Real equiangular lines and related codes

by

Sam Winnick

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2019

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

We consider real equiangular lines and related codes. The driving question is to find the maximum number of equiangular lines in a given dimension. In the real case, this is controlled by combinatorial phenomena, and until only very recently, the exact number has been unknown. The complex case appears to be driven by other phenomena, and configurations are conjectured always to meet the absolute bound of d^2 lines in dimension d. We consider a variety of the techniques that have been used to approach the problem, both for constructing large sets of equiangular lines, and for finding tighter upper bounds. Many of the best-known upper bounds for codes are instances of a general linear programming bound, which we discuss in detail. At various points throughout the thesis, we note applications in quantum information theory.

Acknowledgements

I would like to thank my supervisor Jon Yard for his guidance and support over these last two years and for spending countless hours discussing research with me. I am very lucky to have the opportunity to learn from him. I am very grateful to my thesis readers Jon Yard, Chris Godsil, and David Wagner for taking their time to read this thesis. I also want to acknowledge the support of the Institute for Quantum Computing and the Department of Combinatorics and Optimization. Finally, I would like to thank my wife Yamin Zhou for her love and support.

Table of Contents

0.1 Notations and conventions

We make frequent use of the following notations.

Notation	Meaning
K	Either $\mathbb{R}$ or $\mathbb{C}$
$N(X)$	The maximum cardinality of a simplex in a code space X
$N(X,A)$	The maximum cardinality of an A-code in a code space X
$H_{K,d}, H_d$	The real vector space of $d \times d$ Hermitian resp. symmetric matrices for $K = \mathbb{R}$ resp. $\mathbb{C}$
$v^{\dagger}$	The conjugate transpose of a vector $v \in K^d$
v^*	The dual of a vector v
$[v]$	(for a vector v in a K-vector space): $\mathrm{span}_K(v)$, the line spanned by v
$\mathbb{N}$	Either $\{0, 1, 2, \cdots\}$ or $\{1, 2, 3, \cdots\}$, depending on context.
$[n]$	(for a natural number $n \in \mathbb{N}$): $\{1, \cdots, n\}$

Thoughout the thesis, we use the standard “math conventions” as opposed to “physics conventions” for sesquilinear forms and related things, though we find some of the math conventions to be pairwise incompatible. These conventions are listed below. Asterisks denote standard conventions that nevertheless clash with the others. Sorting out these details is obligatory when doing careful proofs in representation theory, cf. Theorem (4.4.2).

Math convention	Physics convention
$\langle u, v\rangle$ is linear in the left, antilinear in the right	opposite
*Linear functional ϕ written to the left of vector v: $\phi(v)$	same
*Usual matrix notation	same
(should be transposed; multiplication unchanged)	
*Dagger and star are written as right-superscripts $v^{\dagger}$, v^*	same
Hence, confusingly $\langle u, v\rangle = v^*(u) = v^{\dagger}\ u$	
It would be nicer to have $\langle u, v\rangle = (u)^* v = u\ ^{\dagger}v$.	
Identify: $V \otimes V^* = \mathrm{End}(V)$	opposite - $V^* \otimes V$
Peter-Weyl Theorem involves $V \otimes V^*$	opposite - $V^* \otimes V$
*Writing functionals on right, we have $v \otimes \phi \mapsto (v)\phi$.	
Homogeneous spaces have the form G/H w/left G-action	should be $H\backslash G$; right
Reproducing kernels $k_{(x)}$ in right factor: $\langle f, k_{(x)}\rangle = f(x)$	opposite - $\langle k_{(x)}, f\rangle = f(x)$
*Writing functionals on right, have $\langle f, k_{(x)}\rangle = (x)f$	

Throughout the thesis we only use left representations and left group actions.

Chapter 1

Introduction

A set of lines through the origin of $\mathbb{R}^d$ or $\mathbb{C}^d$ is called *equiangular* if each pair of lines meets at the same angle. Real equiangular lines were first investigated in 1948 by Haantjes, who proved that the maximum number of equiangular lines in $\mathbb{R}^3$ is 6, and in $\mathbb{R}^4$ is again 6 [Haa48]. In $\mathbb{R}^3$ the unique such configuration of lines arises as the 6 lines through the 12 vertices, which come in antipodal pairs, of a regular icosahedron centered at the origin. While the 4th dimension does not allow for this configuration to be extended to 7 or more lines, the extra space allows for an entirely different configuration of 6 equiangular lines, which is the smallest real equiangular configuration to exhibit a non-uniform matroid structure.

Quadratic upper bounds for the maximum number of equiangular lines in $\mathbb{R}^d$ and $\mathbb{C}^d$ were discovered by Gerzon, but in the real case this bound is hardly ever tight, and only known to be in dimensions 2,3,7, and 23. In fact, Gerzon's bound must not be tight in any dimension other than 2, 3, or two less than an odd square. The nonexistence of configurations of real equiangular lines attaining Gerzon's bound is principally a question of combinatorics, as configurations of equiangular vectors correspond to graphs; however, it can sometimes also be viewed as a consequence of other phenomena. For example, it turns out that if the number of lines exceeds twice the dimension, then the cosine of the angle is the reciprocal of an odd integer. This fact can be used in conjunction with a so-called *relative bound* to deduce a sharp upper bound for the maximum number of equiangular lines in real dimensions 2-7.

Mysteriously, the maximum number of real equiangular lines in dimensions 7-13 is 28, and in dimension 14, whether there are 28 or 29 equiangular lines remains an open problem. Lemmens and Seidel proposed the method of “pillars”, which purely uses algebraic combinatorics, and found the exact maximum number of equiangular lines in low dimensions [LS73]. Their method can be summed up as “projecting onto a coclique”. In 2016, Balla, Draxler, Keevash, and Sudakov made a breakthrough by projecting onto a *clique* to derive upper bounds on the maximum number of equiangular lines at a fixed angle that are linear in the dimension and do not depend on the angle [BDKS16]. This followed the surprising result of Bukh, who was the first to prove an upper bound that was linear in the dimension, but for a fixed angle [Buk15]. Recently, Jiang, Tidor, Yao, Zhang, and Zhao claimed to have solved the problem [JTY$^+$19]. Their new lower bound holds only for sufficiently large dimensions, and does not resolve questions like the maximum number of equiangular lines in real dimension 14.

Complex equiangular lines behave entirely differently than the real ones. In the complex case, the combinatorial tools analogous to those used in the real case become too complicated, and findings suggest that the combinatorial constraints of the real case do not carry over. Additionally, it was conjectured in the Ph.D. thesis of Zauner that Gerzon’s bound is attained in all complex dimensions [Zau11]. Thus, the emphasis is placed on constructions, especially the promising Heisenberg covariant symmetric informationally complete positive operator-valued measures (Heisenberg SICs), also conjectured by Zauner. Such SICs are generated by a single *fiducial* vector. The coordinates of this vector become extremely complicated in dimensions higher than 3, and the main focus of research in the complex case has been in understanding the underlying number theory of the relevant number fields, which are, apparently, Ray class fields [AFMY17]. We do not venture down that road in this thesis, but we provide the easiest examples of Heisenberg SICs, which are those in dimensions 2 and 3.

Throughout our exposition, we note many connections with quantum information theory. A class of structures that has some similarity with SICs is that of mutually unbiased bases (MUBs), another example we return to throughout the thesis.

Lastly, we explore the linear programming bounds first introduced in [DGS77]. The linear programming bounds give the best known bounds for many codes, and they have been applied in many areas. In Chapter 4, we consider a generalized framework in which the bounds hold and compute some examples.

Chapter 2

Equiangular lines

In this chapter, we discuss the theory of equiangular lines. The concept can be approached using methods from diverse fields of math. The geometric and linear algebraic approaches described in the first two sections is intended to equip the reader with the necessary tools to jump to any other section. A lot could have been included in this chapter, and one method that will not be discussed but deserves mention is an analytical approach introduced by Cohn, Kumar, and Minton in 2016 [CKM16]. They show how one can use a suitable set of equations encoding equiangularity or the tight frame property and an approximate solution to prove the existence of a nearby exact solution. This technique so far has not been fruitful with in the complex case, but it has been used in other settings to find codes, such as quaternionic projective space and the octonionic plane.

2.1 Geometric preliminaries

In this section we introduce the notion of equiangular lines as equidistant points in projective space. Projective space embeds into a submanifold of a hypersphere in the space of Hermitian resp. symmetric (for $K = \mathbb{R}$ resp. $\mathbb{C}$) matrices. From this point of view, the difficult questions about equiangular lines can be thought of as being caused by the complicated geometry of projective space arising from a rank condition. Really, this is an oversimplification, and the way to understand these complications seems to either be through combinatorics (especially in the real case) or algebraic geometry and number theory (especially in the complex case). Many important facts about equiangular lines present themselves more naturally through the language of linear algebra rather than geometry (see Section 2.2), so we omit some details in this section and prove them in the next section.

Nevertheless, approaching the topic geometrically helps build intuition. We end this section and sections to follow with some remarks on the connections to quantum information.

Throughout this thesis, K will always stand for $\mathbb{R}$ or $\mathbb{C}$.

Definition 2.1.1. For a unit vector v in K^d with $K = \mathbb{R}$ or $\mathbb{C}$ we denote $[v] := \operatorname{span}_K\{v\}$. For our purposes, a *line* in K^d is a linear subspace $l = [v]$ (for some unit vector $v \in K^d$).

By definition we only consider lines through the origin.

Definition 2.1.2. The *angle* between two unit vectors $v, v' \in K^d$ is $|\langle v, v' \rangle|$, where $\langle \cdot, \cdot \rangle$ denotes the dot product resp. Hermitian inner product for $K = \mathbb{R}$ resp. $\mathbb{C}$, and $\dagger$ denotes the transpose resp. conjugate transpose. The *angle* between two lines $[v], [v']$ is $|\langle v, v' \rangle|$.

The angle of two lines is well-defined. Really, $\langle v, v' \rangle$ is the cosine of the angle, in the usual usage of the word *angle*, but the definition given above is standard convention in this field of research. Beware that the cosine function decreases on $(0, \pi)$.

Definition 2.1.3. A set (or sequence) of unit vectors or lines is *equiangular* if there exists a real number $\alpha > 0$ such that the angle between any distinct two is α.

We do not consider the degenerate case $\alpha = 0$. Thus $\alpha \in (0, 1)$. Though equiangularity is a property the vectors or lines collectively satisfy, it is common to refer to a *set of equiangular lines* rather than an *equiangular set of lines*, as the language feels more natural.

Definition 2.1.4. Projective $d-1$ space $K\mathbb{P}^{d-1}$ is the quotient $(K^d \setminus \{0\})/K^\times$, that is, the set of lines in K^d. For a fixed hyperplane $P \subseteq K^{d-1}$ that does not contain the origin, each line $l \in K\mathbb{P}^{d-1}$ not parallel to P intersects P at a unique point. By specifying a fixed point $O \in P$, the lines $l \in K\mathbb{P}^{d-1}$ parallel to P can be shifted to a unique line in P through O, and therefore correspond to elements of $K\mathbb{P}^{d-2}$. In equations one might express this as

$$K\mathbb{P}^{d-1} \cong K^{d-1} \sqcup K\mathbb{P}^{d-2} \qquad \text{or recursively} \qquad K\mathbb{P}^{d-1} \cong K^{d-1} \sqcup K^{d-2} \sqcup \cdots \sqcup K^0$$

In other words $K\mathbb{P}^{d-1}$ looks like K^{d-1} with a projective hyperplane at infinity.

