathcal{E}}}$, and define $k_{\tilde{\mathcal{E}}} := z_{\tilde{\mathcal{E}}} + 1$. Furthermore, define $k_{sys} = \min_{\tilde{\mathcal{E}}} k_{\tilde{\mathcal{E}}}$ where $\tilde{\mathcal{E}}$ ranges over all matchings for $\mathcal{G}$ which cover $\mathcal{M}$, and $d_{sys} = n - k_{sys} + 1$.*

Lemma 16. *For a given bipartite graph $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ which merits a matching that covers $\mathcal{M}$, we have*

$$s \leq k_{min} \leq k_{sys} \leq n \tag{5.15}$$

and

$$d_{sys} \leq d_{min}. \tag{5.16}$$

Proof. Let $\mathbf{A}$ be the adjacency matrix of $\mathcal{G}$.

For any subset $\mathcal{M}' \subseteq \mathcal{M}$ we have $d_{min} \leq |\mathcal{N}(\mathcal{M}')| - |\mathcal{M}'| + 1$, and likewise $k_{min} = n - d_{min} + 1 \geq$

$|\mathcal{M}'| + (n - |\mathcal{N}(\mathcal{M}')|)$. Taking $\mathcal{M}' = \mathcal{M}$ (and noting that in our framework, every $c_j \in \mathcal{V}$ is connected to at least one vertex in $\mathcal{M}$, hence $|\mathcal{N}(\mathcal{M})| = n$) we obtain $k_{min} \geq s$.

Now choose a set $\mathcal{M}'$ for which the above relation holds with equality, that is, $k_{min} = |\mathcal{M}'| + (n - |\mathcal{N}(\mathcal{M}')|)$. Since $\mathcal{N}(\mathcal{M}')$ is simply the union of the support sets of the rows of $\mathbf{A}$ corresponding to $\mathcal{M}'$, then each of these rows must have at least $n - |\mathcal{N}(\mathcal{M}')| = |\mathcal{N}(\mathcal{M}')^c|$ zeros. Furthermore, any matching $\tilde{\mathcal{E}}$ which covers $\mathcal{M}$ must identify the rows of $\mathcal{M}'$ with columns of $\mathcal{N}(\mathcal{M}')$. Thus, in the matched adjacency matrix $\mathbf{A}_{\tilde{\mathcal{E}}}$, the row corresponding to $j \in \mathcal{M}'$ must have $|\mathcal{M}'| - 1$ zeros in the columns of $\mathcal{N}(\mathcal{M})$ which are matched to $\mathcal{M}' \setminus \{j\}$, in addition to the $n - |\mathcal{N}(\mathcal{M}')|$ zeros in the columns corresponding to $\mathcal{N}(\mathcal{M}')^c$.

This gives us $k_{\tilde{\mathcal{E}}} \geq |\mathcal{M}'| + (n - |\mathcal{N}(\mathcal{M}')|)$ for each matching $\tilde{\mathcal{E}}$, hence $k_{sys} \geq k_{min}$. It follows directly that $d_{sys} \leq d_{min}$. Finally, it is clear from definition that for any $\mathcal{M}$-covering matching $\tilde{\mathcal{E}}$ we must have that $k_{\tilde{\mathcal{E}}}$ is less than the length of the adjacency matrix $\mathbf{A}$, which is n, hence $k_{sys} \leq n$. $\quad\square$

Corollary 7. *Let $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ be a bipartite graph which merits a systematic linear code. The largest minimum distance obtainable by a systematic linear code is d_{sys}. This distance can be achieved by a subcode of a Reed-Solomon code.*

Proof. Let $\mathcal{C}$ be a systematic linear code which is valid for $\mathcal{G}$. Then $\mathcal{C}$ must have a codeword containing at least $k_{sys} - 1$ zeros, i.e. a codeword of Hamming weight at most $n - k_{sys} + 1 = d_{sys}$. Since the code is linear, this Hamming weight is an upper bound for its minimum distance, so $d(\mathcal{C}) \leq d_{sys}$.

It remains to see that there are systematic linear codes which are valid for $\mathcal{G}$ and achieve a minimum distance of d_{sys}. Let $\tilde{\mathcal{E}}$ be an $\mathcal{M}$-covering matching for G such that $k_{\tilde{\mathcal{E}}} = k_{sys}$. Then for any $k \geq k_{sys}$ and prime power $q \geq n$, we claim that an $[n, k]_q$ Reed-Solomon code contains a systematic linear subcode that is valid for $\mathcal{G}$. Indeed, choose a set of n distinct elements $\{\alpha_i\}_{i=1}^n \subseteq \mathbb{F}_q$ as the defining set of our Reed-Solomon code. Then to form our subcode's generator matrix $\mathbf{G}$, note that (as mentioned before) $\mathbf{G}$ must have zero entries in the same positions as the zero entries of $\mathbf{A}_{\tilde{\mathcal{E}}}$, and indeterminate elements in the remaining positions. There are at most $k_{sys} - 1$ zeros in any row of $\mathbf{A}_{\tilde{\mathcal{E}}}$ (and at *least* $s - 1$ zeros in each row, since there must be s columns which have nonzero entries in exactly one row). For each row $i \in \{1, \ldots, s\}$ of $\mathbf{A}_{\tilde{\mathcal{E}}}$, let $\mathcal{I}_i \subseteq \{1, \ldots, n\}$ be the set of column indices j such that $[\mathbf{A}_{\tilde{\mathcal{E}}}]_{i,j} = 0$. Then form the polynomial $t_i(x) = \prod_{j \in \mathcal{I}_i}(x - \alpha_j)$ and normalize by $t_i(\alpha_{j(i)})$, which accordingly has degree at most k_{sys} (and at least $s - 1$). We now set the i^{th} row of $\mathbf{G}$ to be $(t_i(\alpha_1), \ldots, t_i(\alpha_n))$, and we see that by construction this row has zeros precisely at the indices $j \in \mathcal{I}_i$ as desired.

The rows of $\mathbf{G}$ generate a code with minimum distance at least that of the original Reed-Solomon code, which is $n - k + 1$. Furthermore, by setting $k = k_{sys}$ for our Reed-Solomon code, we see this new code $\mathcal{C}$ has minimum distance at least $n - k_{sys} + 1 = d_{sys}$. Since by our previous argument,

$d(\mathcal{C}) \leq d_{sys}$, the minimum distance of $\mathcal{C}$ must achieve d_{sys} with equality. $\qquad\qquad$ $\square$

5.7 Arbitrary MDS Codes

Up until now, we have been extensively employing Reed-Solomon codes in our constructions and proofs, though this has mainly been due to their familiarity, their ease of discussion, and the fact that they have a number of efficient decoding algorithms. It is worth mentioning, however, that we could have instead used any MDS codes with the same dimensions to satisfy our constraints.

Let $\mathcal{C}_{MDS}$ be an MDS code in $\mathbb{F}_q^n$ of dimension k and generator matrix $\mathbf{G}_{MDS} \in \mathbb{F}_q^{k \times n}$. For any subset $\mathcal{I} \subseteq [n]$ such that $|\mathcal{I}| \leq k - 1$, there is a nonzero codeword $\mathbf{c} = [c_1, ..., c_n] \in \mathcal{C}$ such that $c_i = 0$, $\forall i \in \mathcal{I}$.

To see this, let $\mathbf{g}_i$ be the i^{th} column of $\mathbf{G}_{MDS}$, and define the submatrix $\mathbf{G}_{\mathcal{I}} := [\mathbf{g}_i]_{i \in \mathcal{I}}$. Since $\mathbf{G}_{\mathcal{I}}$ is a tall matrix, there is a nonzero vector $\mathbf{h} \in \mathbb{F}_q^{1 \times k}$ such that $\mathbf{h}\mathbf{G}_{\mathcal{I}} = \mathbf{0} \in \mathbb{F}_q^{|\mathcal{I}|}$. Therefore the desired codeword is $\mathbf{c} = \mathbf{h}\mathbf{G}_{MDS}$.

Now suppose we have a bipartite graph $\mathcal{G} = (\mathcal{M}, \mathcal{V}, \mathcal{E})$ of constraints for which we wish to produce a valid code $\mathcal{C}$. For any integer d such that $d \leq |\mathcal{N}(m_i)|$, $\forall i$, choose $\mathcal{C}_{MDS} \subseteq \mathbb{F}_q^n$ to have dimension $k = n - d + 1$. If $\mathbf{A}$ is the adjacency matrix of $\mathcal{G}$, then each row of $\mathbf{A}$ has at most $n - d = k - 1$ zeros. So for each i, we can find some $\mathbf{h}_i \in \mathbb{F}_q^k$ such that the codeword $\mathbf{c}_i := \mathbf{h}_i \mathbf{G}_{MDS}$ has zeros in the same entries as the zeros of $\mathbf{A}$.

Thus the matrix

$$\mathbf{G} := \begin{bmatrix} \mathbf{c}_1 \\ \vdots \\ \mathbf{c}_s \end{bmatrix}$$

is the generator matrix for a linear code $\mathcal{C}$ which is valid for $\mathcal{G}$. Note that since the minimum distance of $\mathcal{C}$ must be at least that of $\mathcal{C}_{MDS}$, which is d. On the other hand, $d(\mathcal{C})$ is equal to the lowest Hamming weight of any of its codewords, among which are the rows $\mathbf{c}_i$ of $\mathbf{G}$. Row $\mathbf{c}_i$ can have Hamming weight no greater than $|\mathcal{N}(m_i)|$, so we have

$$d \leq d(\mathcal{C}) \leq \min_{i \in [s]} |\mathcal{N}(m_i)|. \tag{5.17}$$

This, of course, is in addition to the other bounds on distance we had from Theorem 26. (In fact, the upper bound in (5.17) is actually one of these, since it can be rewritten as $|\mathcal{N}(m_i)| = |\mathcal{N}(m_i)| - |\{m_i\}| + 1$). Thus, a priori it seems that we can freely control our code's distance by simply choosing $d = \min_{\mathcal{M}' \subseteq \mathcal{M}} |\mathcal{N}(\mathcal{M}')| - |\mathcal{M}'| + 1$. But it is important to remember that the dimension of $\mathcal{C}$ (equal to the rank of $\mathbf{G}$) depends on our particular choice of the rows $\mathbf{c}_i = \mathbf{h}_i \mathbf{G}_{MDS}$. Thus, if we begin with

an MDS code of too high a minimum distance d, it may result in $\mathcal{C}$ having dimension smaller than $s = |\mathcal{M}|$.

One more thing to point out is that in the case where we would like $\mathcal{C}$ to be a *systematic* code, so that s of the columns of its generator matrix $\mathbf{G}$ are a permutation of the columns of the $s \times s$ identity matrix, we have to take slightly more care in choosing the $\mathbf{h}_i$. Indeed, if m_i is matched to $c_{j(i)}$ as described in the previous section, then we must make sure that $\mathbf{h}_i \mathbf{g}_{j(i)} \neq 0$. This is always possible, because any k columns of $\mathbf{G}_{MDS}$ must be full rank by Theorem 25, and our only requirement on $\mathbf{h}_i$ had been that it lie in the null space of a submatrix $\mathbf{G}_{\mathcal{I}}$ of at most $k-1$ columns. In this case, we can be sure $\mathcal{C}$ will have dimension s since $\mathbf{G}$ will have an identity submatrix, and hence will have rank s.

5.8 Example

In this section, we give an explicit example of how to use a Reed-Solomon code to construct a systematic code which is valid for the constraints induced by the bipartite graph $\mathcal{G}$ of Figure 5.1. Corollary 6 bounds the distance of any constrained code $\mathcal{C}$ as $d(\mathcal{C}) \leq 5$, but if $\mathcal{C}$ is required to be *systematic*, Corollary 7 bounds the distance as $(\mathcal{C}) \leq 4$. This is a simple example where the systematic distance bound from Corollary 7 actually proves tighter than that given by our original (general) bound in Corollary 6.

Corollary 7 requires an $\mathcal{M}$-covering matching $\tilde{\mathcal{E}}$, and we see that if we match m_1, m_2, and m_3 to c_1, c_2, and c_3, respectively, and remove all other edges incident to these three c_i from our bipartite graph (i.e. remove edges (m_2, c_1) and (m_2, c_3)), then the matched adjacency matrix becomes

$$\mathbf{A}_{\tilde{\mathcal{E}}} = \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 1 & 1 \\ \mathbf{0} & 1 & \mathbf{0} & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}. \tag{5.18}$$

Here the bold entries correspond to the edges that have been removed from $\mathcal{G}$, or equivalently '1' entries which have been removed from the original adjacency matrix $\mathbf{A}$ from (5.2). By speculation, this matching minimizees the maximum number of zeros in any row of the matched adjacency matrix, and hence it is an optimal matching and yields the bound of Corollary 7.

We can construct a linear code that is vlid for $\mathbf{A}_{\tilde{\mathcal{E}}}$ from a $[7, 4]_7$ Reed-Solomon code. If we fix a primitive element α of $\mathbb{F}_7$ and select the defining set $\{\alpha_1, ..., \alpha_7\}$ to be $\{0, 1, \alpha, ..., \alpha^5\}$ when we form the generator matrix $\mathbf{G}_{RS}$, then using the method described in Section 5.4, we identify the message

symbols m_1, m_2, and m_3, respectively, with the polynomials

$$t_1(x) = \alpha^5(x-1)(x-\alpha), \tag{5.19}$$

$$t_2(x) = \alpha^4 x(x-\alpha)(x-\alpha^2), \tag{5.20}$$

$$t_3(x) = \alpha^3 x(x-1). \tag{5.21}$$

Here, the leading constants of the polynomials are normalizing factors to ensure that our code will be in systematic form. Again following the procedure from Section sec:ReedSolomon, we form the vector $\mathbf{t}_i$ of the coefficients of $t_i(x)$ for $i = 1, 2, 3$, and stack these vectors to form the matrix $\mathbf{T}$. Our subcode's generator matrix will then take the form

$$\mathbf{G} = \mathbf{TG}_{RS} = \begin{bmatrix} 1 & 0 & 0 & \alpha^2 & \alpha^5 & 1 & \alpha^5 \\ 0 & 1 & 0 & 0 & 1 & \alpha^4 & 1 \\ 0 & 0 & 1 & \alpha^5 & \alpha^5 & \alpha^2 & 1 \end{bmatrix}, \tag{5.22}$$

where we can indeed see that the submatrix corresponding to the first three columns is the identity matrix, and hence we have produced a valid systematic code for $\mathcal{G}$.

Chapter 6

Conclusions and Future Work

In this thesis, we have delved into several topics in coding theory in its various forms. We have presented algebraic constructions of Ingleton-violating entropy vectors, low-coherence frames, and constrained codes, and have gained much structural insight into these problems. More importantly, our work has paved the way for some exciting new directions of study in these areas. In this chapter, we suggest several future continuations of our work.

6.1 Characterizing the Entropy Region

The constructions of Ingleton-violating group-characterizable entropy vectors that we presented in Chapter 2 represent a broad class of points in a very interesting part of the entropy region, and in some ways they suggest techniques that can be used to determine points in many other subregions. But there remains a great deal of work to be done to completely characterize the entropy region, or even just the space of entropy vectors which violate the Ingleton inequality. This task requires a much more general approach to studying group characterizable entropy vectors, and would likely demand far more complicated groups than those which can be expressed as small matrices over finite fields, or even those which merit obvious group actions. A more reachable goal would be to more extensively characterize the portion of the entropy region which is achieved by entropy vectors arising from familiar classes of groups. An ultimate goal is to use these groups to construct high-performing network codes, so it makes sense to restrict our attention to groups that easy to work with and facilitate code design. As a first step, it would be interesting to classify networks for which we can construct codes from the projective and general linear groups $PGL(2,q)$ and $GL(2,q)$ discussed in Chapter 2, particularly if they can outperform linear codes on these networks.

6.2 Frame Design

The idea of reducing the number of distinct inner product magnitudes arising in a tight frame is in some sense a natural generalization of a Grassmanian frame. The construction we presented in Chapter 3 achieves this goal, though in a very specific manner. It would be interesting to classify other frames that have this property, or if there are other ways to generalize the notion of a Grassmanian frame. The generalizations of our construction that we gave in Chapter 4 give us much more flexibility in controlling the dimensions and alphabets of our frames, but it remains a challenge to identify groups that work well in this generalized framework. While it may be simple enough to explicitly compute the inner products of these frames directly from the characters of different groups, it is sometimes less trivial to explicitly write out the frame elements themselves, as one might have noticed in the frames constructed from $SL_2(q)$ in Section 4.5. It is conceivable that we can describe how the coherence of these frames will behave in terms of the interaction between various subgroups or generators of our group, which could allow us to produce new groups more easily.

In light of the fact that our frames do contain a great deal of structure, it would be pleasing to identify applications in which our frames outperform other frames with low coherence. In many compressive sensing simulations, or experiments in which we attempt to decode a noise-corrupted spherical code message, it is difficult to perform significantly better than randomly generated frames (though even matching their performance with a deterministically-designed frame is useful). Furthermore, while our frames provably satisfy the Strong Coherence Property in certain regimes, we can see from Figures 4.3(a) and 4.3(b) that this often requires both the number of rows m and columns n of our frame matrices to be very large. In the case where $m = n^\gamma$, for $\gamma < 1$, then for large enough n the bounds on the coherence of our group-based harmonic frames do not outcompete those for random harmonic frames. Thus it would be very useful to find a deterministic set of frames which is guaranteed to have lower coherence in high dimensions, or perhaps in the regime where $m = O(\log n)$. This might involve tightening the bounds we provided in Chapters 3 and 4, identifying groups which work better in our framework, or finding a new approach to constructing frames.

6.3 Constrained Coding

Chapter 5 answered several questions about the connection between the constraining bipartite graph and the minimal distance of a constrained code. A major open question that remains is whether the bound on minimum distance in Theorem 26 always achievable with a subcode of a Reed-Solomon code. Our work answers this question when we demand a systematic subcode, in which case we can only achieve the altered bound of Corollary 7. But ultimately we would like to be able to provide a

Reed-Solomon subcode for any given set of constraints that achieves the optimal minimal distance so that we can take advantage of existing fast Reed-Solomon decoders.

Another open question concerns how to compute the quantity d_{sys} from Corollary 7. This is equivalent to computing the quantity k_{sys} from Definition 9, which boils down to finding the matching $\tilde{\mathcal{E}}$ which minimizes the maximum number of zeros in any row of the matched adjacency matrix $\mathbf{A}_{\tilde{\mathcal{E}}}$ from Definition 8. A priori, it is unclear whether this problem has a computationally fast solution. We can actually relax this problem, however, to obtain the following linear program:

$$\text{minimize}_{\{\mathbf{w}^{(j)}\}_{j=1}^{s} \subset \mathbb{R}^n} \left(\max_i \left[(n-s)\mathbf{1}_{s \times 1} - \mathbf{A} \left(\mathbf{1}_{s \times 1} - \sum_j \mathbf{w}^{(j)} \right) \right]_i \right) \tag{6.1}$$

$$\text{subject to: } 0 \leq [\mathbf{w}^{(j)}]_i \leq 1, \qquad\qquad j = 1, ..., s$$

$$\sum_i [\mathbf{w}^{(j)}]_i = 1, \qquad\qquad j = 1, ..., s$$

$$[\mathbf{A}\mathbf{w}^{(j)}]_j = 1, \qquad\qquad j = 1, ..., s$$

$$0 \leq \left[\sum_j \mathbf{w}^{(j)} \right]_i \leq 1, \qquad\qquad i = 1, ..., n$$

where $\mathbf{1}_{s \times 1}$ is a length-s vector of all 1s, and for a vector $\mathbf{v}$, the notation $[\mathbf{v}]_i$ denotes the i^{th} entry. The idea of the above linear program is that we would like $\mathbf{w}^{(j)}$ to be a vector with all 0s except for a single 1 in the entry corresponding to the index of the code symbol c_i which is matched to the message symbol m_j. In this case, $\mathbf{A} \left(\mathbf{1}_{s \times 1} - \sum_j \mathbf{w}^{(j)} \right)$ is the total number of zeros in each row of the matched adjacency matrix, and the objective function gives us d_{sys}. It is not too difficult to see that if the solution to this optimization problem yields a set of vectors $\mathbf{w}^{(j)}$ with all 0,1 entries, then this certifies that the solution is indeed a matching and that the minimized value of the objective function is the true value of d_{sys}. Otherwise, this linear program only gives us a lower bound on d_{sys}. It remains to characterize the effectiveness of this method of searching for the largest achievable minimum distance for systematic linear codes.

Appendices

Appendix A

Chapter 3 Proofs

A.1 The Fourier Pairing of (3.22) and (3.23) for Cyclic Groups of Prime Order

We will now begin to develop the tools needed to prove Theorems 7, 8, 9 and 10. We will explicitly prove Theorems 7 and 8 and defer the proofs of Theorems 9 and 10 to Appendix B. Let us return to representations of the cyclic group $G = \mathbb{Z}/n\mathbb{Z}$, where $A = \{a_1, ..., a_m\}$ is a subset of G (not necessarily a group), $\mathbf{U} = \mathrm{diag}(\omega^{a_1}, \omega^{a_2}, ..., \omega^{a_m})$, where $\omega = e^{2\pi i/n}$ and the powers ω^{a_i} are distinct, and $\mathcal{U} = \{\mathbf{U}, \mathbf{U}^2, ..., \mathbf{U}^{n-1}, \mathbf{U}^n = \mathbf{I_m}\}$. As before, taking $\mathbf{v} = \frac{1}{\sqrt{m}}\mathbf{1_m}$ the normalized vector of all ones, $\mathbf{U}^\ell \mathbf{v} = \begin{bmatrix} \omega^{a_1 \ell} & \omega^{a_2 \ell} & ... & \omega^{a_m \ell} \end{bmatrix}^T$. Then if we index the columns as $\ell = 0, 1, ..., n-1$, we have $\mathbf{M}$ as in (3.18). The inner product associated to the element $\mathbf{U}^\ell$ takes the form

$$c_\ell := \frac{\mathbf{v}^* \mathbf{U}^\ell \mathbf{v}}{||\mathbf{v}||_2^2} = \frac{1}{m} \sum_{a \in A} \omega^{\ell a}. \tag{A.1}$$

We define $\alpha_\ell := |c_\ell|^2$ to be the squared norm of the ℓ^{th} inner product. If for any $t \in \mathbb{Z}/n\mathbb{Z}$ we define the set $\mathcal{A}_t := \{(k_i, k_j) \in A \times A \mid a_i - a_j \equiv t \mod n\}$ with size $N_t := |A_t|$, then we have the Fourier pairing given by Equations (3.22) and (3.23).

Now consider the framework of Section 3.5 where n is a prime, m is a divisor of $n - 1$, and A is the unique cyclic subgroup of $H = (\mathbb{Z}/n\mathbb{Z})^\times$ of size m. If $\kappa = \frac{n-1}{m}$, then A consists of the nonzero κ^{th} powers in $\mathbb{Z}/n\mathbb{Z}$. Let x be a multiplicative generator of the cyclic group $(\mathbb{Z}/n\mathbb{Z})^\times$. Then the distinct cosets of A in $(\mathbb{Z}/n\mathbb{Z})^\times$ are $\{A, xA, x^2 A, ..., x^{r-1} A\}$. If $\ell \in x^d A$, then we see that $c_\ell = c_{x^d}$ and hence $\alpha_\ell = \alpha_{x^d}$. Likewise, if $t \in x^d A$, it is not too difficult to see that we have a bijection

$$\mathcal{A}_t \to \mathcal{A}_{x^d} : \quad (a_i, a_j) \mapsto (x^d t^{-1} a_i, x^d t^{-1} a_j).$$

It follows that

$$N_t = N_{x^d} \text{ if } t \in x^d A. \tag{A.2}$$

It is straightforward to see from their definitions that $c_0 = \alpha_0 = 1$ and $N_0 = m$. With this in mind, we may write the condensed forms of (3.22) and (3.23):

$$\alpha_\ell = \frac{1}{m^2} \left(a_0 + \sum_{d=0}^{\kappa-1} N_{x^d} \sum_{a \in A} \omega^{x^d \ell a} \right),$$

$$N_t = \frac{m^2}{n} \left(\alpha_0 + \sum_{d=0}^{\kappa-1} \alpha_{x^d} \sum_{a \in A} \omega^{-x^d ta} \right).$$

In particular,

$$\alpha_{x^{d'}} = \frac{1}{m^2} \left(N_0 + \sum_{d=0}^{\kappa-1} N_{x^d} \sum_{a \in A} \omega^{x^{d+d'}a} \right) \tag{A.3}$$

$$= \frac{1}{m} \left(1 + \sum_{d=0}^{\kappa-1} N_{x^d} c_{x^{d+d'}} \right), \tag{A.4}$$

$$N_{x^{d'}} = \frac{m^2}{n} \left(\alpha_0 + \sum_{d=0}^{\kappa-1} \alpha_{x^d} \sum_{a \in A} \omega^{-x^{d+d'}a} \right) \tag{A.5}$$

$$= \frac{m^2}{n} \left(1 + m \sum_{d=0}^{\kappa-1} \alpha_{x^d} c^*_{x^{d+d'}} \right). \tag{A.6}$$

On one final note, since the roots of unity sum to 0:

$$1 + mc_1 + mc_x + mc_{x^2} + ... + mc_{x^{\kappa-1}} = 0. \tag{A.7}$$

A.2 $\kappa = 2$, and Proof of Theorem 7

As before, take n to be a prime, m a divisor of $n-1$, and $A = \{a_1, ..., a_m\}$ the unique multiplicative subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m. Let us examine the case where $\kappa := \frac{n-1}{m} = 2$. Fix a multiplicative generator x for $(\mathbb{Z}/n\mathbb{Z})^\times$. In this case, A has two distinct cosets: A and xA. Our frame will correspondingly have two distinct inner product values: $c_1 = \frac{1}{m} \sum_{a \in A} \omega^a$ and $c_x = \frac{1}{m} \sum_{a \in A} \omega^{xa}$. There are two equations of the form (A.4),

$$\alpha_1 = \frac{1}{m} \left(1 + N_1 c_1 + N_x c_x \right) , \quad \alpha_x = \frac{1}{m} \left(1 + N_1 c_x + N_x c_1 \right).$$

From (B.10), substituting $c_x = -\left(\frac{1}{m} + c_1\right)$ gives us

$$\alpha_1 = \frac{1}{m}\left(1 - \frac{1}{m}N_x + (N_1 - N_x)c_1\right), \tag{A.8}$$

$$\alpha_x = \frac{1}{m}\left(1 - \frac{1}{m}N_1 + (N_x - N_1)c_1\right). \tag{A.9}$$

From (A.8) and (A.9), we can see that since α_1, α_x, N_1, and N_x are real, then c_1 must be real as well (and thus so is c_x). This allows us to write

$$\alpha_1 = c_1^2, \qquad \alpha_x = c_x^2 = \left(\frac{1}{m} + c_1\right)^2. \tag{A.10}$$

Lemma 17. *Let n be a prime, and κ and m satisfy $\kappa = \frac{n-1}{m} = 2$. Let A be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m. As before, let N_t be the number of pairs $(a_1, a_2) \in A \times A$ such that $a_1 - a_2 = t$. Let x be the multiplicative generator of $(\mathbb{Z}/n\mathbb{Z})^\times$. Then*

- *If $n - 1$ is divisible by 4, $N_1 = \frac{1}{2}(m-2)$ and $N_x = \frac{1}{2}m$.*

- *Otherwise, $N_1 = N_x = \frac{1}{2}(m-1)$, $(-1 \notin A)$.*

Proof. Let us first count the number of pairs (a_1, a_2) such that $a_1 - a_2 \in A$, which will give us $\sum_{a \in A} N_a$. From (A.2), this is precisely equal to mN_1. Since A is the group of nonzero squares in $\mathbb{Z}/n\mathbb{Z}$, we can write $a_1 = h_1^2$ and $a_2 = h_2^2$ for some choice of $h_1, h_2 \in (\mathbb{Z}/p\mathbb{Z})^\times$. If we let $x_1 = h_1 - h_2$ and $x_2 = h_1 + h_2$, then $a_1 - a_2 = (h_1 - h_2)(h_1 + h_2) = x_1 \cdot x_2$.