$\mathbb{RP}^1 \cong S^1$ (left) and $\mathbb{RP}^2$ (right). $\mathbb{RP}^2$ is isomorphic to the closed disc with antipodal points of the enclosing circle identified. $\mathbb{RP}^3$ is a closed ball with antipodal points of the enclosing sphere identified. $\mathbb{CP}^1 \cong S^2$.

Projective $d-1$ space can be made into a metric space as follows.

Definition 2.1.5. Let $H_{K,d}$ (or H_d for short) denote the real vector space of symmetric resp. hermitian $d \times d$ matrices for $K = \mathbb{R}$ resp. $\mathbb{C}$, which has dimension $\binom{d+1}{2}$ resp. d^2. H_d is a subspace of the vector space $K^{d\times d}$ of all $d \times d$ matrices over K, which has the inner product

$$\langle A, B\rangle = \operatorname{tr}\bigl(AB^\dagger\bigr) \tag{2.1}$$

H_d inherits this inner product, which induces a norm, which induces a metric. The norm induced by (2.1) is the entrywise 2-norm, and it is the only matrix norm used in this thesis. The distance between two hermitian matrices A and B is thus $\sqrt{\operatorname{tr}((A-B)^2)}$.

Proposition 2.1.1. The elements of $K\mathbb{P}^{d-1}$ correspond uniquely to matrices of the form $vv^\dagger$ where v is a unit vector.

(By the spectral theorem, such matrices are precisely the matrices of rank 1 and trace 1.)

Proof. To see that the mapping $[v] \mapsto vv^\dagger$ is well-defined, note that if $[v] = [v']$, then $v' = cv$ for some $c \in K$ of absolute value 1, and then $v'v'^\dagger = cvv^\dagger\bar{c} = vv^\dagger$. On the other hand, the mapping is invertible because $[v]$ is recovered as the column space of $vv^\dagger$. ■

Thus, we have an injection

$$\gamma : K\mathbb{P}^{d-1} \hookrightarrow H_d \tag{2.2}$$

With *distance* between $l_1, l_2 \in K\mathbb{P}^{d-1}$ defined as $d(\gamma(l_1), \gamma(l_2))$, $K\mathbb{P}^{d-1}$ is a metric space.

Let us note some basic geometrical properties of $\operatorname{im}\gamma$.

Proposition 2.1.2. $K\mathbb{P}^{d-1}$ is contained in the d^2-2 dimensional hypersphere of radius $\sqrt{1-1/d}$ centered at $\frac{1}{d}I$ and contained in the trace 1 hyperplane of H_d.

Proof. Let $[v] \in k\mathbb{P}^{d-1}$. The corresponding element of H_d is $\rho = vv^\dagger$.

$$\operatorname{dist}\left(\rho\,,\,\frac{1}{d}I\right) = \sqrt{\operatorname{tr}(\rho^2) - \frac{2}{d}\operatorname{tr}(\rho) + \frac{1}{d^2}\operatorname{tr}(I)} = \sqrt{\operatorname{tr}(\rho^2) - \frac{1}{d}} = \sqrt{1 - \frac{1}{d}} \tag{2.3}$$

■

Corollary 2.1.0.1. $\mathbb{CP}^1$ comprises the entire sphere centered at $\frac{1}{2}I$ in the trace 1 3-plane of $H_{\mathbb{C},2}$. $\mathbb{RP}^1$ comprises the entire circle centered at $\frac{1}{2}I$ in the trace 1 plane of $H_{\mathbb{R},2}$. In higher dimensions $d > 2$, $K\mathbb{P}^{d-1}$ "embeds" into a proper subset of the d^2-2 dimensional hypersphere centered at $\frac{1}{d}I$ in the trace 1 hyperplane of $H_{K,d}$.

Proof sketch. $\mathbb{RP}^1$ and $\mathbb{CP}^1$ are the one-point compactifications of $\mathbb{R}^1$ and $\mathbb{C}^1$. γ is a homeomorphism onto its image, so $\operatorname{im}\gamma$ is compact, nonempty, and open, and therefore comprises all of S^1 or S^2 respectively. In higher dimensions, $\dim K\mathbb{P}^{d-1} = 2(d-1) < d^2 - 2 = \dim S^{d^2-2}$, but $\operatorname{im}\gamma$ should have dimension equal to $\dim K\mathbb{P}^{d-1}$. □

There are various frameworks in which notions of *dimension* and *embedding* are defined, embeddings preserve dimension, and the above argument can be made rigorous by noting that γ is indeed an embedding. For example, $K\mathbb{P}^{d-1}$ and S^{d^2-2} are differentiable manifolds and algebraic varieties of the specified dimensions.

The distance between lines at angle α is $\sqrt{2-2\alpha^2}$, so finding equiangular lines in K^d is the same thing as finding a regular simplex (*i.e.* a set of equidistant points) in $K\mathbb{P}^{d-1}$.

Proposition 2.1.3. The largest regular simplex in $\mathbb{R}^n$ contains $n+1$ points.

The proposition is somewhat intuitive geometrically, but a rigorous proof uses machinery from the next section, so we include only a sketch here. It suffices to show the following.

1. Any two regular simplices of unit vectors in $\mathbb{R}^n$ with $n+1$ vertices differ only by an orthogonal matrix.
2. There is a regular simplex of unit vectors in $\mathbb{R}^n$ with $n+1$ vertices that spans $\mathbb{R}^n$.

Proof sketch (foreshadowing). A regular simplex of unit vectors in $\mathbb{R}^n$ is an equiangular set. If two such simplices have the same angle, then they have the same Gram matrix, so the first point above is an instance of Proposition (2.2.1) from the next section; the proof is omitted here. We will also later see that since the number of vectors exceeds the dimension and the vectors are equiangular, the angle α is encoded in the yet-to-be defined Seidel matrix, so they must have the same angle. Alternately, we can use our geometric intuition to determine α: consider two distinct vectors v_i and v_j of a regular simplex centered at the origin of $\mathbb{R}^n$. Then $\langle v_i, v_j \rangle = -\alpha$. Intuitively we expect that the centroid of the simplex is the origin, *i.e.* $\sum v_i = 0$, from which it follows that

$$0 = \left|\sum v_i\right|^2 = \sum |v_i|^2 + \sum_{i,j \ i \neq j} \langle v_i, v_j \rangle = (n+1) - \alpha n(n+1) \tag{2.4}$$

This implies $\alpha = 1/n$. As for the second point, this thesis includes three constructions of a regular simplex of $n+1$ vertices centered at the origin of $\mathbb{R}^n$, which span $\mathbb{R}^n$: first we have Example 2.1.1, and in the next section, we provide two methods for recovering equiangular vectors from their Gram matrix. □

Note that even if $\sum v_i$ is not necessarily 0 we obtain the inequality $\alpha \leq 1/n$. These smaller angles (larger in the usual sense of the word *angle*) are obtained by dragging the simplex into a higher dimension and renormalizing the vectors.

Example 2.1.1 (Simplex construction)**.** Let $e_1, \cdots, e_n$ be the standard unit vectors in $\mathbb{R}^n$. We can make this into a regular simplex by adding a vector $w = (a_1, \cdots, a_n)$ and requiring that for each $i = 1, \cdots, n$, $2 = |e_i - w|^2 = (a_i - 1)^2 + (n-1)a_i^2$. Thus a_i is a constant a with respect to i, and satisfies

$$na^2 - 2a - 1 \tag{2.5}$$

The centroid of the simplex is $c = (b, \cdots, b)$, where $b = (a+1)/(n+1)$. Now apply the shift to all the vectors that sends c to the origin to the origin and renormalize. The result is a regular simplex of $n+1$ unit vectors, which are equiangular with angle $-1/n$.

Proposition (2.1.3) leads to an important corollary: the absolute bound.

Corollary 2.1.0.2. The largest regular simplex in the trace 1 hyperplane of $H_{K,d}$ contains $\binom{d+1}{2}$ resp. d^2 elements for $K = \mathbb{R}$ resp. $\mathbb{C}$.

Proof. Choose an isometry $H_{K,d} \to \mathbb{R}^n$ for $n = \binom{d+1}{2}$ resp. d^2. ■

Corollary 2.1.0.3 (The absolute bound)**.** Let $N(K\mathbb{P}^{d-1})$ denote the maximum number of equiangular lines in K^d. Then

$$N(\mathbb{RP}^{d-1}) \leq \binom{d+1}{2} \qquad\qquad N(\mathbb{CP}^{d-1}) \leq d^2$$

We have inequality instead of equality because the vertices of the simplex are required to have rank 1 in addition to trace 1. This rank 1 property can be viewed as the source of a lot of difficulty in understanding equiangular lines.

Corollary 2.1.0.4. $N(\mathbb{RP}^1) = 3$ and $N(\mathbb{CP}^1) = 4$.

Proof. Inscribe a regular triangle in S^1 and a regular tetrahedron in S^2. ■

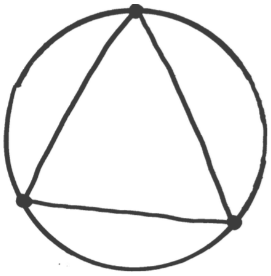

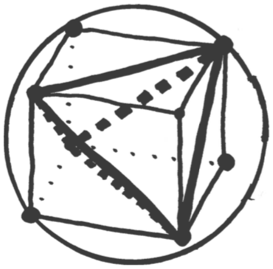

Remark 2.1.1 (Connection to quantum information theory)**.** In quantum information, qu*d*it pure states are represented by elements of $\mathbb{CP}^{d-1}$. The image of γ (defined in equation (2.2)) for $d = 2$ is a 2-sphere of radius $1/\sqrt{2}$. In this context, the renormalized sphere is known as the *Bloch sphere*. A pure state of a 2-qubit system (qu*d*it for $d = 4$) is a line through the origin of $\mathbb{C}^4$, not simply a pair $([v_1], [v_2]) \in S^2 \times S^2$ (for unit vectors v_1, v_2). Such pairs correspond to pure states of the form $[v_1 \otimes v_2]$, or equivalently, matrices of the form $v_1 v_1^\dagger \otimes v_2 v_2^\dagger$. For a bipartite system consisting of a qu$d_1$it and a qu$d_2$it, the function

$$\begin{aligned} \sigma : \mathbb{CP}^{d_1-1} \times \mathbb{CP}^{d_2-1} &\hookrightarrow \mathbb{CP}^{d_1 d_2 - 1} \\ ([v_1], [v_2]) &\mapsto [v_1 \otimes v_2] \end{aligned}$$

is a well-defined injection called the *Segre embedding*, and its image is a determinantal projective variety called the *Segre variety*. (Note that $v_1 v_1^\dagger$ and $v_2 v_2^\dagger$ can be recovered from $v_1 v_1^\dagger \otimes v_2 v_2^\dagger$ using the partial trace). As one might expect, the Segre variety has real dimension, $2d_1 + 2d_2 - 4 = \dim_{\mathbb{R}} \mathbb{CP}^{d_1-1} + \dim_{\mathbb{R}} \mathbb{CP}^{d_2-1} = \dim_{\mathbb{R}}(\mathbb{CP}^{d_1-1} \times \mathbb{CP}^{d_2-1})$. For example, the variety of separable pure states in a 2-qubit system has real dimension 4, whereas the variety of all pure states has real dimension 6, indicating that in general, not all pure states are separable. The rest are *entangled.*

Mixed states are represented by density matrices, *i.e.* convex combinations of rank 1 projection matrices, which in turn represent pure states. A mixed state on a bipartite system is *separable* if it has the form $\rho = \sum_i p_i \rho_i \otimes \rho'_i$ where ρ_i are density matrices acting on the first subsystem and ρ'_i are density matrices on the second. If it is possible to write ρ in this form with $i = 1$, then ρ is a *product state.* A probability distribution (or *ensemble*) of pure or mixed states gives rise to a mixed state by taking convex combinations. Somewhat conversely, by writing a mixed state as a convex combination, one can obtain a probability distribution on the set of pure states, though this distribution is not unique. Still, it turns out that two such distributions have the same measurement statistics.