Equivalently, we may write

$$\begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} h_1 \\ h_2 \end{bmatrix} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}.$$

We see that for any choice of the pair (x_1, x_2), there is a unique pair (h_1, h_2) which maps to it. Since we need to consider only pairs where h_1 and h_2 are nonzero, we must eliminate the cases where $x_1 = x_2$ (corresponding to when $h_2 = 0$) and $x_1 = -x_2$ (corresponding to when $h_1 = 0$).

In order to have $x_1 \cdot x_2 \in A$, we must either have x_1 and x_2 both in A or both in xA. If $-1 \in A$, a quick counting argument shows that there are $2m(m-2)$ valid choices for (x_1, x_2) which satisfy $x_1 \cdot x_2 \in A$, each yielding a pair (h_1, h_2) with h_1 and h_2 nonzero. But we are concerned only with their squares h_1^2 and h_2^2, so we can group these ordered pairs into sets of four, $\{(\pm h_1, \pm h_2)\}$, and the number of distinct pairs (h_1^2, h_2^2) with h_1^2 and h_2^2 nonzero and $h_1^2 - h_2^2 \in A$ is thus

$$mN_1 = \frac{1}{4}(2m(m-2)) = \frac{m}{2}(m-2), \text{ if } -1 \in A.$$

Likewise, $x_1 \cdot x_2 \in xA$ precisely when x_1 and x_2 are in opposite cosets of A. If this is true, and $-1 \in A$, then we cannot have $x_1 = x_2$ or $x_1 = -x_2$, since this would imply that x_1 and x_2 are in the same coset. Thus, any pair (x_1, x_2) in either $A \times xA$ or $xA \times A$ will yield $x_1 \cdot x_2 \in xA$, so there are $2m^2$ possible pairs, each yielding a pair (h_1, h_2). Again, we must divide by $4m$ to get the number of feasible pairs (h_1^2, h_2^2) such that $h_1^2 - h_2^2 = x$, and we find that

$$N_x = \frac{1}{2}m, \ (-1 \in A).$$

If $-1 \notin A$, then the calculations for N_1 and N_x change slightly: Now the condition $x_1 = -x_2$ implies that x_1 and x_2 are in opposite cosets of A. Thus, we have one extra case to consider when calculating N_1, and one less case when calculating N_x, so we find

$$N_1 = N_x = \frac{1}{2}(m - 1), \ (-1 \notin A).$$

Note that $-1 \in A$, or rather -1 is a square modulo n, precisely when $(\mathbb{Z}/n\mathbb{Z})^\times$ contains a fourth root of unity, and since $(\mathbb{Z}/n\mathbb{Z})^\times$ is a cyclic multiplicative group of size $n - 1$, this occurs precisely when $n - 1$ is divisible by 4. $\qquad\square$

We now have all the ingredients to prove Theorem 7, which we restate here for convenience:

Theorem 7: *Let n be a prime, m a divisor of $n-1$, and $\omega = e^{\frac{2\pi i}{n}}$. Let $A = \{a_1, ..., a_m\}$ be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m, and set $\mathbf{U} = \mathrm{diag}(\omega^{a_1}, ..., \omega^{a_m}) \in \mathbb{C}^{m \times m}$, $\mathbf{v} = \frac{1}{\sqrt{m}}[1, ..., 1]^T \in \mathbb{C}^{m \times 1}$, and $\mathbf{M} = [\mathbf{v}, \mathbf{U}\mathbf{v}, ..., \mathbf{U}^{n-1}\mathbf{v}]$.*

If $\kappa := \frac{n-1}{m} = 2$, there are two distinct inner product values between the columns of $\mathbf{M}$, both of which are real. If $n - 1$ is divisible by 4, these inner products are $\frac{-1 \pm \sqrt{1+2m}}{2m}$. In this case, $\mathbf{M}$ has coherence $\sqrt{\frac{n-m-\frac{1}{2}}{m(n-1)} + \frac{1}{2m}}$.

If $n - 1$ is not divisible by 4, then the columns of $\mathbf{M}$ form an equiangular frame. The two inner products are $\pm\sqrt{\frac{1}{m}\left(\frac{1}{2} + \frac{1}{2m}\right)}$, and the coherence is $\sqrt{\frac{n-m}{m(n-1)}}$.

Proof. (Theorem 7) From (A.8), (A.9), (A.10), and Lemma 17, we have that if $n - 1$ is divisible by 4,

$$c_1^2 = \frac{1}{m}\left(\frac{1}{2} - c_1\right),$$

and making the substitution $c_1 = -\left(\frac{1}{m} - c_x\right)$ from (B.10) yields the same quadratic equation in c_x. Solving this reveals that c_1 and c_x will take on the values $\frac{-1 \pm \sqrt{1+2m}}{2m}$, and the solution with the larger norm is $\frac{-1 - \sqrt{1+2m}}{2m}$, which indicates that the coherence is

$$\mu = \left|\frac{-1 - \sqrt{1+2m}}{2m}\right| = \sqrt{\frac{n-m-\frac{1}{2}}{m(n-1)} + \frac{1}{2m}} \quad (n \equiv 1 \mod 4).$$

On the other hand, if $n-1$ is not divisible by 4, then from Lemma 17 equations (A.9) and (A.9) become

$$c_1^2 = c_x^2 = \frac{1}{m}\left(\frac{1}{2} + \frac{1}{2m}\right),$$

so this gives us coherence

$$\mu = \sqrt{\frac{1}{m}\left(\frac{1}{2} + \frac{1}{2m}\right)} = \sqrt{\frac{n-m}{m(n-1)}} \quad (n \not\equiv 1 \mod 4).$$

□

A.3 $\kappa = 3$, and Proof of Theorem 8

Take n to be a prime, m a divisor of $n-1$, and $A = \{a_1, ..., a_m\}$ the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m. We now consider the case where $\kappa = \frac{n-1}{m} = 3$, so that if x is a generator of $(\mathbb{Z}/n\mathbb{Z})^\times$, then A is cyclically generated by x^3, and consists of the cubes of all the nonzero integers modulo n. In this case our distinct inner products will be c_1, c_x, and c_{x^2}, with corresponding squared norms α_1, α_x, and α_{x^2}. Our goal in this section will be to prove Theorem 8.

We first make the following remark:

Lemma 18. *Let n be a prime, $\omega = e^{\frac{2\pi i}{n}}$, and r and m satisfy $\kappa = \frac{n-1}{m} = 3$. If we take A to be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m, then the inner product values $c_\ell = \frac{1}{m}\sum_{a\in A}\omega^{\ell a}$ are all real.*

Proof. A is the set of cubes in $(\mathbb{Z}/n\mathbb{Z})^\times$, and since -1 is its own cube it will lie in A. Multiplication by -1 will therefore permute the elements of A, so we have

$$c_\ell^* = \left(\frac{1}{m}\sum_{a\in A}\omega^{\ell a}\right)^* = \frac{1}{m}\sum_{a\in A}\omega^{-\ell a} = c_\ell. \tag{A.11}$$

□

We begin by making the following definition:

Definition 10. *For any two cosets $t_1 A$ and $t_2 A$, we define the translation degree from $t_1 A$ to $t_2 A$, to be the quantity*

$$N_{t_1 A, t_2 A} := |(1 + t_1 A) \cap t_2 A| = \#\{\alpha \in t_1 A \mid 1 + \alpha \in t_2 A\}.$$

Similarly, for any coset tA, define the translation degree from tA to 0 to be the quantity

$$N_{tK,0} := |(1 + tA) \cap \{0\}| = \begin{cases} 1 & if\ -1 \in tA, \\ 0 & otherwise. \end{cases}$$

We can express our previously defined values N_t in terms of the translation degrees as follows:

Lemma 19. *Let n be a prime, and m and κ satisfy $\kappa = \frac{n-1}{m} = 3$. Let A be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m. Define $N_t = \#\{(a_1, a_2) \in A \times A \mid a_1 - a_2 \equiv t \mod n\}$ Then $N_t = N_{K,tK}$.*

Proof. For every pair $(a_1, a_2) \in A \times A$ we have that $a_1 - a_2 \in tA$ if and only if $1 - a_2 a_1^{-1} \in tA$. There are mN_t such pairs in total (N_t pairs for every element in tA). Note that $-a_2 a_1^{-1} \in A$, since $-1 \in A$. If we select any of the m candidates for $a_1 \in A$, then there are $N_{A,tA}$ choices for a_2 that will satisfy this requirement. Thus, we have $mN_t = mN_{A,tA}$, and the result follows. $\square$

Some other facts about translation degrees:

Lemma 20. *Let n be a prime, m a divisor of $n - 1$ such that $\frac{n-1}{m} = 3$, and A the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m. Then $N_{t_1 A, t_2 A} = N_{t_2 A, t_1 A}$ for all $t_1, t_2 \in \mathbb{Z}/n\mathbb{Z}$.*

Proof. Suppose $b_1 \in t_1 A$ such that $1 + b_1 = b_2 \in t_2 A$. Then $1 - b_2 = -b_1$, with $-b_2 \in t_2 A$ and $-b_1 \in t_1 A$ (since -1 is a cube and is thus in A). In fact, we see that we have a bijection between the sets $\{(b_1, b_2) \in t_1 A \times t_2 A \mid 1 + b_1 = b_2\}$ and $\{(b_1', b_2') \in t_2 A \times t_1 A \mid 1 + b_1' = b_2'\}$ which sends $(b_1, b_2) \mapsto (b_1', b_2') := (-b_2, -b_1)$. This gives us $N_{t_1 A, t_2 A} = N_{t_2 A, t_1 A}$. $\square$

Lemma 21. *Let n be a prime, m a divisor of $n - 1$ such that $\kappa = \frac{n-1}{m}$, and A the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m. If x is the multiplicative generator of $(\mathbb{Z}/n\mathbb{Z})^\times$, then $N_{x^i A, x^j A} = N_{x^{\kappa-i} A, x^{\kappa-i+j} A}$.*

Proof. Let $a \in A$ such that $1 + x^i a = x^j b$, with $b \in A$. Then multiplying both sides of this equation by $x^{\kappa-i}$, we get $x^{\kappa-i} + x^\kappa a = x^{\kappa-i+j} b$. Note that $x^\kappa a \in A$. Now, multiplying both sides of this equation by $(x^\kappa a)^{-1} \in A$, we obtain $1 + x^{\kappa-i}(x^\kappa a)^{-1} = x^{\kappa-i+j} b (x^\kappa a)^{-1}$, where $x^{\kappa-i}(x^\kappa a)^{-1} \in x^{\kappa-i} A$ and $x^{\kappa-i+j} b (x^\kappa a)^{-1} \in x^{\kappa-i+j} A$. We see that we in fact have a bijection between the sets $\{(x^i a, x^j b) \in x^i A \times x^j A \mid 1 + x^i a = x^j b\}$ and $\{(x^{\kappa-i} c, x^{\kappa-i+j} d) \in x^{\kappa-i} A \times x^{\kappa-i+j} A \mid 1 + x^{\kappa-i} c = x^{\kappa-i+j} d\}$ which sends $(x^i a, x^j b) \mapsto (x^{\kappa-i}(x^\kappa a)^{-1}, x^{\kappa-i+j} b (x^\kappa a)^{-1})$. $\square$

Lemma 22. *Let n be a prime, m a divisor of $n - 1$ such that $\kappa = \frac{n-1}{m}$, and A the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m. Set $H = (\mathbb{Z}/n\mathbb{Z})^\times$, with multiplicative generator x. For any coset $t_0 A$, we have $N_{t_0 A, 0} + \sum_{i=1}^\kappa N_{t_0 A, x^i A} = |t_0 A|$.*

Proof. This simply follows from the observation that any element of $t_0 A$, when translated by 1, must be sent to either 0 or exactly one of the cosets $x^i A \in H/A$. $\square$

Lemma 23. *Let n be a prime, and m a divisor of $n-1$ such that $\kappa := \frac{n-1}{m} = 3$. Take A to be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m, and x a multiplicative generator for $H := (\mathbb{Z}/n\mathbb{Z})^\times$. Then*

$$N_{xA,x^2A} - N_{A,A} = 1. \tag{A.12}$$

Proof. We prove this by counting the size of the set

$$\mathcal{A}_K := \{(a_1, a_2) \in A \times A \mid a_1 - a_2 \in A\}$$

in two ways. First, using Equation (A.2), we can simply count the elements in this set as

$$|\mathcal{A}_A| = \sum_{a \in A} N_a = m N_1. \tag{A.13}$$

Alternatively, we note that when $\kappa = 3$, the difference between any two elements in A takes the form

$$b_1^3 - b_2^3 = (b_1 - b_2)(b_1 - \zeta b_2)(b_1 - \zeta^2 b_2),$$

where ζ is a primitive third root of unity, and b_1 and b_2 are nonzero elements of $\mathbb{Z}/n\mathbb{Z}$. Let us define

$$x_1 := b_1 - b_2, \ x_2 := b_1 - \zeta b_2, \ x_3 := b_1 - \zeta^2 b_2. \tag{A.14}$$

We can express this using matrices as

$$\begin{bmatrix} 1 & -1 \\ 1 & -\zeta \\ 1 & -\zeta^2 \end{bmatrix} \cdot \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}.$$

In this form we can see that b_1 and b_2, and x_3 are uniquely determined by x_1 and x_2. In particular,

$$x_3 = -\zeta(x_1 + \zeta x_2). \tag{A.15}$$

Now, if $b_1^3 - b_2^3 \in A$, then Table A.1 lists the possibilities for the cosets of A to which x_1, x_2, and x_3 must belong (up to a permutation of the cosets).

Table A.1: $b_1^3 - b_2^3 \in A$

x_1	x_2	x_3	Multiplicity
A	A	A	1
xA	xA	xA	1
x^2A	x^2A	x^2A	1
A	xA	x^2A	6

The last case is representative of six possible cases which we obtain by permuting the order of the cosets (thus it has "multiplicity 6"). In short, we must have x_1, x_2, and x_3 all in the same coset, or all in different cosets of A in order to have $b_1^3 - b_2^3 \in A$. Let us attempt to count the quantity

$$\#\{(x_1, x_2) \in A \times A \mid x_3 = -\zeta(x_1 + \zeta x_2) \in A, \ b_1 \neq 0, b_2 \neq 0\}.$$

Since x generates $(\mathbb{Z}/n\mathbb{Z})^\times \cong \mathbb{F}_n^\times$, the multiplicative group of the finite field with n elements, and κ divides $n - 1$, the order of this group, then any κ^{th} root of unity will be contained in $\mathbb{F}_n^\times$, so ζ will lie in one of the cosets of A.

We will first consider the case where $\zeta \in A$. Since $\kappa = 3$, $-1 \in A$, so $-\zeta \in A$. Thus, the condition that $-\zeta(x_1 + \zeta x_2) \in A$ is equivalent to the condition that $x_1 + \zeta x_2 \in A \iff 1 + \zeta x_2 x_1^{-1} \in A$. If we fix x_1 to be any one of the m elements of A, we have exactly $N_{A,A}$ choices for x_2 which satisfy this condition (for each $a \in A$ such that $1 + a \in A$, simply set $x_2 = ax_1\zeta^{-1}$). This gives us a total of $mN_{A,A}$ ordered pairs $(x_1, x_2) \in A \times A$, each corresponding to a unique pair (b_1, b_2) with $b_1^3 - b_2^3 \in A$. But we must rule out those which have either b_1 or b_2 equal to zero. If $b_1 = 0$, then any choice of $b_2 \in A$ will satisfy that all the x_i are in A. Likewise, if $b_2 = 0$, then any choice of $b_1 \in A$ will do the same. Thus, there are $2m$ cases to eliminate, so

$$\#\{(x_1, x_2) \in A \times A \mid x_3 \in A, \ b_1 \neq 0, b_2 \neq 0\}$$
$$= mN_{A,A} - 2m. \tag{A.16}$$

By mimicking these calculations, it is not too difficult to see that we also have

$$\#\{(x_1, x_2) \in xA \times xA \mid x_3 \in xA, \ b_1 \neq 0, b_2 \neq 0\} \tag{A.17}$$
$$= \#\{(x_1, x_2) \in x^2A \times x^2A \mid x_3 \in x^2A, \ b_1 \neq 0, b_2 \neq 0\} \tag{A.18}$$
$$= mN_{A,A} - 2m. \tag{A.19}$$

Now consider the case where x_1, x_2, and x_3 are each in different cosets of A. We see that this rules out the case where either b_1 or b_2 is zero, since this would force all the x_i to be in the same coset.

Suppose $x_1 \in A$, $x_2 \in xA$, and $x_3 \in x^2 A$. Since $x_3 = -\zeta(x_1 + \zeta x_2)$, we must have $1 + \zeta x_2 x_1^{-1} \in x^2 A$, where we note that $x_2 x_1^{-1} \in xA$. For any fixed $x_1 \in A$, there are $N_{xA, x^2 A}$ choices for x_2 that satisfy this constraint. Thus, we arrive at

$$\#\{(x_1, x_2) \in A \times xA \mid x_3 \in x^2 A, \ b_1 \neq 0, b_2 \neq 0\}$$
$$= m N_{xA, x^2 A}. \tag{A.20}$$

With a little work exploiting Lemma 20, we see that we will arrive at the same result for any of the six permutations of the cosets corresponding to x_1, x_2, and x_3.

We comment that for any ordered pair $(a_1, a_2) \in A \times A$ such that $a_1 - a_2 \in A$ the nine pairs $(b_1, b_2) = (\zeta^{n_1} a_1^{1/3}, \zeta^{n_2} a_2^{1/3})$, for n_1 and n_2 ranging independently between 0 and 2, will all satisfy $(b_1^3, b_2^3) = (a_1, a_2)$. Thus, in counting the size of $\mathcal{A}_A$, we will have to add up our previous quantities from (A.16), (A.17), (A.18), and (A.20) (with multiplicities) and then divide by 9. This gives us

$$|\mathcal{A}_A| = \frac{1}{9}\left(3(m N_{A,A} - 2m) + 6m N_{xA, x^2 A}\right). \tag{A.21}$$

Finally, combining (A.13) and (A.21), and using Lemma 19 to make the substitution $N_1 = N_{A,A}$, we obtain the result for the case where $\zeta \in A$.

For the case where $\zeta \notin A$ we can verify that the relation does in fact still hold. It suffices to prove the result for when $\zeta \in xA$, for the result will also hold when $\zeta \in x^2 A$ due to the interchangeability of xA and $x^2 A$ which arises from both being multiplicative generators of H/A. In this case, we can show using similar counting arguments as before that for $d = 0, 1, 2$,

$$\#\{(x_1, x_2) \in x^d A \times x^d A \mid x_3 \in x^d A, b_1 \neq 0, b_2 \neq 0\}$$
$$= m N_{xA, x^2 A} - m, \tag{A.22}$$
$$\#\{(x_1, x_2) \in x^d A \times x^{d+1} A \mid x_3 \in x^{d+2} A, b_1 \neq 0, b_2 \neq 0\}$$
$$= m N_{x^2 A, xA} - m, \tag{A.23}$$
$$\#\{(x_1, x_2) \in x^d A \times x^{d+2} A \mid x_3 \in x^{d+1} A, b_1 \neq 0, b_2 \neq 0\}$$
$$= m N_{A,A}. \tag{A.24}$$

Summing these values up for $d = 1, 2, 3$, and again dividing by 9 and equating the value to (A.13), we

obtain

$$mN_1 = \frac{1}{9}\left(3(mN_{xA,x^2A} - m)\right)$$
$$+ \frac{1}{9}\left(3(mN_{x^2A,xA} - m) + 3mN_{A,A}\right), \qquad (A.25)$$

which after substituting $N_1 = N_{A,A}$ and $N_{x^2A,xA} = N_{xA,x^2A}$ (from Lemmas 19 and 20) reduces to the desired relation $N_{xA,x^2A} - N_{A,A} = 1$. $\qquad \square$

Lemma 24. *Let n be a prime, m a divisor of $n-1$ such that $\kappa := \frac{n-1}{m} = 3$, and A the subgroup of $(\mathbb{Z}/n\mathbb{Z})^{\times}$ of size m. Then if x is a multiplicative generator for $(\mathbb{Z}/n\mathbb{Z})^{\times}$, $\omega := e^{\frac{2\pi i}{n}}$, and $c_{\ell} = \frac{1}{m}\sum_{a\in A}\omega^{\ell a}$ is the inner product value corresponding to $\ell \in \mathbb{Z}/n\mathbb{Z}$, then*

$$\mathbf{cc}^* = \frac{1}{m}[I - \mathrm{diag}(\mathbf{c}) + P(I + B)C], \qquad (A.26)$$

where $\mathbf{c} = [c_1, c_x, c_{x^2}]^T$, I is the 3×3 identity matrix, and

$$B = \begin{bmatrix} N_1 & N_{x^2} & N_x \\ N_x & N_1 & N_{x^2} \\ N_{x^2} & N_x & N_1 \end{bmatrix}, \quad C = \begin{bmatrix} c_1 & c_{x^2} & c_x \\ c_x & c_1 & c_{x^2} \\ c_{x^2} & c_x & c_1 \end{bmatrix},$$

$$P = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}.$$

Proof. The terms of $\mathbf{cc}^*$ will take the form

$$c_{x^i}c_{x^j}^* = \frac{1}{m^2} \sum_{(a_1,a_2)\in A\times A} \omega^{x^i a_1 - x^j a_2}. \qquad (A.27)$$

If $i \neq j$, note that, $x^i a_1 - x^j a_2 \in x^d A$ if and only if $1 - x^{j-i}a_2 a_1^{-1} \in x^{d-i}A$, and there are $mN_{x^{j-i}A, x^{d-i}A}$ choices for (a_1, a_2) that satisfy this. Thus, we obtain

$$c_{x^i}c_{x^j}^* = \frac{1}{m}\sum_{d=0}^{\kappa-1} N_{x^{j-i}A, x^{d-i}A} c_{x^d}, \quad (i \neq j). \qquad (A.28)$$

If $i = j = d'$, (A.27) becomes $\frac{1}{m^2}\sum_{(a_1,a_2)\in A\times A}\omega^{x^{d'}(a_1 - a_2)}$. Separating the terms where $a_1 = a_2$, we

can apply the same reasoning as above and use Lemma 19 to obtain

$$|c_{x^{d'}}|^2 = \frac{1}{m}\left(1 + \sum_{d=0}^{\kappa-1} N_{x^d} c_{x^{d+d'}}\right). \tag{A.29}$$

Equation (A.26) can now be verified from (A.28) and (A.29) using Lemmas 19, 20, 21, 28, and 23. □

We are now ready to prove Theorem 8, which we restate here:

Theorem 8: *Let n be a prime, m a divisor of $n - 1$, and $\omega = e^{\frac{2\pi i}{n}}$. Let $A = \{a_1, ..., a_m\}$ be the unique subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$ of size m, and set $\mathbf{U} = diag(\omega^{a_1}, ..., \omega^{a_m}) \in \mathbb{C}^{m\times m}$, $\mathbf{v} = \frac{1}{\sqrt{m}}[1, ..., 1]^T \in \mathbb{C}^{m\times 1}$, and $\mathbf{M} = [\mathbf{v}, \mathbf{Uv}, ..., \mathbf{U}^{n-1}\mathbf{v}]$.*

If $\kappa := \frac{n-1}{m} = 3$, then the coherence of $\mathbf{M}$ will satisfy

$$\mu \le \frac{1}{3}\left(2\sqrt{\frac{1}{m}\left(3 + \frac{1}{m}\right)} + \frac{1}{m}\right) \approx \sqrt{\frac{4}{3m}}, \tag{A.30}$$

and for large enough m, we will asymptotically have the following lower bound on coherence:

$$\mu \ge \frac{1}{\sqrt{m}} \text{ (asymptotically)}, \tag{A.31}$$

which is strictly greater than the Welch bound.

Proof. (Theorem 8) Notice in (A.26) that B and C are circulant matrices (as is $I + B$), and hence they can be diagonalized by Fourier matrices. Let $\gamma = e^{2\pi i/3}$ and

$$F = \begin{bmatrix} 1 & 1 & 1 \\ 1 & \gamma & \gamma^2 \\ 1 & \gamma^2 & \gamma^4 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & \gamma & \gamma^{-1} \\ 1 & \gamma^{-1} & \gamma \end{bmatrix},$$

so that $\frac{1}{\sqrt{3}}F$ is the 3×3 discrete Fourier matrix. We first note that the matrix P from above is simply $\frac{1}{3}F^2 = \frac{1}{3}F^{*2}$. Now it is easy to verify that since $\mathbf{c}$ has real components by Lemma 18, then if we write $F\mathbf{c} =: [w_1, w_2, w_3]^T$, then we have that w_1 is real and $w_2 = w_3^*$. So we may write $w_1 = \alpha$, $w_2 = \beta e^{i\theta}$, and $w_3 = \beta e^{-i\theta}$, where α and β are real and β is nonnegative. If we let $\mathbf{a} = [N_1, N_x, N_{x^2}]^T$, then we can easily verify that by pre-multiplying Equation (A.26) by F and post-multiplying by F^*, noting that $FF^* = 3I$, $FPF^* = 3P$, $FCF^* = 3\,diag(F\mathbf{c})$ and $FBF^* = diag(F\mathbf{a})$, we can rewrite it as

$$(F\mathbf{c})(F\mathbf{c})^* = \frac{1}{m}[3I - F\,diag(\mathbf{c})F^*$$
$$+ 27P(I + diag(F\mathbf{a}))\,diag(F\mathbf{c})]. \tag{A.32}$$

One can further check that $F \operatorname{diag}(\mathbf{c})F^*$ is circulant with first column $F\mathbf{c}$, and if we write $F\mathbf{a} = [y_1, y_2, y_3]^T$, then (A.32) becomes

$$
\begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} [w_1^*, w_2^*, w_3^*] = \frac{1}{m} \left[\begin{bmatrix} 3 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{bmatrix} - \begin{bmatrix} w_1 & w_3 & w_2 \\ w_2 & w_1 & w_3 \\ w_3 & w_2 & w_1 \end{bmatrix} \right. \tag{A.33}
$$

$$
+ 27 \left. \begin{bmatrix} (1+y_1)w_1 & 0 & 0 \\ 0 & 0 & (1+y_3)w_3 \\ 0 & (1+y_2)w_2 & 0 \end{bmatrix} \right].
$$

If we consider only the coordinates of the above matrices which do not involve y_1, y_2 or y_3, then after substituting $w_1 = \alpha$, $w_2 = \beta e^{i\theta}$ and $w_3 = \beta e^{-i\theta}$, we can solve the resulting equations to obtain the relations

$$
\alpha = -\frac{1}{m}, \qquad \beta = \sqrt{\frac{1}{m}\left(3 + \frac{1}{m}\right)}. \tag{A.34}
$$

We can use these to bound the coherence as follows:

$$
\begin{bmatrix} c_1 \\ c_x \\ c_{x^2} \end{bmatrix} = F^{-1} \begin{bmatrix} \alpha \\ \beta e^{i\theta} \\ \beta e^{-i\theta} \end{bmatrix} = \frac{1}{3} \begin{bmatrix} \alpha + 2\beta \cos(\theta) \\ \alpha + 2\beta \cos(\theta - \frac{2\pi}{3}) \\ \alpha + 2\beta \cos(\theta + \frac{2\pi}{3}) \end{bmatrix}. \tag{A.35}
$$

$$
\min_{\theta} \max\{|c_1|, |c_x|, |c_{x^2}|\} \leq \mu \leq \max_{\theta} \max\{|c_1|, |c_x|, |c_{x^2}|\} \tag{A.36}
$$

From (A.34), we know that α is negative, and β is positive by definition. Since $|\alpha| < |\beta|$, then by inspection we have

$$
\max_{\theta} \max\{|c_1|, |c_x|, |c_{x^2}|\} = \frac{1}{3}|\alpha + 2\beta(-1)| \tag{A.37}
$$

$$
= \frac{1}{3}\left(2\sqrt{\frac{1}{m}\left(3 + \frac{1}{m}\right)} + \frac{1}{m}\right). \tag{A.38}
$$

This gives us our upper bound. Asymptotically, we can ignore the term $\alpha = -\frac{1}{m}$ in our expressions for c_1, c_x, and c_{x^2}, and if we do so, we find that

$$
\arg\min_{\theta} \max\{|c_1|, |c_x|, |c_{x^2}|\} \approx \frac{\pi}{2},
$$

which follows from noting that since $|c_1|, |c_x|$, and $|c_{x^2}|$ are continuous functions of θ, the smallest value of their maximum must occur when two of them are set equal to each other (in this case, when $|c_x| = |c_{x^2}|$, so that asymptotically $|\cos(\theta + \frac{2\pi}{3})| = |\cos(\theta - \frac{2\pi}{3})|$). Substituting $\frac{\pi}{2}$ for θ gives us our (asymptotic) lower bound on μ:

$$\min_{\theta} \max\{|c_1|, |c_x|, |c_{x^2}|\} \approx \frac{1}{\sqrt{m}}. \tag{A.39}$$

We easily verify that this is greater than the Welch bound, which in this case becomes

$$\sqrt{\frac{n - m}{m(n - 1)}} = \sqrt{\frac{2}{3m} + \frac{1}{3m^2}}.$$

$\square$

Appendix B

Chapter 4 Proofs

B.1 Universal Upper Bound On Our Frame Coherence: Proof of Theorems 9, 10, and 18

In this section, we return to the framework of Theorems 9, 10, and 18. Let p be a prime and r a a positive integer, and set our group G (in Theorem 17) to be the finite field $\mathbb{F}_{p^r}$.