When ρ is a mixed state, the left three expressions of (2.3) are still equal. The *purity* of ρ is defined as $\mathrm{tr}(\rho^2)$, which is roughly the distance from the *maximally mixed state* $\frac{1}{d}I$.

2.2 Linear algebraic approach

In the last section we gained some geometrical intuition about equiangular lines, but we skipped over a few important details when proving the absolute bound. It turns out that there is a succinct proof of the absolute bound using basic linear algebra. We also prove the relative bound, which, combined with a number theoretic result about the angle we prove in Chapter 3, gives us the exact number of real equiangular lines in dimensions 2-7.

In the previous section we considered *sets* of equiangular lines, whereas in this section we will consider *sequences.* (Our *sequences* will always be finite and have distinct terms.) There is no philosophical or combinatorial reason for doing this; our use of sequences is merely an artifact of linear algebra. The concepts discussed here apply just as well to sets of equiangular vectors or lines. For instance, if G is the Gram matrix of a *tight frame* (a certain kind of sequence of vectors), then so is $\Sigma G \Sigma^{-1}$ for any permutation matrix Σ.

This is not to say that the distinction of ordered and unordered objects is unimportant. It is just not relevant in this section. In the next chapter we will describe two very different notions of *switching class*, one for graphs on vertex set $[n]$ and one for isomorphism classes of graphs. For nearly every term we define in this section, there is a unordered version, whose definition we tacitly omit, e.g. *ordered* vs. *unordered configuration.*

If $v_1, \cdots, v_n$ are nonzero vectors, then $\Pi_1, \cdots, \Pi_n$ will always denote the corresponding rank 1 projectors. We identify a sequence $(v_1, \cdots, v_n)$ of vectors in K^d with the matrix $[v_1 \cdots v_n]$ whose columns are the coordinate vectors relative to the standard basis.

Definition 2.2.1. Let $L = (v_1, \cdots, v_n)$ be a sequence of vectors in an inner product space $(V, \langle \cdot, \cdot \rangle)$ over $\mathbb{R}$ or $\mathbb{C}$. The *Gram matrix* for L is the matrix with ijth entry $\langle v_i, v_j \rangle$.

Definition 2.2.2. Two sequences L_1, L_2 of vectors are orthogonally/unitarily equivalent if there exists an orthogonal/unitary matrix U such that $L_2 = UL_1$.

Proposition 2.2.1. Two sequences L_1, L_2 of vectors in K^d are orthogonally resp. unitarily equivalent (for $K = \mathbb{R}$ resp. $\mathbb{C}$) if and only if they have the same Gram matrix.

Proof. If $L_2 = UL_1$ for an orthogonal/unitary matrix U, then $\langle v_i^2, v_j^2 \rangle = \langle Uv_i^1, Uv_j^1 \rangle = \langle v_i^1, v_j^1 \rangle$. Conversely, suppose $L_1^\dagger L_1 = L_2^\dagger L_2$. Use the singular value decompositions $L_1 = U_1 \Sigma_1 V_1^\dagger$, $L_2 = U_2 \Sigma_2 V_2^\dagger$. Then by the existence and uniqueness of positive semi-definite square roots,

$$V_1 \Sigma_1 V_1^\dagger = \sqrt{V_1 \Sigma_1^2 V_1^\dagger} = \sqrt{V_1 \Sigma_1 U_1^\dagger U_1 \Sigma_1 V_1^\dagger} = \sqrt{L_1^\dagger L_1} = \sqrt{L_2^\dagger L_2} = \cdots = V_2 \Sigma_2 V_2^\dagger$$

Thus,

$$L_2 = U_2 \Sigma_2 V_2^\dagger = U_2 V_2^\dagger V_2 \Sigma_2 V_2^\dagger = U_2 V_2^\dagger V_1 \Sigma_1 V_1^\dagger = U_2 V_2^\dagger V_1 U_1^\dagger U_1 \Sigma_1 V_1^\dagger = UL_1$$

for the unitary $U = U_2 V_2^\dagger V_1 U_1^\dagger$. In the real case, the matrices U_1, U_2, V_1, V_2 are real, and so U is orthogonal. ■

Proposition 2.2.2. Let G be a symmetric resp. hermitian $n \times n$ matrix over K for $K = \mathbb{R}$ resp. $\mathbb{C}$ of rank d. Then G is positive semidefinite if and only if G is the Gram matrix of some sequence L of n vectors in K^d.

Proof. Since G is positive semi-definite, its eigenvalues are non-negative real numbers. Since G is normal, it is unitarily diagonalizable:

$$G = UDU^\dagger$$

Now, $n-d$ of the eigenvalues are 0. Let U' be the matrix obtained from U by deleting the ith row whenever the ith eigenvalue occuring in D is 0, and let D' be obtained from D by deleting the ith row and column whenever the ith eigenvalue occuring in D is 0. Then $L = (U'\sqrt{D'})^\dagger$ is a sequence of vectors in K^d satisfying $G = L^\dagger L$. ■

Definition 2.2.3. We call a sequence of vectors or lines modulo action by orthogonal resp. unitary matrices for $K = \mathbb{R}$ resp. $\mathbb{C}$ a *configuration* (of vectors or lines in K^d).

Proposition (2.2.1) tells us that configurations of equiangular vectors correspond to Gram matrices. Note that if L' is obtained from L by replacing some of the vectors v_j with $c_j v_j$, ($c_j \in K$, $|c_j| = 1$,) then L and L' have different Gram matrices even though they determine the same sequences of *lines*. We will later see that the *triple products* are a total invariant of line configurations, just as the Gram matrix is a total invariant of vector configurations. (Especially in the real case), making such replacements $v_j \mapsto -v_j$ is called *Seidel switching*.

Proposition 2.2.3 (The Cauchy-Binet formula)**.** Let A be an $n \times d$ matrix and let B be a $d \times n$ matrix over a commutative ring. Then

$$\det(AB) = \sum_M \det(A_M)\det(B_{M^T})$$

where the sum is over all $d \times d$ “submatrix masks” (akin to a “bitmask”), A_M is the $d \times d$ matrix corresponding to the mask M, and M^T is the “transposed mask”.

The Cauchy-Binet formula can be proved using the exterior algebra [Kon13].

Theorem 2.2.1 (Gram matrix theorem)**.** Let $(V, \langle \cdot, \cdot \rangle)$ be an inner product space over $\mathbb{R}$ or $\mathbb{C}$. Let L be a sequence of n vectors in V. Then L is linearly independent if and only if the Gram matrix for L is nonsingular.

Proof. Let $d = \dim(V)$ and let G be the Gram matrix for L. Choose an orthonormal basis for V so that we may regard L as the $d \times n$ matrix whose columns are coordinate vectors. Then $G = L^\dagger L$. By the Cauchy-Binet formula,

$$\det(G) = \sum_M \det(L_M)\det(L^\dagger_{M^T}) = \sum_M \det(L_M)\overline{\det(L_M)} = \sum_M |\det(L_M)|^2$$

The right-hand-side is zero if and only if $d \geq m$ and every $d \times d$ minor of L vanishes, which is the case if and only if L is linearly dependent. ■

Corollary 2.2.1.1. Let L be a rank r sequence of n vectors and $G = L^\dagger L$. Then $\operatorname{rk} G = r$.

Proof. G has rank at least r if and only if G has a nonsingular $r \times r$ submatrix. Such a submatrix is the is the Gram matrix of a subsequence R of L. By the Gram matrix theorem R is linearly independent. Such r is maximized when R has maximum rank. ■

We now give a second proof of the absolute bound.

Proof. Let $L = [v_1 \cdots v_n]$ be a sequence of equiangular unit vectors with angle α in K^d. Then

$$\langle \Pi_i, \Pi_j \rangle = \operatorname{tr}(\Pi_i \Pi_j) = \begin{cases} 1 & i = j \\ \alpha^2 & i \neq j \end{cases}$$

By the Gram matrix theorem, $\{\Pi_1, \cdots, \Pi_n\}$ is linearly independent in H_d if and only if the Gram matrix $G_{ij} = \langle \Pi_i, \Pi_j \rangle$ is nonsingular. Observe that $G = \alpha^2 J + (1 - \alpha^2) I$ where J is the all-ones matrix. The eigenvalues of G are $1 - \alpha^2$ and $(n-1)\alpha^2 + 1$, which are both nonzero. Hence G is nonsingular, so $\{\Pi_1, \cdots, \Pi_n\}$ is linearly independent. Finally recall that $\dim(H_{\mathbb{R},d}) = \binom{d+1}{2}$ and $\dim(H_{\mathbb{C},d}) = d^2$. ■

The absolute bound is not attained in general. Later we will show that $N(\mathbb{RP}^3) = 6 < 10 = \binom{5}{2}$. In the complex case, on the other hand, Zauner's conjecture asserts that $N(\mathbb{CP}^{d-1}) = d^2$ in every dimension d. This has been verified in dimensions up to 21, and many others, and numerical solutions exist in dimensions up to 139, [AFMY17]. In all of these cases, Heisenberg covariant configurations have been found, and only in dimension 8 a different type was found, known as the *Hoggar lines.*

The following proposition is a precursor to the relative bound.