Our frame matrix $\mathbf{M}$ from (4.44) will take the form

$$\mathbf{M} = \begin{bmatrix} \omega^{\mathrm{Tr}(a_1 x_1)} & \omega^{\mathrm{Tr}(a_1 x_2)} & \ldots & \omega^{\mathrm{Tr}(a_1 x_n)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \omega^{\mathrm{Tr}(a_m x_1)} & \omega^{\mathrm{Tr}(a_m x_2)} & \ldots & \omega^{\mathrm{Tr}(a_m x_n)} \end{bmatrix}, \tag{B.1}$$

where $\omega = e^{\frac{2\pi i}{p}}$ and we have expressed the elements of $\mathbb{F}_{p^r}$ as $\{x_i\}_{i=1}^n$. In terms of powers of x, we may relabel these elements as $x_1 = 0$, and $x_i = x^{i-1}$, $i = 2, ..., n = p^r$. Note that with this relabeling, the first column of $\mathbf{M}$ is all 1's.

As we commented before Equation (4.46), if $x_j - x_i = x_\ell$, the inner product between the i^{th} and j^{th} columns is

$$\sum_{a_t} \left(\omega^{\mathrm{Tr}(a_t x_i)}\right)^* \left(\omega^{\mathrm{Tr}(a_t x_j)}\right) = \sum_{a_t} \omega^{\mathrm{Tr}(a_t(x_j - x_i))} \tag{B.2}$$

$$= \sum_{a_t} \omega^{\mathrm{Tr}(a_t x_\ell)}. \tag{B.3}$$

We will be making extensive use of the sums in (B.3) in this section, so we will make the following definition:

Definition 11. *For any $z \in \mathbb{F}_{p^r}$ and A a subgroup of $\mathbb{F}_{p^r}^{\times}$, we will define c_z to be the normalized inner*

product sum corresponding to z, that is,

$$c_z = \frac{1}{m} \sum_{a \in A} \omega^{\text{Tr}(az)}, \tag{B.4}$$

where $\omega = e^{2\pi i/p}$.

The following property of the values c_z is simple, but worth establishing:

Lemma 25. *For any* $z \in \mathbb{F}_{p^r}$, *we have* $c_z^* = c_{-z}$.

Proof. Expanding c_z as a sum, we have

$$(c_z)^* = \left(\frac{1}{m} \sum_{a \in A} \omega^{\text{Tr}(za)} \right)^* \tag{B.5}$$

$$= \frac{1}{m} \sum_{a \in A} \omega^{-\text{Tr}(za)} \tag{B.6}$$

$$= c_{-z}. \tag{B.7}$$

$\square$

Recall that the set of nonzero field elements $\mathbb{F}_{p^r}^{\times}$ is a cyclic group under multiplication, so let x be a multiplicative generator. The elements of $\mathbb{F}_{p^r}$ can now be expressed as $\{0, 1, x, ..., x^{p^r-1}\}$. The inner product corresponding to 0 is simply $c_0 = 1$, which arises only when taking the inner product of a frame element with itself. The nontrivial inner products are thus c_{x^i}, for $i = 0, ..., p^r - 1$.

We point out that if A is a size-m multiplicative subgroup, it is unique (since $\mathbb{F}_{p^r}$ is cyclic) and is a cyclic group generated by x^κ, where $\kappa = \frac{p^r-1}{m}$. The cosets of A are $A, xA, ..., x^{\kappa-1}A$. One interesting observation is that elements in the same coset of A give rise to the same inner product value:

Lemma 26. *If* z *is in the coset* $x^i A$, *then* $c_z = c_{x^i}$.

Proof. Write $z = x^i a_z$, for some $a_z \in A$. Then,

$$c_z = \sum_{a \in A} \omega^{\text{Tr}(x^i a_z \cdot a)} \tag{B.8}$$

$$= \sum_{a \in A} \omega^{\text{Tr}(x^i a)}, \tag{B.9}$$

where $\omega = e^{2\pi i/p}$ and the last equality follows from the fact that since A is a group, multiplication by a_z simply permutes its elements. $\square$

In light of Lemma 26, we see concretely that there is indeed only a single nontrivial inner product value for each coset of A, and each arises with the same multiplicity (because each coset has the same

number of elements). Furthermore, since $\{1, x, ..., x^{\kappa-1}\}$ is a set of representatives for each of the cosets of A, we only need to be concerned with the values c_{x^i}, $i = 0, 1, ..., \kappa - 1$. The largest absolute value of these will be the coherence.

Lemma 27. *The values $c_1, c_x, ..., c_{x^{\kappa}-1}$ satisfy the equation*

$$1 + mc_1 + mc_x + ... + mc_{x^{\kappa}-1} = 0, \tag{B.10}$$

where m is the size of A.

Proof. If we expand the sum in (B.10) using the fact that each of the m elements $z \in x^d A$ satisfies $c_z = c_{x^d}$, and that $c_0 = 1$, we get

$$1 + \sum_{d=1}^{\kappa-1} mc_{x^d} = \sum_{z \in \mathbb{F}_{p^r}} c_z \tag{B.11}$$

$$= \sum_{a \in A} \sum_{z \in \mathbb{F}_{p^r}} \omega^{\mathrm{Tr}(za)}, \tag{B.12}$$

where $\omega = e^{\frac{2\pi i}{p}}$. But the function $\chi_{reg}(y) := \sum_{z \in \mathbb{F}_{p^r}} \omega^{\mathrm{Tr}(zy)}$ which arises as the internal sum in (B.12) is the well-known character of the "regular representation" of $\mathbb{F}_{p^r}$, which is equal to p^r if $y = 0$ and 0 otherwise [87]. Since no elements of A are 0, we see that (B.12) sums to zero. $\qquad\square$

Our following work will involve taking many sums and products of field elements, and determining in which coset of A they lie. While it is in general easy to determine in which coset a product lies (for example, if $z_1 \in x^{i_1} A$ and $z_2 \in x^{i_2} A$, then $z_1 z_2 \in x^{i_1+i_2} A$), it is often not obvious in which coset a sum lies. To get around this problem, we will make use of the following quantities, which are the natural generalization of the translation degrees we defined in Definition 10 of Appendix A:

Definition 12. *Given two cosets $x_1 A$ and $x_2 A$, we define the translation degree from $x_1 A$ to $x_2 A$ to be the quantity*

$$N_{x_1 A, x_2 A} = \#\{z \in x_1 A \mid 1 + z \in x_2 A\} = |1 + x_1 A \cap x_2 A|. \tag{B.13}$$

Likewise, we define $N_{x_1 A, 0}$ and $N_{0, x_2 A}$ (the translation degrees from $x_1 A$ to 0 and from 0 to $x_2 A$, respectively) to be

$$N_{x_1 A, 0} = |\{-1\} \cap x_1 A|, \tag{B.14}$$

$$N_{0, x_2 A} = |\{1\} \cap x_2 A|. \tag{B.15}$$

We will quickly point out a simple property of the translation degrees:

Lemma 28. *Set $H = \mathbb{F}_{p^r}^\times$. For any coset $x_0 A$, we have*

$$N_{x_0 K, 0} + \sum_{x_i K \in H/A} N_{x_0 A, x_i A} = |x_0 A|.$$

In particular, if $-1 \in x_0 A$, this equation reduces to

$$1 + N_{x_0 A, A} + N_{x_0 A, xA} + N_{x_0 A, x^2 A} + \ldots + N_{x_0 A, x^{\kappa-1} A} = m,$$

and if $-1 \notin x_0 A$, this equation becomes

$$N_{x_0 A, A} + N_{x_0 A, xA} + N_{x_0 A, x^2 A} + \ldots + N_{x_0 A, x^{\kappa-1} A} = m.$$

Proof. This simply follows from the observation that any of the m elements of $x_0 A$, when added to 1, must either be equal to 0 or lie in exactly one of the cosets $x_i A \in H/A$. $\square$

The following lemma will be instrumental in bounding these inner product values.

Lemma 29. *Let $\mathbf{c} = [c_1, c_x, c_{x^2}, \ldots, c_{x^{\kappa-1}}]^T$, and let F be the scaled $\kappa \times \kappa$ Fourier matrix with entries defined by $F_{ij} = \gamma^{(i-1)(j-1)}$, where $\gamma = e^{2\pi i/\kappa}$. Then, if we let $\mathbf{w} := [w_1, \ldots, w_\kappa]^T = F\mathbf{c}$ so that $w_{d+1} = \sum_{t=0}^{\kappa-1} \gamma^{td} c_{x^t}$ for $d = 0, 1, \ldots, \kappa - 1$, we have*

$$w_1 = -\frac{1}{m}, \tag{B.16}$$

$$|w_i| = \sqrt{\frac{1}{m}\left(\kappa + \frac{1}{m}\right)}, \qquad i \neq 1. \tag{B.17}$$

Proof. For any $d = 0, 1, \ldots, \kappa - 1$, we have

$$|w_{d+1}|^2 = \left(\sum_{t=0}^{\kappa-1} \gamma^{td} c_{x^t} \right) \left(\sum_{\ell=0}^{\kappa-1} \gamma^{\ell d} c_{x^\ell} \right)^* \tag{B.18}$$

$$= \left(\sum_{t=0}^{\kappa-1} \gamma^{td} c_{x^t} \right) \left(\sum_{\ell=0}^{\kappa-1} \gamma^{-\ell d} c_{-x^\ell} \right) \tag{B.19}$$

$$= \sum_{t=0}^{\kappa-1} \sum_{\ell=0}^{\kappa-1} \gamma^{(t-\ell)d} c_{x^t} c_{-x^\ell} \tag{B.20}$$

$$= \sum_{s=0}^{\kappa-1} \sum_{\ell=0}^{\kappa-1} \gamma^{sd} c_{x^{s+\ell}} c_{-x^\ell} \tag{B.21}$$

$$= \sum_{s=0}^{\kappa-1} \gamma^{sd} \sum_{\ell=0}^{\kappa-1} c_{x^{s+\ell}} c_{-x^\ell}. \tag{B.22}$$

Now, we note that

$$m^2 c_{x^{s+\ell}} c_{-x^\ell} = \left(\sum_{a \in A} \omega^{\mathrm{Tr}(x^{s+\ell} a)} \right) \left(\sum_{a' \in A} \omega^{\mathrm{Tr}(-x^\ell a')} \right) \tag{B.23}$$

$$= \sum_{a,a' \in A} \omega^{\mathrm{Tr}(x^{s+\ell} a - x^\ell a')} \tag{B.24}$$

$$= \sum_{a,a' \in A} \omega^{\mathrm{Tr}(-x^\ell a'(1 - x^s a a'^{-1}))} \tag{B.25}$$

$$= \sum_{a',a'' \in A} \omega^{\mathrm{Tr}(-x^\ell a'(1 - x^s a''))} \tag{B.26}$$

$$= \sum_{t=0}^{\kappa-1} \sum_{\substack{\{a',a'' \in A \ : \\ 1 - x^s a'' \in x^t A\}}} \omega^{\mathrm{Tr}(-x^\ell a'(1 - x^s a''))} \tag{B.27}$$

$$+ \sum_{\substack{\{a',a'' \in A \ : \\ 1 - x^s a'' = 0\}}} 1$$

$$= \sum_{t=0}^{\kappa-1} N_{-x^s A, x^t A} \left(\sum_{a''' \in A} \omega^{\mathrm{Tr}(-x^t x^\ell a''')} \right) \tag{B.28}$$

$$+ \sum_{a' \in A} N_{-x^s A, 0} \tag{B.29}$$

$$= m \sum_{t=0}^{\kappa-1} N_{-x^s A, x^t A} \cdot c_{-x^{t+\ell}} + m N_{-x^s A, 0}. \tag{B.30}$$

Now we can substitute this into (B.22), and we obtain:

$$|w_{d+1}|^2 = \sum_{s=0}^{\kappa-1} \gamma^{sd} \sum_{\ell=0}^{\kappa-1} \left(\frac{1}{m} \left(\sum_{t=0}^{\kappa-1} N_{-x^s A, x^t A} \cdot c_{-x^{t+\ell}} + N_{-x^s A, 0} \right) \right) \tag{B.31}$$

$$= \sum_{s=0}^{\kappa-1} \gamma^{sd} \frac{1}{m} \left(\sum_{t=0}^{\kappa-1} N_{-x^s A, x^t A} \sum_{\ell=0}^{\kappa-1} c_{-x^{t+\ell}} + \sum_{\ell=0}^{\kappa-1} N_{-x^s A, 0} \right) \tag{B.32}$$

$$= \sum_{s=0}^{\kappa-1} \gamma^{sd} \frac{1}{m} \left(\sum_{t=0}^{\kappa-1} N_{-x^s A, x^t A} \left(-\frac{1}{m} \right) + \kappa N_{-x^s A, 0} \right) \tag{B.33}$$

$$= -\frac{1}{m^2} \sum_{s=0}^{\kappa-1} \gamma^{sd} \sum_{t=0}^{\kappa-1} N_{-x^s A, x^t A} + \frac{1}{m} \sum_{s=0}^{\kappa-1} \gamma^{sd} \kappa N_{-x^s A, 0} \tag{B.34}$$

$$= -\frac{1}{m^2} \sum_{s=0}^{\kappa-1} \gamma^{sd} \left(m - N_{-x^s A, 0} \right) + \frac{\kappa}{m} \sum_{s=0}^{\kappa-1} \gamma^{sd} N_{-x^s A, 0}, \tag{B.35}$$

where (B.33) follows from Equation (B.10), and (B.35) follows from Lemma 28. Note that $N_{-x^s A, 0}$ is equal to 1 if $s = 0$ and equal to 0 otherwise. Thus, (B.35) becomes

$$|w_{d+1}|^2 = -\frac{1}{m^2} \left((m-1) + m \sum_{s=1}^{\kappa-1} \gamma^{sd} \right) + \frac{\kappa}{m}. \tag{B.36}$$

Now, if $d \neq 0$, then $\sum_{s=1}^{\kappa-1} \gamma^{sd} = -1$, and after rearranging terms we obtain

$$|w_{d+1}|^2 = \frac{1}{m} \left(\kappa + \frac{1}{m} \right). \tag{B.37}$$

If $d = 0$, then $\sum_{s=1}^{\kappa-1} \gamma^{sd} = m$, and (B.36) gives us $|w_1|^2 = \frac{1}{m^2}$. In fact, in this case, we can compute w_1 explicitly, since

$$w_1 = \sum_{t=0}^{\kappa-1} c_{x^t} = -\frac{1}{m}. \tag{B.38}$$

$\square$

We can now use Lemma 29 to bound the coherence of our frames constructed from finite fields.

Theorem 29. *Let $G = \mathbb{F}_{p^r}$ be the finite field with elements $\{x_1, ..., x_{p^r}\}$, and $H = \mathbb{F}_{p^r}^\times$ the (cyclic) multiplicative group of the nonzero field elements. If A is the unique subgroup of H of size m, with elements $\{a_1, ..., a_m\}$, and $\mathbf{M}$ is the frame with columns defined in (B.1), then the coherence μ of $\mathbf{M}$ is upper-bounded by*

$$\mu \leq \frac{1}{\kappa} \left((\kappa - 1) \sqrt{\frac{1}{m} \left(\kappa + \frac{1}{m} \right)} + \frac{1}{m} \right). \tag{B.39}$$

Proof. The proof follows from Lemma 29. Using the notation of this lemma, we may write $\mathbf{c} = \frac{1}{\kappa} F^* \mathbf{w}$, so that

$$|c_{x^d}| = \frac{1}{\kappa} \left| \sum_{j=1}^{\kappa} \gamma^{d(j-1)} w_j \right| \tag{B.40}$$

$$\leq \frac{1}{\kappa} \sum_{j=1}^{\kappa} |w_j| \tag{B.41}$$

$$= \frac{1}{\kappa} \left((\kappa - 1) \sqrt{\frac{1}{m} \left(\kappa + \frac{1}{m} \right)} + \frac{1}{m} \right), \tag{B.42}$$

where (B.41) follows from the triangle inequality and (B.42) follows from Lemma 29. Since the coherence is equal to the largest value among the $|c_{x^d}|$, $d = 0, ..., p^r - 1$, the result now follows immediately. $\square$

Recall from Theorems 10 and 18 that when the size m of A happens to be an odd integer, we can derive even tighter bounds on coherence, provided that p is an odd prime. (Note that in our original framework of Theorem 6, when m was taken to be a divisor of $p - 1$, the only case where p could be even was when $p = 2$ and $m = 1$ in which case our frames would be 1-dimensional and trivially have coherence equal to 1.) We will prove this result shortly, but we first present the following equivalent condition on A for when its size is even or odd.

Lemma 30. *Let p be a prime, r an integer, m a divisor of $p^r - 1$, and $\kappa := \frac{p^r - 1}{m}$. Let $\mathbb{F}_{p^r}$ be the finite field with p^r elements, whose multiplicative group $\mathbb{F}_{p^r}^{\times}$ has cyclic generator x, and let A be the unique subgroup of $\mathbb{F}_{p^r}^{\times}$ of size m. Then $-1 \in A$ if and only if either p or m is even. If p and m are both odd, then κ is even and $-1 \in x^{\frac{\kappa}{2}} A$.*

Proof. If p is even, that is $p = 2$, then $-1 \equiv 1$ in $\mathbb{F}_{p^r}$, so trivially $-1 \in A$. If p is odd, then the order m of A is even if and only if A has a subgroup of size 2, which means there is a nontrivial element in A which is a root of the polynomial $X^2 - 1$. The element -1 is the only such root.

If both m and p are odd, then $p^r - 1$ must be even, hence so is $\kappa = \frac{p^r - 1}{m}$. In this case, since the square of -1 obviously lies in A (which is equal to $x^{\kappa} A$), we must have $-1 \in x^{\frac{\kappa}{2}} A$. $\square$

We need one more tool before we can prove our tighter bound:

Lemma 31. *Let p, r, m, κ, x, and A be defined as in Lemma 30 and $\mathbf{w} = [w_1, ..., w_{\kappa}]^T$ be defined as in Lemma 29. If either p or m is even ($-1 \in A$) then for any $d = 0, 1, ..., \kappa - 1$, we have $c_{x^d} = c_{x^d}^*$, and for any $i = 2, 3, ..., \kappa$ we have $w_i^* = w_{\kappa - i + 2}$. If p and m are both odd ($-1 \in x^{\frac{\kappa}{2}} A$), then $c_{x^d} = c_{x^{d + \kappa/2}}^*$ and $w_i^* = (-1)^{i-1} w_{\kappa - i + 2}$.*

Proof. As usual, set $\omega = e^{2\pi i/p}$ and $\gamma := e^{2\pi i/\kappa}$. If $-1 \in A$, then multiplication by -1 permutes the elements of A, so we have

$$c_{x^d}^* = \left(\frac{1}{m} \sum_{a \in A} \omega^{x^d a} \right)^* \tag{B.43}$$

$$= \frac{1}{m} \sum_{a \in A} \omega^{-x^d a} \tag{B.44}$$

$$= \frac{1}{m} \sum_{a \in A} \omega^{x^d a} \tag{B.45}$$

$$= c_{x^d}. \tag{B.46}$$

It follows that c_{x^d} is real. Furthermore, in this case we have

$$w_i^* = \left(\sum_{j=1}^{\kappa} \gamma^{(i-1)(j-1)} c_{x^{j-1}} \right)^* \tag{B.47}$$

$$= \sum_{j=1}^{\kappa} \gamma^{-(i-1)(j-1)} c_{x^{j-1}}^* \tag{B.48}$$

$$= \sum_{j=1}^{\kappa} \gamma^{(-i+1)(j-1)} c_{x^{j-1}} \tag{B.49}$$

$$= \sum_{j=1}^{\kappa} \gamma^{(\kappa-i+1)(j-1)} c_{x^{j-1}} \tag{B.50}$$

$$= \sum_{j=1}^{\kappa} \gamma^{((\kappa-i+2)-1)(j-1)} c_{x^{j-1}} \tag{B.51}$$

$$= w_{\kappa-i+2}. \tag{B.52}$$

Now, if instead $-1 \in x^{\frac{\kappa}{2}} A$, then multiplication by $-x^{\frac{\kappa}{2}}$ permutes the elements of A, and we have

$$c_{x^d} = \frac{1}{m} \sum_{a \in A} \omega^{x^d a} \tag{B.53}$$

$$= \frac{1}{m} \sum_{a \in A} \omega^{-x^d x^{\frac{\kappa}{2}} a} \tag{B.54}$$

$$= \left(\frac{1}{m} \sum_{a \in A} \omega^{x^{d+\frac{\kappa}{2}} a} \right)^* \tag{B.55}$$

$$= c_{x^{d+\kappa/2}}^*. \tag{B.56}$$

Also in this case, we may write

$$w_i^* = \left(\sum_{j=1}^{\kappa} \gamma^{(i-1)(j-1)} c_{x^{j-1}} \right)^* \tag{B.57}$$

$$= \sum_{j=1}^{\kappa} \gamma^{-(i-1)(j-1)} c_{x^{j-1}}^* \tag{B.58}$$

$$= \sum_{j=1}^{\kappa} \gamma^{-(i-1)(j-1)} c_{x^{j-1+\frac{\kappa}{2}}} \tag{B.59}$$

$$= \sum_{j=1}^{\kappa} \gamma^{-(i-1)(j-1+\frac{\kappa}{2})} \gamma^{(i-1)\frac{\kappa}{2}} c_{x^{j-1+\frac{\kappa}{2}}} \tag{B.60}$$

$$= \gamma^{(i-1)\frac{\kappa}{2}} \sum_{j=1}^{\kappa} \gamma^{(\kappa-i+1)(j-1+\frac{\kappa}{2})} c_{x^{j-1+\frac{\kappa}{2}}} \tag{B.61}$$

$$= \gamma^{(i-1)\frac{\kappa}{2}} w_{\kappa-i+2} \tag{B.62}$$

$$= (-1)^{i-1} w_{\kappa-i+2}, \tag{B.63}$$

where the last line follows from the fact that $\gamma^{\frac{\kappa}{2}} = -1$. This completes the proof of the lemma. $\square$

We are now equipped to prove the second part of Theorem 18, which we restate here for convenience:

Theorem 30. *Let p be a prime, r a positive integer, m a divisor of $p^r - 1$, and $A = \{a_i\}_{i=1}^m$ the unique subgroup of $\mathbb{F}_{p^r}^\times$ of size m. Then setting $\omega = e^{\frac{2\pi i}{p}}$ and $\kappa := \frac{p^r - 1}{m}$, if both p and m are odd, the coherence μ of our frame $\mathbf{M}$ in (B.1) satisfies*

$$\mu \le \frac{1}{\kappa} \sqrt{\left(\frac{1}{m} + \left(\frac{\kappa}{2} - 1 \right) \beta \right)^2 + \left(\frac{\kappa}{2} \right)^2 \beta^2}, \tag{B.64}$$

where $\beta = \sqrt{\frac{1}{m} \left(\kappa + \frac{1}{m} \right)}$.

Proof. Since both p and m are odd, then from Lemma 30 we know that κ is even and -1 lies in the coset $x^{\frac{\kappa}{2}} A$. There is a 1-1 correspondence between the set of integers $\{1, ..., \kappa\}$ and itself which sends $j \mapsto \kappa - j + 2 \mod \kappa$, for $j = 1, ..., \kappa$. This mapping fixes the singletons $\{1\}$ and $\{\frac{\kappa}{2} + 1\}$ and interchanges the elements in the pairs $\{j, \kappa - j + 2\}$ for $j = 2, ..., \frac{\kappa}{2}$.

As in Lemma 29, set $\mathbf{c} = [c_1, c_x, c_{x^2}, ..., c_{x^{\kappa-1}}]^T$ and $\mathbf{w} := [w_1, ..., w_\kappa]^T = F\mathbf{c}$, where F is the scaled $\kappa \times \kappa$ Fourier matrix with entries $F_{ij} = \gamma^{(i-1)(j-1)}$, where $\gamma = e^{2\pi i/\kappa}$. Since $-1 \in x^{\frac{\kappa}{2}} A$, then by

Lemmas 29 and 31 we have

$$w_j \cdot w_{\kappa-j+2} = w_j \left(\left((-1)^{j-1} \right)^{-1} w_j^* \right) \tag{B.65}$$

$$= (-1)^{j-1} |w_j|^2 \tag{B.66}$$

$$= (-1)^{j-1} \beta^2, \tag{B.67}$$

where $\beta = \sqrt{\frac{1}{m} \left(\kappa + \frac{1}{m} \right)}$.

We quickly note that given integers i and j, the conjugate of $\gamma^{-(i-1)(j-1)}$ can be expressed as

$$\left(\gamma^{-(i-1)(j-1)} \right)^* = \gamma^{-(i-1)(-j+1)} \tag{B.68}$$

$$= \gamma^{-(i-1)(\kappa-j+1)} \tag{B.69}$$

$$= \gamma^{-(i-1)((\kappa-j+2)-1)}. \tag{B.70}$$

Note that the inverse of F is $\frac{1}{\kappa} F^*$. From the equation $\mathbf{c} = \frac{1}{\kappa} F^* \mathbf{w}$, we may write

$$c_{x^{i-1}} = \frac{1}{\kappa} \sum_{j=1}^{\kappa} \gamma^{-(i-1)(j-1)} w_j. \tag{B.71}$$

Now we can group the terms of the summation of $c_{x^{i-1}}$ by our above subsets of indices ($\{j, \kappa-j+2\}$ for $j = 2, \ldots, \frac{\kappa}{2}$) as follows:

$$c_{x^{i-1}} = \frac{1}{\kappa} \left[w_1 + \gamma^{-(i-1)\frac{\kappa}{2}} w_{\frac{\kappa}{2}+1} + \sum_{j=2}^{\frac{\kappa}{2}} \left(\gamma^{-(i-1)(j-1)} w_j + \gamma^{-(i-1)((\kappa-j+2)-1)} w_{\kappa-j+2} \right) \right]. \tag{B.72}$$

We know from Lemma 29 that $w_1 = -\frac{1}{m}$ and $|w_{\frac{\kappa}{2}+1}| = \beta$. Also,

$$\gamma^{-(i-1)\frac{\kappa}{2}} = (\gamma^{\frac{\kappa}{2}})^{-(i-1)} = (-1)^{-(i-1)} = (-1)^{i-1},$$

and from Lemma 31 we know that $w_{\frac{\kappa}{2}+1} = (-1)^{\frac{\kappa}{2}} w_{\frac{\kappa}{2}+1}^*$. Thus, if $\frac{\kappa}{2}$ is even we have that $w_{\frac{\kappa}{2}+1}$ is purely real, so $\gamma^{-(i-1)\frac{\kappa}{2}} w_{\frac{\kappa}{2}+1} = \pm\beta$. And if $\frac{\kappa}{2}$ is odd, we have that $w_{\frac{\kappa}{2}+1}$ is purely imaginary, in which case $\gamma^{-(i-1)\frac{\kappa}{2}} w_{\frac{\kappa}{2}+1} = \pm i\beta$.