Proposition 2.2.4. Let $L = [v_1 \cdots v_n]$ be a sequence of equiangular unit vectors in K^d with angle α. Suppose A is a matrix over K that can be expressed as a linear combination $A = \sum a_i \Pi_i$. Then the coefficients a_i are given by the formula

$$a_i = \frac{\operatorname{tr}(A\Pi_i) - \alpha^2 \operatorname{tr}(A)}{1 - \alpha^2}$$

Proof. The formula follows from the calculation

$$\operatorname{tr}(A\Pi_j) = \operatorname{tr} \sum a_i \Pi_i \Pi_j = a_j + \alpha^2 \sum_{i \neq j} a_i = a_j + \alpha^2(\operatorname{tr}(A) - a_j) = \alpha^2 \operatorname{tr}(A) + a_j(1 - \alpha^2)$$

■

Corollary 2.2.1.2. Suppose that the identity matrix can be expressed as a linear combination $I = \sum a_i \Pi_i$. Then the angle is given by

$$\alpha^2 = \frac{n-d}{d(n-1)} \tag{2.6}$$

Proof. Plug in $A = I$ to Proposition (2.2.4) and take the trace. ■

Corollary 2.2.1.3. Suppose that the identity matrix can be expressed as a linear combination $I = \sum a_i \Pi_i$. Then the coefficients are given by $a_i = d/n$

In the complex case and in the context of quantum mechanics, the fact that an equiangular tight frame satisfies $\frac{d}{n} \sum \Pi_i = I$ is interpreted as the matrices $\frac{d}{n} \Pi_i$ constituting a symmetric positive operator-valued measure (symmetric POVM). The probability that a state described by density matrix ρ is observed in state i is

$$p_i = \frac{d}{n} \operatorname{tr}(\rho \Pi_i) = \frac{d}{n}(\alpha^2 + a_i(1-\alpha^2)) \tag{2.7}$$

When the absolute bound is attained, $n = d^2$, the POVM elements span H_d, so ρ is a linear combination $\rho = \sum a_i \Pi_i$, and the coefficients a_i can be experimentally recovered via (2.7) by performing the measurement many times on identically prepared states ρ. The symmetric POVM is thus called *informationally complete*, or a *SIC POVM*, or simply *SIC*.

Corollary 2.2.1.4. A set of equiangular lines attaining the absolute bound is a tight frame with angle $\alpha = 1/\sqrt{d+2}$ resp. $\alpha = 1/\sqrt{d+1}$ in the real resp. complex case.

Proof. When the absolute bound is attained, $\{\Pi_i\}$ is a basis for H_d, so in particular the identity matrix is spanned by $\{\Pi_i\}$ and we may apply the corollary (2.2.1.2). ■

The relative bound states that in general, equation (2.6) is an inequality.

Theorem 2.2.2 (The relative bound)**.** Let L be an equiangular sequence of n unit vectors in K^d with angle α. Then regardless of whether $K = \mathbb{R}$ or $\mathbb{C}$,

$$\alpha^2 \geq \frac{n-d}{d(n-1)} \qquad \text{equivalently} \qquad n \leq \frac{d(1-\alpha^2)}{1-d\alpha^2} \tag{2.8}$$

Equality holds if and only if $\frac{d}{n}\sum \Pi_i = I$.

Consequently, we get the following inequality, which occasionally shows up in calculations.

$$\sum_{ij} |\langle v_i, v_k\rangle|^2 = n(n-1)\alpha^2 + n \geq \frac{n^2}{d} \tag{2.9}$$

[CKM16] offered the following elegant proof.

Proof. The relative bound follows from the non-negativity of the following expression.

$$\left\langle I - \frac{d}{n}\sum \Pi_i \ , \ I - \frac{d}{n}\sum \Pi_i \right\rangle = d - 2d + \frac{d^2}{n^2}(n(n-1)\alpha^2 + n)$$

The bound is attained if and only if the above expression is 0, that is $I = \frac{d}{n}\sum \Pi_i$. ■

Definition 2.2.4. Let L be a sequence of n equiangular unit vectors in K^d at angle α. If the following equivalent conditions hold, then L is an *equiangular tight frame* or *ETF*:

1. The relative bound is attained.
2. $\sum_{ij} |\langle v_i, v_j\rangle|^2 = \frac{n^2}{d}$.
3. $\frac{d}{n}\sum_i \Pi_i = I$.
4. $I \in \mathrm{span}_{\mathbb{R}}\{\Pi_1, \cdots, \Pi_n\}$.
5. The Gram matrix G has exactly two eigenvalues.
6. The *Seidel matrix* X has exactly two eigenvalues. ($X := \frac{1}{\alpha}(G - I)$.)

In such a case, the eigenvalues of G are 0 with multiplicity $n - d$ and n/d with multiplicity d, and the eigenvalues of X are $-1/\alpha$ and $(n-d)/(d\alpha)$ with the same multiplicities.

It only remains to show that 1,2,3,4 are equivalent to 5,6, and to derive the formulas for the eigenvalues. Note that the eigenvalues of G and X differ only by a linear relation.

Proof. Suppose 3 holds. Then

$$G^2 = L^\dagger L L^\dagger L = L^\dagger \sum \Pi_i L = \frac{n}{d} G$$

By the Caley-Hamilton theorem, the eigenvalues of G are in $\{0, \frac{n}{d}\}$. The multiplicity of 0 is $n-d$ where $d = \operatorname{rk}(G) = \operatorname{rk}(L)$. The multiplicity of $\frac{n}{d}$ is determined by the fact $\operatorname{tr}(G) = n$. On the other hand, if G has three or more eigenvalues, then G does not satisfy a quadratic relation (since each eigenvalue is a root of the minimal polynomial), so 3 does not hold. Finally, note that if G only has one eigenvalue, then G is a multiple of the identity, so the off-diagonal entries are 0. This is the degenerate case of orthogonal vectors, which we disregard. ■

We will later discuss more conditions equivalent to the ETF property in the real case. It is worth noting that the *tight frame* property (3) can be extended to sequences of vectors that are not necessarily equiangular, or even unit vectors. We introduce these more general *frames* in section 2.4. *Tight frame* is a property of configurations of vectors, that is, like equiangularity, it is invariant under orthogonal/unitary matrices.

Theorem 2.2.2 shows that the ETF property depends only on the angle, dimension, and number of lines. Example 5.2.2 provides distinct complex ETF configurations with the same number of lines in the same dimension, despite of course having the same angle. These configurations even attain the absolute bound. 3 is the lowest complex dimension where this phenomenon arises. It would be interesting to know if this phenomenon arises with fewer lines in $\mathbb{C}^3$ and what the smallest real example is. The latter could be checked exhaustively (in principle), as we will see that configurations of real equiangular vectors correspond to graphs, though if the dimension is large, the search may be intractable.

The proof of Proposition 2.2.2 shows how to obtain a sequence of equiangular vectors from its Gram matrix G, and thereby from its graph or Seidel matrix, but this construction is not unique, requiring a choice of orthonormal eigenbasis for G. The eigenspaces typically have high dimension, thus requiring a "creative" choice. We now show how a *canonical* sequence of equiangular vectors can be constructed. In this setting, it is convenient to note a relationship to the graph isomorphism problem and define the complement of an ETF.

We begin with an $n \times n$ *Seidel matrix* X, (any symmetric matrix with 0s on the diagonal and ± 1 elsewhere). Let α be the negative reciprocal of the least eigenvalue of X. Set $G = \alpha X + I$ and let $L = [w^1 \cdots w^n]$ be the $n \times n$ orthogonal projection matrix of K^n onto $\operatorname{im}(G)$. The columns of L are the projections of the standard basis vectors onto $\operatorname{im}(G)$, and G is the Gram matrix for the L with respect to the inner product $\langle \cdot, \cdot \rangle_G$ defined by

$$\langle v, u \rangle_G := u^\dagger G v$$

Alternately, G is the Gram matrix of $\widetilde{L} = \sqrt{G}L$ with respect to the standard inner product on K^n, where $\sqrt{G}$ is the unique positive-semidefinite square root of G.

Now let X_1, X_2 be Seidel matrices and let $G_1, G_2, L_1, L_2, \widetilde{L}_1, \widetilde{L}_2$ be determined as above. Then it is easy to prove that the following are equivalent:

1. $X_2 = \Sigma X_1 \Sigma^{-1}$ (The graphs Y_1, Y_2 are *isomorphic* with isomorphism $\Sigma : Y_1 \to Y_2$.)
2. $G_2 = \Sigma G_1 \Sigma^{-1}$
3. $L_2 = \Sigma L_1 \Sigma^{-1}$
4. $\widetilde{L}_2 = \Sigma \widetilde{L}_1 \Sigma^{-1}$

The interpretation of 3 and 4 from the sequence of equivalent conditions above is that isomorphic graphs define the same equiangular vectors up to permutation, but this permutation is applied both to the sequence of vectors and to the coordinates of the vectors themselves, thus preserving the hardness of the graph isomorphism problem. (If the permutation were only applied to one of these things, the problem would be as easy as sorting.)

Now let us introduce the *complement* of an ETF. Let G be the Gram matrix of an ETF of n unit vectors of rank d. Recall that $G^2 = \frac{n}{d}G$. Thus $L = \frac{d}{n}G$ is the orthogonal projector onto $\operatorname{im}(G)$. Consider the orthogonal projector $L' = I - L$ onto $\operatorname{im}(G)^\perp$. Let $G' = \frac{n}{n-d}L'$. Then G' is positive semi-definite, so it is a Gram matrix. The diagonal entries are $G'_{ii} = \frac{n}{n-d}(1 - \frac{d}{n}) = 1$ and the off diagonal entries are

$$G'_{ij} = \frac{n}{n-d}\left(-\frac{d}{n}G_{ij}\right) = \mp \frac{d}{n-d}\sqrt{\frac{n-d}{d(n-1)}} = \mp \sqrt{\frac{n-(n-d)}{(n-d)(n-1)}}$$

so G' is the Gram matrix of an ETF of n unit vectors of rank $n-d$. This configuration is called the *complementary ETF*. Note that L' is a sequence of vectors with Gram matrix G'. In other words, the columns v_i of L are the projections of the standard basis vectors e_i of $\mathbb{R}^n$ onto $\operatorname{im}(G)$, and the columns u_i of L' are the projections of e_i onto $\operatorname{im}(G)^\perp$. If X and Y are the Seidel matrix and graph of an ETF, then the Seidel matrix and graph of the complementary ETF are $-X$ and the complement graph Y^C. The simplex configuration is therefore the unique configuration whose complement is the degenerate configuration.

In Section 2.4 we will introduce the *canonical tight frame* of a given frame. In this language, L is almost the canonical tight frame of $\widetilde{L}$. More precisely, let B be an orthonormal basis for $\operatorname{im} G$, and let L_0 be the $d \times n$ matrix whose columns are the coordinate vectors relative to B for the columns of L. Similarly define $\widetilde{L}_0$ in terms of $\widetilde{L}$. Then $L_0 = (\widetilde{L}_0 \widetilde{L}_0^\dagger)^{-1/2} \widetilde{L}_0 = \widetilde{L}_0^{\mathrm{ctf}}$.

2.3 Triple products

We have seen that the Gram matrix $\langle v_i, v_j \rangle$ is a total invariant of a configuration of nonzero vectors. Similarly, it turns out that the 3-tensor of *triple products* $T_{ijk} = \operatorname{tr} \Pi_i \Pi_j \Pi_k = \langle v_i, v_j \rangle \langle v_j, v_k \rangle \langle v_k, v_i \rangle$ is a total invariant of a configuration of lines provided that none of the pairwise inner products are 0. This requirement is satisfied when the lines are equiangular. (Recall that by convention we do not consider orthogonal lines to be equiangular). Chien and Waldron showed that the complete sequence of m-products $T_{i_1 \cdots i_m} = \operatorname{tr} \Pi_{i_1} \cdots \Pi_{i_m} = \langle v_{i_1}, v_{i_2} \rangle \cdots \langle v_{i_{m-1}}, v_{i_m} \rangle \langle v_{i_m}, v_{i_1} \rangle$, for $m = 1, \cdots, n$, is a total invariant of any ordered configuration of n lines [CW16]. The m-products are also called the *Bargmann invariants.*

Proposition 2.3.1. The m-products are invariants of a configuration of lines.