From these observations and Lemma 31, we have

$$\gamma^{-(i-1)(j-1)} w_j + \gamma^{-(i-1)((\kappa-j+2)-1)} w_{\kappa-j+2}$$

$$= \gamma^{-(i-1)(j-1)} w_j + (-1)^{j-1} \gamma^{-(i-1)((\kappa-j+2)-1)} w_j^* \tag{B.73}$$

$$= \gamma^{-(i-1)(j-1)} w_j + (-1)^{j-1} \left(\gamma^{-(i-1)(j-1)} w_j \right)^*. \tag{B.74}$$

If j is even (B.74) becomes $2i\Im(\gamma^{-(i-1)(j-1)}w_j)$ and if j is odd it becomes $2\Re(\gamma^{-(i-1)(j-1)}w_j)$, where $\Im(z)$ and $\Re(z)$ denote the imaginary and real parts of the complex number z respectively. If we define the phase θ_j such that $w_j = \beta e^{i\theta_j}$, we can further express these as

$$2i\Im(\gamma^{-(i-1)(j-1)}w_j) = 2i\beta \sin\left(\theta_j - \frac{2\pi}{\kappa}(i-1)(j-1)\right) \tag{B.75}$$

and

$$2\Re(\gamma^{-(i-1)(j-1)}w_j) = 2\beta \cos\left(\theta_j - \frac{2\pi}{\kappa}(i-1)(j-1)\right). \tag{B.76}$$

To simplify our notation, we will define

$$\tilde{\theta}_j := \theta_j - \frac{2\pi}{\kappa}(i-1)(j-1),$$

allowing us to write the summation in (B.72) as

$$\sum_{j=2}^{\frac{\kappa}{2}}\left(\gamma^{-(i-1)(j-1)}w_j + \gamma^{-(i-1)((\kappa-j+2)-1)}w_{\kappa-j+2}\right)$$

$$= \sum_{j \text{ even}} 2i\beta \sin(\tilde{\theta}_j) + \sum_{j \text{ odd}} 2\beta \cos(\tilde{\theta}_j).$$

Now, we can bound the coherence by

$$\mu \leq \max_{\{\theta_j\}} \max_i |c_{x^{i-1}}| \leq \max_i \max_{\{\theta_j\}} |c_{x^{i-1}}|,$$

and from our above discussion this becomes

$$\max_{\{\tilde{\theta}_j\}} \frac{1}{\kappa}\left| -\frac{1}{m} \pm \beta + \sum_{j \text{ even}} 2i\beta \sin(\tilde{\theta}_j) + \sum_{j \text{ odd}} 2\beta \cos(\tilde{\theta}_j)\right| \tag{B.77}$$

if $\frac{\kappa}{2}$ is even, and

$$\max_{\{\tilde{\theta}_j\}} \frac{1}{\kappa}\left| -\frac{1}{m} \pm i\beta + \sum_{j \text{ even}} 2i\beta \sin(\tilde{\theta}_j) + \sum_{j \text{ odd}} 2\beta \cos(\tilde{\theta}_j)\right| \tag{B.78}$$

if $\frac{\kappa}{2}$ is odd.

If we set

$$n_e := \#\{j \text{ even} \mid 2 \leq j \leq \frac{\kappa}{2}\}$$

and

$$n_o := \#\{j \text{ odd} \mid 2 \leq j \leq \frac{\kappa}{2}\},$$

then by speculation, (B.77) becomes bounded by

$$\frac{1}{\kappa}\left|\frac{1}{m} + \beta + n_e 2i\beta + n_o 2\beta\right| = \frac{1}{\kappa}\sqrt{\left(\frac{1}{m} + \beta(1+2n_o)\right)^2 + (2n_e\beta)^2} \tag{B.79}$$

and (B.78) becomes bounded by

$$\frac{1}{\kappa}\left|\frac{1}{m} + i\beta + n_e 2i\beta + n_o 2\beta\right| = \frac{1}{\kappa}\sqrt{\left(\frac{1}{m} + 2n_o\beta\right)^2 + \beta^2(1+2n_e)^2}. \tag{B.80}$$

Finally, we note that when $\frac{\kappa}{2}$ is even, then $n_e = \frac{\kappa}{4}$ (half the numbers between 1 and $\frac{\kappa}{2}$, inclusive, are even), and hence $n_o = \left(\frac{r}{\kappa} - 1\right) - n_e = \frac{\kappa}{4} - 1$. Thus, (B.79) becomes

$$\frac{1}{\kappa}\sqrt{\left(\frac{1}{m} + \beta\left(1 + 2\left(\frac{\kappa}{4} - 1\right)\right)\right)^2 + \left(2 \cdot \frac{\kappa}{4} \cdot \beta\right)^2} \tag{B.81}$$

$$= \frac{1}{\kappa}\sqrt{\left(\frac{1}{m} + \left(\frac{\kappa}{2} - 1\right)\beta\right)^2 + \left(\frac{\kappa}{2}\right)^2 \beta^2}. \tag{B.82}$$

We get the same bound when $\frac{\kappa}{2}$ is odd. Indeed, in this case $n_e = \frac{1}{2}\left(\frac{\kappa}{2} - 1\right) = \frac{\kappa}{4} - \frac{1}{2}$ (now half the numbers between 1 and $\frac{\kappa}{2} - 1$, inclusive, are even), and $n_o = \left(\frac{\kappa}{2} - 1\right) - n_e = \frac{\kappa}{4} - \frac{1}{2}$. Then (B.80) also becomes

$$\frac{1}{\kappa}\sqrt{\left(\frac{1}{m} + 2\left(\frac{\kappa}{4} - \frac{1}{2}\right)\beta\right)^2 + \beta^2\left(1 + 2\left(\frac{\kappa}{4} - \frac{1}{2}\right)\right)^2} \tag{B.83}$$

$$= \frac{1}{\kappa}\sqrt{\left(\frac{1}{m} + \left(\frac{\kappa}{2} - 1\right)\beta\right)^2 + \left(\frac{\kappa}{2}\right)^2 \beta^2}. \tag{B.84}$$

This concludes the proof. □

Remark: If were to mimic the proof of Theorem 10 in the case when m is even (so $-1 \in A$ and the c_x^d are real), then we would arrive at the same bound as in Theorem 29. Indeed, in this case from Lemma 31 we have

$$\left(\gamma^{-(i-1)(j-1)}w_j\right)^* = \gamma^{-(i-1)((\kappa-j+2)-1)}w_{\kappa-j+2}, \tag{B.85}$$

for $j = 2, ..., \kappa$, and hence if κ is odd we have

$$c_{x^{i-1}} = \frac{1}{\kappa} \sum_{j=1}^{\kappa} \gamma^{-(i-1)(j-1)} w_j \tag{B.86}$$

$$= \frac{1}{\kappa} \left[w_1 + \sum_{j=2}^{\frac{\kappa+1}{2}} \left(\gamma^{-(i-1)(j-1)} w_j + \gamma^{-(i-1)((\kappa-j+2)-1)} w_{\kappa-j+2} \right) \right] \tag{B.87}$$

$$= \frac{1}{\kappa} \left[-\frac{1}{m} + \sum_{j=2}^{\frac{\kappa+1}{2}} 2\beta \cos(\tilde{\theta}_j) \right]. \tag{B.88}$$

If κ is even, we note that our above condition (B.85) implies that $\left(\gamma^{-(i-1)\frac{\kappa}{2}} w_{\frac{\kappa}{2}+1} \right)^* = \gamma^{-(i-1)\frac{\kappa}{2}} w_{\frac{\kappa}{2}+1}$, so we must have that $\gamma^{-(i-1)\frac{\kappa}{2}} w_{\frac{\kappa}{2}+1}$ is real, and hence equal to $\pm\beta$. Thus we get

$$c_{x^{i-1}} = \frac{1}{\kappa} \left[w_1 + \gamma^{-(i-1)\frac{\kappa}{2}} w_{\frac{\kappa}{2}+1} + \sum_{j=2}^{\frac{\kappa}{2}} \left(\gamma^{-(i-1)(j-1)} w_j + \gamma^{-(i-1)((\kappa-j+2)-1)} w_{\kappa-j+2} \right) \right] \tag{B.89}$$

$$= \frac{1}{\kappa} \left[-\frac{1}{m} \pm \beta + \sum_{j=2}^{\frac{\kappa}{2}} 2\beta \cos(\tilde{\theta}_j) \right]. \tag{B.90}$$

In either case, maximizing over $\{\theta_j\}$ gives us an upper bound of

$$\mu \leq \frac{1}{\kappa} \left(\frac{1}{m} + (\kappa - 1)\beta \right), \tag{B.91}$$

which matches with our bound from Theorem 29.

Bibliography

[1] W. U. Bajwa, R. Calderbank, and S. Jafarpour. Why Gabor frames? two fundamental measures of coherence and their role in model selection. *J. Commun. Netw.*, 12:289–307, 2010.

[2] S. Bandyopadhyay, P. O. Boykin, V. Roychowdhury, and F. Vatan. A new proof of the existence of mutually unbiased bases. *Algorithmica*, 34:512–528, 2002.

[3] L. D. Baumert. Cyclic Difference Sets, Lecture Notes in Mathematics. 182, 1971.

[4] J. Benedetto and M. Fickus. Finite normalized tight frames. *Advances in Computational Mathematics*, 18(2-4):357–385, 2003.

[5] E. R. Berlekamp. Nonbinary BCH decoding. *ISIT*, 1967.

[6] E. R. Berlekamp. The weight enumerators for certain subcodes of the second order binary Reed-Muller codes. *Inform. Control*, 17:485–500, 1970.

[7] E. R. Berlekamp and L. Welch. Error correction of algebraic block codes, US patent, 1986.

[8] T. Beth, D. Jungnickel, and H. Lenz. *Design Theory.* Cambridge University Press, Cambridge, U.K., 1999.

[9] L. Bos and S. Waldron. Some remarks on Heisenberg frames and sets of equiangular lines. *N. Z. Jour. Math.*, 36:113–137, 2007.

[10] E. J. Candes. The restricted isometry property and its implications in compressed sensing. *C. R. Acad. Sci. Paris S'er. I Math.*, 346:589–592, 2008.

[11] E. J. Candes, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59:1208–1223, 2006.

[12] E. J. Candes and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Th.*, 51:4203–4215, 2005.

[13] P. G. Casazza. The art of frame theory. *Taiwanese Journal of Mathematics*, 4(2):129–201, June 2000.

[14] P. G. Casazza, M. Fickus, D. G. Mixon, Y. Wang, and Z. Zhou. Constructing tight fusion frames. *Applied and Computational Harmonic Analysis*, 30(2):175–187, 2011.

[15] P. G. Casazza and J. Kovačević. Equal-norm tight frames with erasures. *Advances in Computational Mathematics*, 2003.

[16] P. G. Casazza and G. Kutyniok. Frames of subspaces. *Contemporary Mathematics*, 345:87–114, 2004.

[17] P. G. Casazza and G. Kutyniok. *Finite Frames: Theory and Applications. Chapter 5 (written by S. Waldron)*. Springer, 2012.

[18] P. G. Casazza, G. Kutyniok, and S. Li. Fusion frames and distributed processing. *Applied and Computational Harmonic Analysis*, 25(1):114–132, July 2008.

[19] T. Chan, A. Grant, and D. Pflüger. Truncation technique for characterizing linear polymatroids. *IEEE Trans. Inform. Theory*, 57(10):6364–6378, October 2011.

[20] T. H. Chan. On the optimality of group network codes. *Proc. of the 2005 IEEE ISIT*, pages 1992–1996, September 2005.

[21] T. H. Chan. Capacity regions for linear and abelian network codes. *Proc. of the 2007 ITA Workshop*, pages 73–78, January and February 2007.

[22] T. H. Chan. Group characterizable entropy functions. *Proc. of the 2007 IEEE ISIT*, pages 506–510, June 2007.

[23] T. H. Chan and R. W. Yeung. On a relation between information inequalities and group theory. *IEEE Trans. Inform. Theory*, 48(7):1992–1995, July 2002.

[24] Minghua Chen, Cheng Huang, and Jin Li. On the Maximally Recoverable Property for Multi-Protection Group Codes. In *2007 IEEE Int. Symp. Inf. Theory*, pages 486–490. IEEE, June 2007.

[25] T. Chien and S. Waldron. A classification of the harmonic frames up to unitary equivalence. *Appl. Comput. Harmon. Anal.*, 30:307–318, 2011.

[26] D. Chu. Polyphase codes with good periodic correlation properties (Corresp.). *IEEE Trans. Inform. Theory*, 18(4):531–532, July 1972.

[27] J. H. Conway, R. H. Harding, and N. J. A. Sloane. Packing lines, planes, etc.: Packings in Grassmannian spaces. *Exp. Math.*, 5(2), 1996.

[28] Son Hoang Dau, Wentu Song, Zheng Dong, and Chau Yuen. Balanced Sparsest generator matrices for MDS codes. In *Inf. Theory Proc. (ISIT), 2013 IEEE Int. Symp.*, pages 1889–1893, 2013.

[29] Son Hoang Dau, Wentu Song, and Chau Yuen. On Simple Multiple Access Networks. *IEEE J. Sel. Areas Commun.*, 8716(0733):1–1, 2014.

[30] Son Hoang Dau, Wentu Song, and Chau Yuen. On the existence of MDS codes over small fields with constrained generator matrices. In *Inf. Theory (ISIT), 2014 IEEE Int. Symp.*, pages 1787–1791, June 2014.

[31] P. Delsarte, J. M. Goethals, and J. J. Seidel. Spherical codes and designs. *Geometriae Dedicata*, 6(3):363–388, 1977.

[32] T. K. Dikaliotis, T. Ho, S. Jaggi, S. Vyetrenko, H. Yao, M. Effros, J. Kliewer, and E. Erez. Multiple-access network information flow and correction codes. *IEEE Trans. Inf. Theory*, 57(2):1067–1079, February 2011.

[33] Theodoros K. Dikaliotis, Tracey Ho, Sidharth Jaggi, Svitlana Vyetrenko, Hongyi Yao, Michelle Effros, Jörg Kliewer, and Elona Erez. Multiple-Access Network Information-Flow and Correction Codes. *IEEE Trans. Inf. Theory*, 57(2):1067–1079, February 2011.

[34] C. Ding. Complex codebooks from combinatorial designs. *IEEE Trans. Inf. Theory*, 52(9):4229–4235, September 2006.

[35] C. Ding and T. Feng. A generic construction of complex codebooks meeting the Welch bound. *IEEE Trans. Inf. Theory*, 53, 2007.

[36] C. Ding, M. Golin, and T. Klø ve. Meeting the Welch and Karystinos-Pados bounds on DS-CDMA binary signature sets. *Designs, Codes, Cryptogr.*, 30:73–84, 2003.

[37] D. L. Donoho and M. Elad. Optimally sparse representations in general (non-orthogonal) dictionaries via ℓ_1 minimization. *Proc. Nat. Acad. Sci.*, 100:2197–2202, 2002.

[38] D. L. Donoho and X. Huo. Uncertainty principles and ideal atomic decompositions. *IEEE Trans. Inform. Theory*, 47:2845–2862, 2001.

[39] R. Dougherty. Computations of linear rank inequalities on six variables. *Proc. of the 2014 IEEE ISIT*, pages 2819–2823, June 2014.

[40] R. Dougherty, C. Freiling, and K. Zeger. Insufficiency of linear network coding in network information flow. *IEEE Trans. on Inform. Theory*, pages 2745–2759, 2005.

[41] R. Dougherty, C. Freiling, and K. Zeger. Linear rank inequalities on five or more variables (preprint), 2010.

[42] R. Dougherty, C. Freiling, and K. Zeger. Characteristic-dependent linear rank inequalities and network coding applications. *Proc. of the 2014 IEEE ISIT*, pages 101–105, June 2014.

[43] T. C. Eldar and G. D. Forney, Jr. Optimal tight frames and quantum measurement. *IEEE Trans. Inform. Theory*, 48:599–610, 2002.

[44] J. L. Fan. Array codes as Low-Density Parity Check codes. *Proc. Int'l. Symp. on Turbo Codes*, pages 543–546, September 2000.

[45] M. Fickus, D. G. Mixon, and J. C. Tremain. Steiner equiangular tight frames. *Linear Algebra and its Applications*, 436(5):1014–1027, 2012.

[46] J. P. Gabardo and D. Han. Frame representations for group-like unitary operator systems. *Journal of Operator Theory*, 49:223–244, 2003.

[47] S. W. Golomb. Cyclic Hadamard difference sets—constructions and applications. *Sequences and their Applications*, pages 39–48, 1999.

[48] Parikshit Gopalan, Cheng Huang, Huseyin Simitci, and Sergey Yekhanin. On the Locality of Codeword Symbols. *IEEE Trans. Inf. Theory*, 58(11):6925–6934, November 2012.

[49] W. Halbawi, M. Thill, and B. Hassibi. Coding with constraints: Minimum distance bounds and systematic constructions. *Proc. of ISIT*, 2015.

[50] Wael Halbawi, Tracey Ho, and Iwan Duursma. Distributed gabidulin codes for multiple-source network error correction. In *2014 Int. Symp. Netw. Coding*, pages 1–6. IEEE, June 2014.

[51] Wael Halbawi, Tracey Ho, Hongyi Yao, and Iwan Duursma. Distributed reed-solomon codes for simple multiple access networks. In *2014 IEEE Int. Symp. Inf. Theory*, pages 651–655. IEEE, June 2014.

[52] P. Hall. On representatives of subsets. *J. London Math. Soc.*, 10(1):26–30, 1935.

[53] D. Hammer, A. Romashchenko, A. Shen, and N. Vereshchagin. Inequalities for Shannon entropy and Kolmogorov complexity. 60(2):442–464, April 2000.

[54] Junsheng Han and Luis Alfonso Lastras-Montano. Reliable Memories with Subline Accesses. In *2007 IEEE Int. Symp. Inf. Theory*, pages 2531–2535. IEEE, June 2007.

[55] N. Hay and S. Waldron. On computing all harmonic frames of n vectors in $\mathbb{C}^d$. *Appl. Comput. Harmon. Anal.*, 21:168–181, 2006.

[56] R. W. Heath, Jr., H. Bolcskei, and A. J. Paulraj. Space-time signaling and frame theory. *IEEE Proceedings of ICASSP*, 4:2445–2448, 2001.

[57] R. W. Heath, Jr. and A. J. Paulraj. Linear dispersion codes for MIMO systems based on frame theory. *IEEE Trans. on Sig. Proc.*, 50(10):2429–2441, 2002.

[58] Cheng Huang, Minghua Chen, and Jin Li. Pyramid Codes: Flexible Schemes to Trade Space for Access Efficiency in Reliable Data Storage Systems. In *Sixth IEEE Int. Symp. Netw. Comput. Appl. (NCA 2007)*, pages 79–86. IEEE, July 2007.

[59] A. Ingleton. Representation of matroids. *Combinatorial Mathematics and its Applications*, pages 149–167, 1971.

[60] F. Kahn. *LTE for 4G Mobile Broadband.* Cambridge University Press, New York, 2009.

[61] D. Kalra. Complex equiangular cyclic frames and erasures. *Linear Algebra Appl.*, 419:373–399, 2006.

[62] G. M. Kamath, N. Prakash, V. Lalitha, and P. V. Kumar. Codes with local regeneration. In *2013 Inf. Theory Appl. Work.*, pages 1–5. IEEE, February 2013.

[63] M. Khatirinejad. On Weyl-Heisenberg orbits of equiangular lines. *J. Algebr. Comb.*, 28:333–349.

[64] R. Kinser. New inequalities for subspace arrangements. *Journal of Combinatorial Theory*, 118(1):152–161, January 2011.

[65] A. Klappenecher and M. Rötteler. Constructions of mutually unbiased bases. *Finite Fields and Applications, 7th International Conference, Fq7*, May 2003.

[66] H. Li and E. K. P. Chong. On connections between group homomorphisms and the Ingleton inequality. *Proc. of the 2007 IEEE ISIT*, pages 1996–2000, June 2007.

[67] J. H. Lint and R. M. Wilson. *A Course in Combinatorics.* Cambridge University Press, Cambridge, 1992.

[68] W. Mao and B. Hassibi. Violating the Ingleton inequality with finite groups. *Proc. of Allerton*, 2009.

[69] W. Mao, M. Thill, and B. Hassibi. On group network codes: Ingleton bound violations and independent sources. *Proc. of IEEE ISIT*, pages 2388–2392, 2010.

[70] W. Mao, M. Thill, and B. Hassibi. On the Ingleton violations in finite groups, November 2014.

[71] J. L. Massey. Shift-register synthesis and BCH decoding. *IEEE Trans. Inform. Theory*, 15(1):122–127, 1969.

[72] F. Matúš. Conditional independences among four random variables III: Final conclusion. *Combinatorics, Probability and Computing*, 8:269–276, 1999.

[73] F. Matúš. Infinitely many information inequalities. *IEEE Proceedings of ISIT*, 2007.

[74] Arya Mazumdar. Storage Capacity of Repairable Networks. *arXiv:1408.4862*, August 2014.

[75] D. G. Mixon, W. U. Bajwa, and R. Calderbank. Frame coherence and sparse signal processing. *Proceedings of ISIT*, 2011.

[76] Dimitris S. Papailiopoulos and Alexandros G. Dimakis. Locally repairable codes. In *2012 IEEE Int. Symp. Inf. Theory Proc.*, pages 2771–2775. IEEE, July 2012.

[77] W. W. Peterson. Encoding and error correction procedures for the Bose-Chaudhuri codes. *IRE Transactions on Information Theory*, IT-6:459–470, 1960.

[78] N. Prakash, Govinda M. Kamath, V. Lalitha, and P. Vijay Kumar. Optimal linear codes with a local-error-correction property. In *2012 IEEE Int. Symp. Inf. Theory Proc.*, pages 2776–2780. IEEE, July 2012.

[79] A. Prasad. Representations of $GL_2(\mathbb{F}_q)$ and $SL_2(\mathbb{F}_q)$, and some remarks about $GL_n(\mathbb{F}_q)$. *Lecture Notes, Advanced Instructional School on Representation Theory and Related Topics held at the Bhaskaracharya Pratishthana and the University of Pune*, July 2007.

[80] H. Rauhut. Compressive sensing and structured random matrices.

[81] Ankit Singh Rawat, O. Ozan Koyluoglu, Natalia Silberstein, and Sriram Vishwanath. Optimal Locally Repairable and Secure Codes for Distributed Storage Systems. October 2012.

[82] I. Reed and G. Solomon. Polynomial codes over certain finite fields. *J. Soc. Ind. Appl. Math.*, 1960.

[83] M. Reeder. Characters of $SL_2(q)$. *https://www2.bc.edu/ reederma/SL(2,q).pdf?*, December 2008.

[84] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves. Symmetric informationally complete quantum measurements. *J. Math. Phys*, 45:2171–2180, 2004.

[85] A. J. Scott. Tight informationally complete quantum measurements. *J. Phys. A: Math. Gen.*, 39(43):13507, 2006.

[86] A. J. Scott and M. Grassl. SIC-POVMs: A new computer study. *arXiv:0910.5784v2 [quant-ph]*, 2009.

[87] J. P. Serre and L. L. Scott. *Linear Representations of Finite Groups.* Graduate Texts in Mathematics, Springer, New York, 1977.

[88] A. Shokrollahi, B. Hassibi, B. M. Hochwald, and W. Sweldens. Representation theory for high-rate multiple-antenna code design. *IEEE Transactions on Information Theory*, 47(6):2335–2367, September 2001.

[89] R. C. Singleton. Maximum distance q-nary codes. *IEEE Trans. Inform. Theory*, 10(2):116–118, 1964.

[90] D. Slepian. Group codes for the Gaussian channel. *Bell Sys. Tech. J.*, 47:575–602, April 1968.

[91] T. Strohmer and R. W. Heath Jr. Grassmannian frames with applications to coding and communication. *Appl. Comput. Harmon. Anal.*, 14:257–275, 2003.

[92] Itzhak Tamo, Dimitris S. Papailiopoulos, and Alexandros G. Dimakis. Optimal locally repairable codes and connections to matroid theory. In *2013 IEEE Int. Symp. Inf. Theory*, pages 1814–1818. IEEE, July 2013.

[93] S. Tanaka. Construction and classification of irreducible representations of special linear group of the second order over a finite field. *Osaka J. Math*, 4:65–84, 1967.

[94] A. Terras. *Fourier Analysis on Finite Groups and Applications.* Cambridge University Press, Cambridge, 1999.

[95] M. Thill and B. Hassibi. Low-coherence frames from group Fourier matrices. *In Preparation.*

[96] M. Thill and B. Hassibi. Frames, Group Codes, and Subgroups of $(\mathbb{Z}/p\mathbb{Z})^{\times}$. *Proc. of Allerton*, 2012.

[97] M. Thill and B. Hassibi. Frames from groups: Generalized bounds and dihedral groups. *Proc. of ICASSP*, 2013.

[98] M. Thill and B. Hassibi. On frames from abelian group codes. *Proceedings of ISIT*, 2013.

[99] M. Thill and B. Hassibi. Group frames with few distinct inner products and low coherence. *IEEE Transactions on Signal Processing*, July 2015.

[100] M. Thill, V. Muthakumar, and B. Hassibi. Frames from generalized group Fourier transforms and $SL_2(\mathbb{F}_q)$. *Proc. of ICASSP*, 2014.

[101] J. Tropp and A. Gilbert. Signal recovery from partial information via orthogonal matching pursuit. *IEEE Trans. Inform. Theory*, 53(12):4655–4666, 2007.

[102] R. Vale and S. Waldron. Tight frames and their symmetries. *Constr. Approx.*, 21:83–112, 2005.

[103] R. Vale and S. Waldron. Tight frames generated by finite nonabelian groups. *Numer. Algorithms*, 48:11–27, 2008.

[104] R. Vale and S. Waldron. The symmetry group of a finite frame. *Linear Algebra Appl.*, 433:248–262, 2010.

[105] B. V. Vasic and O. Milenkovic. Combinatorial constructions of low-density parity-check codes for iterative decoding. *IEEE Trans. Inform. Theory*, 50(6):1156–1176, 2004.

[106] S. Waldron. *An Introduction to Finite Tight Frames.* Springer, New York, 2011.

[107] W. D. Wallis. *Combinatorial Designs.* Marcel Dekker, New York, 1988.

[108] L. Welch. Lower bounds on the maximum cross correlation of signals. *IEEE Trans. Inform. Theory*, 20(3):397–399, May 1974.

[109] W. K. Wootters and B. D. Fields. Optimal state-determination by mutually unbiased measurements. *Ann. Physics*, 191:363–381, 1989.

[110] P. Xia, S. Zhou, and G. B. Giannakis. Achieving the Welch bound with difference sets. *IEEE Trans. Inform. Theory*, 51:1900–1907, 2005.

[111] Muxi Yan and A Sprintson. Weakly Secure Network Coding for Wireless Cooperative Data Exchange. In *Glob. Telecommun. Conf. (GLOBECOM 2011), 2011 IEEE*, pages 1–5, December 2011.

[112] Muxi Yan, Alex Sprintson, and Igor Zelenko. Weakly Secure Data Exchange with Generalized Reed Solomon Codes. pages 1366–1370, 2014.

[113] N. Y. Yu and G. Gong. A new binary sequence family with low correlation and large size. *IEEE Trans. Inform. Theory*, 52:1624–1636, 2006.

[114] Z. Zhang and R. W. Yeung. A non-Shannon-type conditional inequality of information quantities. *IEEE Trans. Inf. Theory*, 43(6):1982–1986, November 1997.