This proof is rather trivial. Towards proving that the triple products are a complete invariant we first consider the angle 2-tensor θ_{ij}, defined by $\langle v_i, v_j \rangle = |\langle v_i, v_j \rangle| e^{i\theta_{ij}}$.

Proposition 2.3.2. Let θ^1_{ij} and θ^2_{ij} be angle 2-tensors from two sequences L_1, L_2 of vectors. Then L_1, L_2 determine the same configuration of lines if and only if there is a function $i \mapsto \phi_i$ such that for each i and j, $\theta^1_{ij} = \theta^2_{ij} + \phi_i - \phi_j$.

Proof. Let G_1 and G_2 be the gram matrices for sequences of vectors $L_1 = [v^1_1 \cdots v^1_n]$ and $L_2 = [v^2_1 \cdots v^2_n]$. Then L_1 and L_2 determine the same configurations of n lines if and only if there exists a diagonal matrix $D = \operatorname{diag}(e^{i\phi_1}, \cdots, e^{i\phi_n})$ with complex numbers of norm 1 on the diagonal such that $G_2 = D G_1 D^\dagger$, if and only if the ijth entry of G_2 is $e^{i\theta^2_{ij}} = e^{i(\theta^1_{ij} + \phi_i - \phi_j)}$. ■

The angle m-tensor $\theta_{i_1, \cdots, i_m}$ is defined by $T_{i_1, \cdots, i_m} = |T_{i_1, \cdots, i_m}| e^{i\theta i_1, \cdots, i_m}$. Note that $\theta_{i_1, \cdots, i_m} = \theta_{i_1, i_2} + \cdots + \theta_{i_{m-1}, i_m} + \theta_{i_m, i_1}$.

Theorem 2.3.1. Two sequences of vectors L_1, L_2, each with all nonzero pairwise inner products, determine the same configuration of lines if and only if they have the same triple products.

Proof. Since no pairwise inner product is zero, $T^1_{ijk} = T^2_{ijk}$ if and only if $\theta^1_{ijk} = \theta^2_{ijk}$. In such a case, for a fixed k (eg. $k = 1$), θ^2_{ij} has the form $\theta^1_{ij} + \phi_i - \phi_j$ where for each i, $\phi_i = \theta^1_{ki} - \theta^2_{ki}$. Note that $\phi_j = \theta^1_{kj} - \theta^2_{kj} = -(\theta^1_{jk} - \theta^2_{jk})$. The result now follows from lemma (2.3.2). ■

Note that the tensors T_{ijk} and θ_{ijk} are cyclic in their subscripts, and that $T_{kji} = \overline{T}_{ijk}$. In the real case, $T_{ijk} = T_{kji}$, so for distinct i, j, k, the triple product T_{ijk} depends only on the set of indices $\{i, j, k\}$. Moreover, in the real case, if Y is the graph associated to a sequence of equiangular unit vectors (cf. Section 3.1) with triple products T_{ijk}, then for distinct i, j, k,

$$T_{ijk} = \begin{cases} +\alpha^3 & : Y[\{i,j,k\}] \text{ has an odd number of edges} \\ -\alpha^3 & : Y[\{i,j,k\}] \text{ has an even number of edges} \end{cases}$$

where $Y[\{i, j, k\}]$ is the subgraph induced by $\{i, j, k\}$. The set of 3-sets $\{i, j, k\}$ such that $T_{ijk} = -\alpha^3$ is called the *two-graph* of Y, which we denote $\delta(Y)$. Note that any 4-set of vertices of $\delta(Y)$ contains an even number of 3-sets of $\delta(Y)$. By Theorem (2.3.1), $\delta(Y)$ is a total invariant of the ordered switching class of Y. We will discuss two-graphs more in the section on the chain of graph-like structures.

We have seen that when no two lines are orthogonal, the triple products determine the configuration, and thus the m-products depend on the triple products. The following proposition explicitly describes this dependence. The proof is trivial.

Proposition 2.3.3. $T_{i_1,\cdots,i_m} = T_{i_1,i_2,i_3} T_{i_1,i_3,i_4} T_{i_1,i_4,i_5} \cdots T_{i_1,i_{m-1},i_m}$.

Additional properties of the triple products arise when the lines are equiangular and the rank 1 projection matrices $\Pi_1, \cdots, \Pi_n$ span the vector space $\mathfrak{gl}(d, K)$ of $d \times d$ matrices over K. Since this space has K-dimension d^2 and the Π_i are elements of the $\binom{d+1}{2}$ or d^2 dimensional real vector space H_d, this only happens in the complex case, and happens if and only if $\{\Pi_i\}$ is a SIC-POVM. It was shown in [AFF11] that the structure constants for $\mathfrak{gl}(d, \mathbb{C})$ as an algebra or Lie algebra are closely relate to the triple products.

Proposition 2.3.4. Let $\{\Pi_1, \cdots, \Pi_{d^2}\}$ be a SIC-POVM. Then

$$\Pi_i \Pi_j = \frac{d+1}{d} \sum_{i=1}^{d^2} T_{ijk} \Pi_k - \frac{d\delta_{ij} + 1}{d+1} I$$

Proof. Plug $A = \Pi_i \Pi_j$ into proposition (2.2.4) and note that $\operatorname{tr}(\Pi_i \Pi_j) = \frac{d\delta_{ij}+1}{d+1}$. ■

Appleby, Flammia, and Fuchs [AFF11] define

$$\begin{aligned} J_{ijk} &= \frac{d+1}{d}(T_{ijk} - \overline{T}_{ijk}) \\ R_{ijk} &= \frac{d+1}{d}(T_{ijk} + \overline{T}_{ijk}) \end{aligned}$$

Proposition 2.3.5. Let $\{\Pi_1, \cdots, \Pi_{d^2}\}$ be a SIC-POVM. Then

$$[\Pi_i, \Pi_j] = \sum_{k=1}^{d^2} J_{ijk}\Pi_k$$

Proof. This is a straightforward calculation using (2.3.4). ■

[AFF11] go on to derive an algebraic characterization of SIC-POVMs using triple products: There exists a basis for $H_{d,\mathbb{C}}$ with the so-called Q-Q^T property if and only if there exists a SIC in dimension d. This basis may or may not itself be the SIC.

2.4 Frames

Frames arise in various fields including quantum information theory. Tight frames generalize our previously defined equiangular tight frames, and normalized tight frames define positive operator-valued measures (POVMs), which are used for quantum state tomography. Every frame gives rise to a canonical (normalized) tight frame (CTF) in the same dimension. In this way one obtains a rank 1 informationally complete POVM from any equiangular frame, called a *pretty good measurement*. In fact, we saw such CTFs at the end of Section 2.2 when we showed how to reconstruct a set of equiangular vectors from a graph. It turns out that this CTF can be read directly off a *graph* as follows: First compute the least eigenvalue $-1/\alpha$ of the Seidel matrix. Then, L^{ctf} is simply the sequence of columns of the $n \times n$ orthogonal projection matrix of K^n onto $\text{im}(\alpha X + I)$. We also introduce mutually unbiased bases in this section, which arise in quantum theory from complementary observables and have applications in quantum cryptography and state tomography. Waldron has several papers and a textbook on frame theory, and for the most part we use his terminology [Wal18].

Definition 2.4.1. A *frame* for a vector space V is a sequence of vectors spanning V.

A frame of unit vectors is often called a *unit frame*. Aside from the example given in the introductory paragraph to this section, we will always work in the vector space $V = K^d$.

Proposition 2.4.1. Let A and B be rectangular matrices such that AB and BA are defined. Then AB and BA have the same nonzero eigenvalues.

Proof. Let λ be a nonzero eigenvalue of AB with eigenvector v. Then $Bv \neq 0$ and $\lambda Bv = BABv$, so λ is an eigenvalue of BA (with eigenvector Bv). ■

Definition 2.4.2. A frame $L = [v_1 \cdots v_n]$ is a *tight frame* (with constant c) if $LL^\dagger = cI$ for some $c \in k$. Taking trace of both sides, it follows that

$$c = \frac{1}{d} \sum_i |v_i|^2 \tag{2.10}$$

A tight frame is *normalized* if $c = 1$. This can be achieved by replacing each vector v_i of a tight frame with $\sqrt{d/(n|v_i|^2)}v_i$. A unit normalized tight frame is an orthonormal basis.

By Proposition 2.4.1, $LL^\dagger$ is invertible and its inverse is positive semi-definite, so define:

Definition 2.4.3. The *canonical tight frame* associated to a frame L is $L^{\text{ctf}} = (LL^\dagger)^{-1/2}L$.

Using the trace inner product, $\langle L^{\text{ctf}}L^{\text{ctf}\dagger}, L^{\text{ctf}}L^{\text{ctf}\dagger}\rangle = 1$ and $\langle L^{\text{ctf}}L^{\text{ctf}\dagger}, I\rangle = 1$, so it follows that L^{ctf} is a normalized tight frame. Also note that if L is a tight frame, then $L^{\text{ctf}} = \frac{1}{\sqrt{c}}L$.

Remark 2.4.1. A frame L is a tight frame with constant c if and only if its Gram matrix $L^\dagger L$ equals c times an orthogonal projection matrix. Proposition (2.4.1) shows that this holds if and only if $L^\dagger L$ has ≤ 2 eigenvalues. (If there are two then one of them is 0.)

Definition 2.4.4. Frames L_1 and L_2 are *complementary* if $L_1^{\text{ctf}\dagger}L_1^{\text{ctf}} + L_2^{\text{ctf}\dagger}L_2^{\text{ctf}} = I$.

Proposition 2.4.2. Every frame has a complement, and if L_2 and L_2' are complements of L_1, then L_2 and L_2' are similar ($L_2' = AL_2$ for some $A \in \text{GL}(n, K)$). Moreover, complementary tight frames are orthogonally/unitarily equivalent.

Proof. Let L_1 be a frame. Then $G = I - L_1^{\text{ctf}\dagger}L_1^{\text{ctf}}$ is an orthogonal projection matrix, so it is positive semidefinite. By Proposition (2.2.2), G is the gram matrix of some frame L_0, which is unique up to orthogonal/unitary equivalence. By Remark (2.4.1), L_0 is a tight frame. Now if $L_0 = L_2^{\text{ctf}}$ for some frame L_2 then clearly L_2 and L_0 are similar. ■

Common synonyms for *complementary frame* include *Naimark complement* and *Gale dual.* If Y is the graph corresponding to an ETF (cf. Section 3.1), then the graph corresponding to its complement frame is the graph complement of Y, so we prefer the term *complement.* The term *Gale dual* also clashes with Waldron's *dual frame* [Wal18].

Next we present a generalization of the relative bound for equiangular lines, called the *first Welch bound.* It can be viewed as an application of the Cauchy-Schwarz inequality.

Proposition 2.4.3 (Cauchy-Schwarz inequality)**.** Let u, v be vectors of an inner product space, (for our purposes finite dimensional). Then

$$|\langle u, v\rangle|^2 \leq |u|^2 |v|^2$$

with equality if and only if u and v are linearly dependent.

Proposition 2.4.4. Let A be a $d \times d$ matrix with real eigenvalues. Then $(\operatorname{tr} A)^2 \leq d \operatorname{tr}(A^2)$ with equality if and only if A is a multiple of the identity.

Proof. Let $\lambda_1, \cdots, \lambda_d$ be the eigenvalues of A. Let $\lambda = (\lambda_1, \cdots, \lambda_d)$ and let j be the all 1s vector with d entries. Then by the Cauchy-Schwarz inequality,

$$(\operatorname{tr} A)^2 = \left(\sum \lambda_i\right)^2 = \langle j, \lambda\rangle^2 \leq |j|^2 |\lambda|^2 = d \sum \lambda_i^2 = d \operatorname{tr}\big(A^2\big)$$

with equality if and only if $\lambda_1 = \cdots = \lambda_d$. ■

Theorem 2.4.1 (First Welch bound)**.** Let $L = [v_1 \ \cdots \ v_n]$ be a unit frame for K^d. Then

$$\max_{i \neq j} |\langle v_i, v_j\rangle|^2 \geq \frac{n-d}{d(n-1)}$$

with equality if and only if L is an ETF.

Proof. Apply proposition (2.4.4) with $A = LL^\dagger$ to get

$$\left(\sum_i |v_i|^2\right)^2 \leq d \sum_{ij} |\langle v_i, v_j\rangle|^2 \tag{2.11}$$

with equality if and only if L is a tight frame. The assumption that L is a unit frame simplifies the left-hand-side of (2.11):

$$\frac{n^2}{d} \leq \sum_{ij} |\langle v_i, v_j \rangle|^2$$

Comparing to the maximum remaining terms apart from the diagonal terms, we obtain

$$\frac{n^2}{d} - n \leq \sum_{i \neq j} |\langle v_i, v_j \rangle|^2 \leq n(n-1) \max_{i \neq j} |\langle v_i, v_j \rangle|^2$$

which simplifies to the Welch bound. The left inequality is sharp if and only if L is a tight frame, and the right inequality is sharp if and only if L is equiangular. ■

Next we introduce a concrete non-equiangular example of a class of frames for $\mathbb{C}^d$: MUBs.

Definition 2.4.5. Orthonormal bases $\mathcal{B}_1, \mathcal{B}_2$ for $\mathbb{C}^d$ are called *mutually unbiased bases, (MUBs),* if there exists a real number c such that for each $u \in \mathcal{B}_1$ and $v \in \mathcal{B}_2$, $|\langle u, v \rangle|^2 = c$. It follows that c must equal $1/d$.

There is an absolute bound for MUBs. The proof we give below resembles that of the absolute bound for equiangular lines in the sense that we reason in the inner product space of hermitian $d \times d$ matrices.

Theorem 2.4.2. The maximum number of MUBs in $\mathbb{C}^d$ is less than or equal to $d+1$.

Proof. Let $\{\mathcal{B}_1, \cdots, \mathcal{B}_t\}$ be a set of MUBs for $\mathbb{C}^d$. For a unit vector $u \in \mathbb{C}^d$, define

$$\Gamma_u = \sqrt{\frac{d}{d-1}} \left(uu^\dagger - \frac{1}{d} I \right)$$

Then for $u \in \mathcal{B}_i$ and $v \in \mathcal{B}_j$ we have

$$\langle \Gamma_u, \Gamma_v \rangle = \frac{d}{d-1} \left(|\langle u, v \rangle|^2 - \frac{1}{d} \right) = \begin{cases} 0 & i \neq j \\ -\frac{1}{d-1} & i = j \text{ and } u \neq v \\ 1 & i = j \text{ and } u = v \end{cases}$$

Thus, the Gram matrix for $(\Gamma_u \,|\, u \in \mathcal{B}_i)$ is that of a regular simplex with d vertices centered at the origin of $\mathbb{R}^{d-1}$. Thus if $W_i = \operatorname{span}\{\Gamma_u \,|\, u \in B_i\}$, then $\dim W_i = d-1$, and for $i \neq j$, $W_i \perp W_j$. Since the W_i are subspaces of the d^2-1 real dimensional inner product space of traceless hermitian $d \times d$ matrices, $t \leq (d^2-1)/(d-1) = d+1$. ■

The absolute bound is attained in prime power dimension (see Example 5.3.2). It is unknown whether the bound is always attained or not, or even whether it is ever attained in non-prime-power dimension. In fact, the maximum number of MUBs for $\mathbb{C}^6$ is known only to be at least 3.

Remark 2.4.2 (Connections to quantum information theory)**.** An *observable* of a d level quantum system is a hermitian $d \times d$ matrix $A = \sum \lambda_i \Pi_i$ where Π_i are pairwise orthogonal projectors. If ρ is the density matrix representing the state of the system, then the probability of observing outcome λ_i is $\operatorname{tr}(\rho\Pi_i)$. In such a case, the system collapses into the state $\rho' = \Pi_i \rho \Pi_i / \operatorname{tr}(\Pi_i \rho)$.

A *positive operator-valued measure* (POVM) for a d level quantum system is a set of complex positive semidefinite $d \times d$ matrices $\{E_i\}$ such that $\sum E_i = I$, such as the orthogonal projectors onto the vectors of a normalized tight frame are a POVM. Measurement of a POVM $\{E_i\}$ with respect to a density matrix ρ yields outcome i with probability $p_i = \operatorname{tr}(\rho E_i)$. If $\{E_i\}$ spans the real inner product space of hermitian $d \times d$ matrices, then ρ is determined by the probabilities p_i and the numbers $\operatorname{tr}(E_i E_j)$.

Consider the problem of *quantum hypothesis testing*: A quantum state ρ is promised to be one of the states $\sigma_1, \cdots, \sigma_n$ with respective nonzero probabilities $p_1, \cdots, p_n$. We have full mathematical description of the σ_i and p_i, and we are tasked with performing a single POVM measurement to determine ρ with maximum success probability. This problem is unsolved for $n \geq 3$, but there is a known measurement that gives high success probability, and is therefore dubbed the *pretty good measurement*. The measurement operators are $E_i = S^{-1/2} p_i \sigma_i S^{-1/2}$ for each $i \in \{1, \cdots, n\}$ where $S = \sum_{i=1}^n p_i \sigma_i$. If $\operatorname{span}\{\sigma_1, \cdots, \sigma_n\} = H_{\mathbb{C},d}$, then the linear map $A \mapsto S^{-1/2} A S^{-1/2}$ is an automorphism of $H_{\mathbb{C},d}$, and so the pretty good measurement is informationally complete. For a uniform distribution of pure states $v_1 v_1^\dagger, \cdots, v_n v_n^\dagger$ such that $\operatorname{span}_{\mathbb{C}}\{v_1, \cdots, v_n\} = \mathbb{C}^d$, before simplifying we see that $\sum_{i=1}^n E_i = L^{\mathrm{ctf}} L^{\mathrm{ctf}\dagger}$, illustrating the connection to canonical tight frames, and then of course this quantity is equal to the identity matrix. By its nature, the pretty good measurement disregards the ordering of the measurement operators, that is, a POVM is a *set* of measurement operators, not a *sequence*. Perhaps this property could be exploited in such a way to address the graph isomorphism problem in a quantum setting.

Roughly speaking, two observables A, A' are *complementary* if complete knowledge of one of them implies that the measurement probabilities of the other one are uniform among all outcomes. Thus complementary observables A, A' admit mutually unbiased eigenbases. According to [KR05, Ivo81], at minimum, $d+1$ observables are needed for complete state reconstruction. This could conceivably be achieved by a set of $d+1$ complementary observables $A_1, \cdots, A_{d+1}$, say with respective mutually unbiased eigenbases $\mathcal{B}_1, \cdots, \mathcal{B}_{d+1}$. This is because, for each vector u in each basis $\mathcal{B}_i$, the quantity $\operatorname{tr}\left(\rho u u^\dagger\right)$ is approximated by repeated measurements of observable A_i. From this, one can calculate $\operatorname{tr}(\rho \Gamma_u)$ with Γ_u as defined above. Since $\{\Gamma_u \mid u \in \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_{d+1}\}$ spans the space of traceless $d \times d$ hermitian matrices, ρ can be recovered. MUBs and SICs have the same strange property: the minimum number of objects necessary for complete state reconstruction is the maximum conceivable number of such objects, and moreover, it is elusive in both cases whether such complete sets can always exist. See Example 5.3.1 for the qubit case and the usual state tomography method for a system of n qubits ($d = 2^n$) that uses $4^n - 1 (\geq 2^n + 1)$ observables. This method of state and process tomography is commonly employed by experimentalists, but it requires more data to be collected than is need when working with MUBs or SICs. We also show in Example 5.3.1 the BB84 quantum key exchange, which makes use of two MUBs for $\mathbb{C}^2$, and has generalizations that use more MUBs in higher dimensions.

Chapter 3

Real equiangular lines

In the first and second section we discuss how real equiangular lines can be studied combinatorially and present two more formulations of the ETF property, which are valid in the real case only. Seidel matrices with three or more eigenvalues have been studied extensively, as they are a natural progression from studying ETFs, but we do not discuss them here. In Section 3.3, we review the pillar method of Lemmens and Seidel, and retrace their steps in proving the maximum number $N(\mathbb{RP}^{d-1}) = N(S^{d-1})$ of real equiangular lines (or vectors) in dimensions $d = 2, \cdots, 13$.

Brute force computations are possible in the real case, but are typically inefficient, requiring an iteration over all graphs on vertex set $[n] = \{1, \cdots, n\}$. A significant part of the computation for proving the nonexistence of a n equiangular lines in low dimensions d can be diverted into proofs of Ramsey numbers, or even bounds on them. On the other hand, to demonstrate existence of an equiangular configuration by exhaustive search, one might find computational savings by building onto potential subconfigurations. We have omitted most of our computational work as it has not been very fruitful.

In this chapter we also discuss some generalized graph-like structures, and we prove a theorem about the number of unordered structures depending on the number of generalized edges. It is also worth noting two other important recent developments in the field that we do not have time or space to include: In 2016, it was proved that for each (fixed) α, $N(S^{d-1}, \pm\alpha)$ is of linear order, and explicit coefficients were given. [BDKS16] Then, in 2019, an explicit formula for $N(S^{d-1}, \pm\alpha)$ was given for sufficiently large d [JTY+19].

3.1 Seidel switching

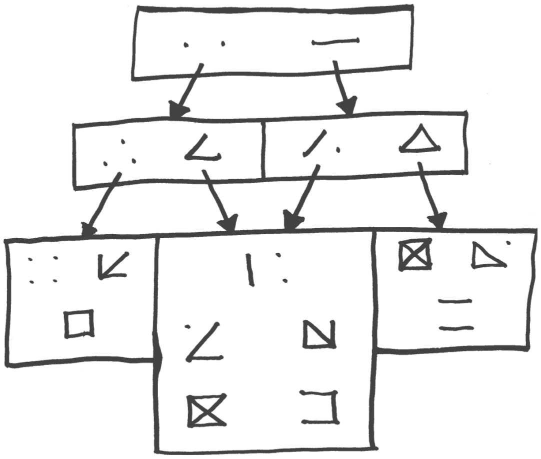

(3.1)

We begin by introducing the most basic operation involving real equiangular lines: Seidel switching. In some situations, it is important to distinguish between *graphs* (pairs (V, E) where V is a finite set and E is a set of two-element subsets of V) on vertex set $V = [n]$ and isomorphism classes of graphs, and to make similar distinctions for other kinds of objects between their ordered and unordered counterparts. For example, Figure (3.1) shows unordered switching classes, which behave entirely differently from ordered switching classes.

Let G be the Gram matrix for a sequence L of n equiangular unit vectors with angle α. Recall that $X = \frac{1}{\alpha}(G - I)$ is called the *Seidel matrix* for L. The matrix X is symmetric with all 0s on the diagonal and ± 1 elsewhere. The Seidel matrix can be defined in the same way in the complex case, where it is sometimes called the *signature matrix.* In the complex case, X is Hermitian, has 0s on the diagonal complex numbers of absolute value 1 elsewhere. This "$U(1)$ phase" makes the complex case much harder to analyze combinatorially, and for that reason, we restrict our attention to the real case and present this section in the current chapter. Understanding the number fields these phases come from in the complex case is an active area of research, and beyond the scope of this thesis.

Conversely, given a symmetric matrix X with 0s on the diagonal and ± 1 elsewhere, by proposition (2.2.2), $G = \alpha X + I$ is the Gram matrix of a sequence of real equiangular equiangular unit vectors. We associate to X the graph Y on vertex set $[n]$ with edge ij if and only if $X_{ij} = +1$ and note that X and Y encode exactly the same data.

Definition 3.1.1. We call an $n \times n$ Gram matrix or sequence L of n equiangular unit vectors *rigid* if 0 is an eigenvalue of the Gram matrix $G = L^\dagger L$.

The multiplicity of 0 is $n - d$ where $d = \operatorname{rk} G = \operatorname{rk} L$. We define the *dimension* of a Seidel matrix X or a graph Y to be this number d and denote it by $d(X)$ or $d(Y)$.

Among all positive semidefinite matrices of the form $tX + I$, where X is the Seidel matrix of a sequence of equiangular unit vectors and $t \in \mathbb{R}$, the only rigid matrix is $G = \alpha X + I$, and it is the only one with rank $< n$. Thus rigid configurations (and not all configurations) of equiangular unit vectors correspond to Seidel matrices and graphs on $[n]$.

We now define Seidel switching for graphs on $[n]$ (and for isomorphism classes of graphs), though it can just as well be defined for Seidel matrices and configurations of equiangular unit vectors. Let Y be a graph on $[n]$ and let v be a vertex. The graph Y_v obtained from Y by *Seidel switching* v has edge set $E \triangle E_v$ where E is the edge set of Y, E_v is the edge neighborhood of v in Y, and $\triangle$ denotes the symmetric difference of sets. Seidel switching is involutionary and commutative. Thus we may define the graph Y_θ obtained by Seidel switching a set $\theta \subseteq [n]$ of vertices, and $(Y_\theta)_\tau = Y_{\theta \triangle \tau}$ for $\theta, \tau \subseteq [n]$. Eg. $Y = Y_\emptyset = Y_{[n]}$. Seidel switching works exactly the same with Seidel matrices and Gram matrices. For a sequence L of vectors, Seidel switching works slightly differently, as $L_{[n]} = -L \neq L$. This is why we deem two sequences L, L' *of the same configuration* if they differ by some $A \in \mathrm{O}(n)$ and not $\mathrm{SO}(n)$ (and similarly with unitary equivalence in the complex case).

Now, for the first time we must carefully distinguish between ordered and unordered structures. If Z is a graph on $[n]$ and Σ is an $n \times n$ permutation matrix, then we denote by $\Sigma Z \Sigma^{-1}$ the graph obtained from Z by permuting the vertex set according to Σ. Two graphs Y, Y' on $[n]$ are *(ordered) switching equivalent* if there exists $\theta \subseteq [n]$ such that $Y' = Y_\theta$. These graphs are *unordered switching equivalent* if there exist $\theta \subseteq [n]$ and a permutation matrix Σ such that $Y' = \Sigma Y_\theta \Sigma^{-1}$. Two isomorphism classes of graphs $\mathcal{Y}, \mathcal{Y}'$ on n vertices are *switching equivalent* if, after labeling the vertices of each with $[n]$, the resulting graphs are unordered switching equivalent. These are of course equivalence relations. Note that even for an isomorphism class of graphs $\mathcal{Y}$, if certain vertices satisfy an "identifiable property" θ, such as having even degree, then we can make sense of the notation $\mathcal{Y}_\theta$ (the isomorphism class of graphs obtained by switching all vertices satisfying the property θ).

Given a graph Y on $[n-1]$, the set of graphs on $[n]$ having Y as the subgraph induced by $[n-1]$ is an ordered switching class. Thus, for graphs on $[n]$, there are $2^{\binom{n-1}{2}}$ ordered switching classes, each of size 2^{n-1}. The unordered switching classes are much more interesting. The number of classes is unknown, and they vary in size for a fixed n. The following theorem (and its proof) make sense for both graphs and isomorphism classes graphs, but we include it especially for the perspective it offers on unordered switching classes.

Theorem 3.1.1. If E is the set of vertices of even degree in Y and n is odd, then Y_E is the unique even graph switching equivalent to Y.

Note: If there is a regular graph of even degree switching equivalent to Y, it must be Y_E. We omit the proof for space, but leave the following picture as a hint.

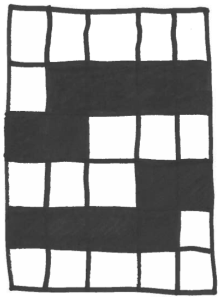

When n is odd, we have a correspondence of unordered switching classes and isomorphism classes of even graphs. It turns out that these sets are equinumerous even when n is even, though this fact is nontrivial, and a nice bijection is not known [MS75]. The number of nonisomorphic even graphs was calculated independently by Robinson and Liskovec. [MS75, Rob69]. We copy the following complicated formula out of [MS75].

Theorem 3.1.2 (Robinson-Liskovec formula)**.** The number of isomorphism classes of even graphs on n vertices is

$$e_n = \sum_{\sigma} \frac{2^{v(\sigma)-\lambda(\sigma)}}{\prod_i i^{\sigma_i}\sigma_i!}$$

where the sum ranges over integer n-tuples $(\sigma_1, \cdots, \sigma_n)$ such that $n = \sum_i i\sigma_i$ where

$$v(\sigma) = \sum_{i<j} \sigma_i\sigma_j \gcd(i,j) + \sum_i i\left(\sigma_{2i} + \sigma_{2i+1} + \binom{\sigma_i}{2}\right)$$

and $\lambda(\sigma) = \sum_i \sigma_i - \operatorname{sgn}\sum_i \sigma_{2i+1}$, and $\operatorname{sgn}(x) = 0$ if $x = 0$ and 1 if $x > 0$.

The first 21 values of e_n are listed in [MS75].

The complement of a switching class is again a switching class. Some unordered switching classes are self-complementary. Two (complementary) switching classes are the same in the ordered and unordered case, and are available for all n: the complete bipartite graphs, and the disjoint unions of two cliques. (A partition or clique might have size 0 or n).

In psychology, Heider's *balance theory* gives a somewhat coarse model describing attitudes a person P holds towards another person O and an object X given the attitude of O towards X. Each of these attitudes is *positive* or *negative.* The attitudes are collectively deemed *balanced* if 1 or 3 of the attitudes are positive, (which, ostensibly, can be rationalized on a case-by-case basis) [Har59]. Though the model seems to intend different roles of P, O, and X, and a "directedness" of the attitudes, the "balanced" condition ignores these. Thus we can represent a positive attitude as an (undirected) edge between two of the nodes P, O, and X, and a negative attitude as a lack thereof. The *balanced* and *unbalanced* regimes then correspond to the two unordered switching classes of graphs on three vertices.

Now consider a community (graph) Y of n people (nodes), with each pair of people having a mutually positive (edge) or negative (no edge) attitude toward each other. Then Y must be considered *balanced* if and only if in every *trio* (subgraph induced by three vertices), there are 1 or 3 positive attitudes. Equivalently, Y is switching equivalent to the complete graph, *i.e.* Y is a disjoint union of two cliques, representing two opposing factions. (One single clique is also balanced.) Mathematically, it would be interesting to know the "fate distributions" of graphs in other switching classes in a dynamical setting where attitudes between randomly selected unbalanced trios are resolved, say, by local complementation. Must they eventually coalesce into two cliques? Another variation is to resolve a randomly selected pair P, Q that is included in an unbalanced trio as follows: if PQ is an edge, replace the neighborhoods $N(P)$ and $N(Q)$ with either $N(P) \cap N(Q)$ (by deleting edges) or $N(P) \cup N(Q)$ (by adding edges), and do something similar when PQ is a nonedge.

Figure (3.1) depicts the following surjective function.

$$\{\text{iso. classes of graphs on } n \text{ vertices}\} \to \{\text{unordered switching classes on } n+1 \text{ vertices}\}$$
$$\mathcal{Y} \mapsto \delta(\mathcal{Y}^+)$$

where $\mathcal{Y}^+$ is the graph obtained from $\mathcal{Y}$ by adding an isolated vertex, and $\delta(\mathcal{Z})$ denotes the switching class of an isomorphism class of graphs $\mathcal{Z}$. The map is not generally injective, as seen in Figure (3.1). We will investigate generalizations of δ and of the function $\mathcal{Y} \mapsto \delta(\mathcal{Y}^+)$ in Section 3.4. We can also just as well consider a ordered version $Y \mapsto \delta(Y^+)$.

3.2 Combinatorial methods and real equiangular tight frames

In this section we analyze the case of real equiangular vectors, especially ETFs, using graph theory. We begin by reviewing the many equivalent structures we have encountered, and argue that the problems of testing graph isomorphism and unordered switching equivalence are polynomially equivalent. We also define switching graphs, strongly regular graphs, and state two formulations of the ETF property, proving one of them.

	ordered structures	unordered structures
type 1	· ordered config. of n rigid e.a. vecs. · rigid e.a. $n \times n$ Gram matrix · $n \times n$ Seidel matrix · graph on $[n]$	· unordrd. config. of n rigid e.a. vecs. · unordrd. rigid e.a. $n \times n$ Gram mat. · unordrd. $n \times n$ Seidel matrix · unordrd. graph on n vertices
type 2	· ordered config. of n rigid e.a. lines · switching class of graphs on $[n]$ · triple products of seq. of n e.a. vecs. · two-graph on n vertices · graph on $[n-1]$	· unordrd. config. of rigid e.a. lines · unordrd. switching class on n vertices · triple products of set of n e.a. vecs. · unordrd. two-graph on n vertices

Two-graph is defined in at least two ways in the literature: either as a synonym for *switching class*, or as the set of 3-sets of vertices such that every 4-set contains an even number of such 3-sets. We use the latter definition. The two-graph $\delta(Y)$ generated by a graph Y is the set of 3-sets that contain an even number of edges in Y. (See section 3.4).

The *switching neighborhoods* of a graph Y are the graphs $Y_{N(v)} - \{v\}$ obtained from Y by isolating and then deleting a vertex. The unordered type 2 structure analogous to a graph on $[n-1]$ is the n-multiset of switching neighborhoods of an isomorphism class of graphs $\mathcal{Y}$ on n vertices. This construction gives a polynomial time reduction from testing graph isomorphism to testing unordered switching equivalence. Let $\mathcal{N}(Y)$ denote the n-multiset of switching neighborhoods of a graph Y. Unordered switching equivalence of Y_1 and Y_2 is determined by *searching* for a bijection $f : \mathcal{N}(Y_1) \to \mathcal{N}(Y_2)$ such that for each $H \in \mathcal{N}(Y_1)$, $H \cong f(H)$. By the following observation, there is also a polynomial time reduction in the other direction.

Proposition 3.2.1. Let Y_1, Y_2 be graphs on $[n]$, and let Y_1^+, Y_2^+ be the graphs on $[n+1]$ obtained by adding an isolated vertex. Let $\simeq$ denote unordered switching equivalence and let $\cong$ denote isomorphism of graphs. Then

$$Y_1 \simeq Y_2 \implies (Y_1^+ \simeq Y_2^+ \text{ if and only if } Y_1 \cong Y_2)$$

Thus, to test whether graphs Y_1, Y_2 are isomorphic given an oracle that decides unordered switching equivalence, it suffices to ask whether $Y_1 \simeq Y_2$ and $Y_1^+ \simeq Y_2^+$. Therefore deciding graph isomorphism and unordered switching equivalence are polynomially equivalent. It is also not hard to show that the canonical labeling problems for graph isomorphism classes and unordered switching equivalence classes are polynomially equivalent.

Another construction available in the real case is the switching graph [GR01].

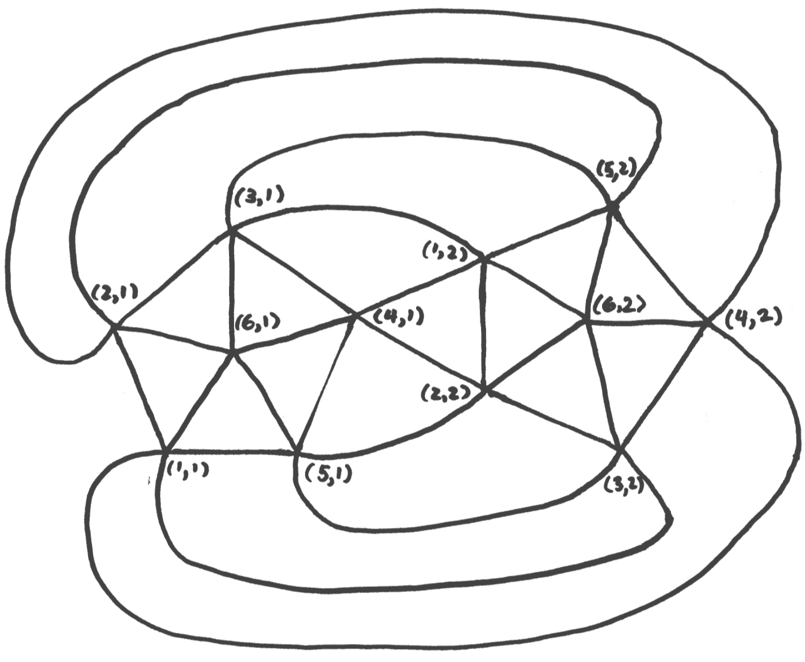

The switching graph of the icosahedral cap $(C_5 \sqcup K_1)_6$ is the icosahedral graph.

Definition 3.2.1. Let Y be a graph on $[n]$. The *switching graph* of Y, $\mathrm{Sw}(Y)$, is the graph on $[n] \times [2]$ with edge relation

$$\begin{aligned}(u,i) \sim (v,i) &\text{ iff } u \sim v \\ (u,i) \sim (v, 2-i) &\text{ iff } u \not\sim v \text{ and } u \neq v\end{aligned}$$

Up to certain isomorphisms, the switching graph is a type 2 structure. More precisely,

Proposition 3.2.2. Let Y and Y' be graphs on $[n]$. Then $Y' = Y_\theta$ for some $\theta \subseteq [n]$ if and only if $\mathrm{Sw}(Y') = f_\theta(\mathrm{Sw}(Y))$, where f_θ is the permutation of $[n] \times [2]$ transposing $(v,1)$ with $(v,2)$ for each $v \in \theta$ and leaving every other element fixed.

The map $\mathrm{Sw}(Y) \to K_n$ given by $(u,i) \mapsto u$ is a double cover, that is, it is a graph homomorphism and each vertex of K_n is mapped by two non-adjacent vertices in $\mathrm{Sw}(Y)$. Let A be the (ordinary) adjacency matrix of Y and let $\overline{A}$ be that of its complement. (The adjacency matrix relates to the Seidel matrix X of Y by $X = A - \overline{A}$ and $A = \frac{1}{2}(X + J - I)$.) Then the adjacency matrix of $\mathrm{Sw}(Y)$ is

$$\begin{pmatrix} A & \overline{A} \\ \overline{A} & A \end{pmatrix}$$

Using the property

$$\det \begin{pmatrix} A & B \\ B & A \end{pmatrix} = \det(A - B)\det(A + B) \tag{3.2}$$

and the relations $X = A - \overline{A}$ and $A + \overline{A} = J - I$, the eigenvalues of $\mathrm{Sw}(Y)$ are the eigenvalues $\lambda_1, \cdots, \lambda_n$ of X together with $n-1$, of multiplicity 1, and -1, of multiplicity $n-1$. The eigenvalues of X are called the *nontrivial eigenvalues* of $\mathrm{Sw}(Y)$. A switching graph is *regular* if it has exactly two nontrivial eigenvalues. (Recall that this means Y corresponds to a tight configuration of equiangular vectors.) In order to state the characterization theorem for for regular switching graphs, we first need some definitions from graph theory.

Definition 3.2.2. The *distance* between two vertices in a graph is the length of a shortest path between them. The *diameter* of a graph is the greatest distance between any two vertices. A graph is *distance regular* if for any vertices u, v, the number of vertices at distance i from u and distance j from v depends only on i and j. A graph of diameter d is *antipodal* if the vertices at distance d from any vertex are at distance d from each other.

Theorem 3.2.1 (Godsil and Hensel)**.** The switching graph of a graph Y on $[n]$ is regular if and only if it is an antipodal distance regular double cover of K_n.

The reader is refered to [GH92] for the proof.

Another class of structures closely related to ETFs is that of strongly regular graphs (SRGs). We can use them to state yet another equivalent formulation of the ETF property.

Definition 3.2.3. A graph Z is *strongly regular* with parameters (n, k, a, c), (also expressed "Z is an SRG(n, k, a, c)") if Z has n vertices, is k-regular, any two adjacent vertices have a common neighbors, and any two non-adjacent vertices have c common neighbors.

It is easy to show that the parameters of an SRG(n, k, a, c) satisfy the *SRG relation*

$$k(k - a - 1) = (n - k - 1)c \tag{3.3}$$

Strong regularity can also be expressed algebraically.

Proposition 3.2.3. Let Z be a graph on $[n]$ with (ordinary) adjacency matrix A. The adjacency matrix of the complement $\overline{Z}$ is $\overline{A} = J - I - A$. Then Z is strongly regular if and only if A^2 is a linear combination of A, I, and $\overline{A}$ (equivalently of A, I, and J), and the parameters (n, k, a, c) are determined by this linear combination.

Proof. The ijth entry of A^2 is the number of walks of length 2 from i to j. Then Z is an SRG(n, k, a, c) if and only if there are k such walks from any vertex back to itself, there are a such walks between adjacent vertices, and there are c such walks between non-adjacent vertices. Equivalently,

$$A^2 = kI + aA + c\overline{A} \tag{3.4}$$

Note also that the parameters are determined by such an expression of A^2. This is evident if Z is complete or empty, and otherwise true since I, A, and $\overline{A}$ are linearly independent. ■

Proposition 3.2.4. A connected k-regular graph is strongly regular if and only if its adjacency matrix has at most three eigenvalues.

Proof. Let Z be a k-regular graph on $[n]$ with adjacency matrix A. If Z is strongly regular, then it satisfies the following equation, equivalent to (3.4).

$$A^2 + (c-a)A + (c-k)I = cJ \tag{3.5}$$

Suppose v is an eigenvector for an eigenvalue other than $\theta \neq k$. A is normal so its eigenspaces are orthogonal, so $v \perp j$. Multiplying (3.5) on the right by v we see that θ is a root of the quadratic polynomial

$$x^2 + (c-a)x + (c-k) = 0 \tag{3.6}$$

Conversely, suppose Z is connected and A has eigenvalues k, θ, τ. Since Z is connected, j is the only eigenvector (up to scalar) for eigenvalue k. Since A is diagonalizable, $\{j\}$ can be completed to an eigenbasis $\mathcal{E}$ for A. The following trick is taken from [GR01]. Consider the matrix

$$M = \frac{1}{(k-\theta)(k-\tau)}(A-\theta I)(A-\tau I)$$

Then clearly $Me = 0$ for all $e \in \mathcal{E} \setminus \{j\}$ and $Mj = 1$. This determines that $M = J$, so A^2 is a linear combination of A, I, and J. ■

The eigenvalues of an SRG and their multiplicities can be determined via Equation (3.6), and clearly depend only on the parameters (n, k, a, c). The integrality of these multiplicities can be used to rule out the existence of SRGs with certain parameters, though the existence of such an SRG can depend on purely combinatorial considerations. Deciding the existence of an SRG with certain parameters is typically difficult. The parameters of an SRG are obviously invariant under isomorphism, however, there exist SRGs with the same parameters that are not isomorphic (or even unordered switching equivalent). Famous examples include the Shrikhande graph and the 4 × 4 rook's graph, or the Chang graphs. See [Wal09, zotc, zotd] for lists of known strongly regular graphs. We now show how SRGs relate to ETFs.

Theorem 3.2.2. Let Y be a graph on $[n]$ and let $1 \leq i \leq n$. Let $Y_{(i)} = Y_{N(i)\cup\{i\}} - \{i\}$ be the graph obtained by *switching on* vertex i. Then Y is the graph of an ETF if and only if $Y_{(i)}$ is an SRG$(n-1, k, a, c)$ for some k, a, and c such that $k = 2c$.

Proof. Without loss of generality switch on vertex 1. Let X be the Seidel matrix of Y and let $X_{(1)}$ be the Seidel matrix of $Y_{(1)}$. Y is the graph of an ETF if and only if there exist

λ_1, λ_2 such that

$$0 = X^2 - (\lambda_1 + \lambda_2)X + \lambda_1\lambda_2 I$$

Expanding this out we obtain

$$0 = \begin{pmatrix} n-1 & j^T X_{(1)} \\ X_{(1)}j & J + X_{(1)}^2 \end{pmatrix} - (\lambda_1 + \lambda_2)\begin{pmatrix} 0 & j^T \\ j & X_{(1)} \end{pmatrix} + \lambda_1\lambda_2 \begin{pmatrix} 1 & 0 \\ 0 & I \end{pmatrix}$$

By equating the blocks, we see that Y is an ETF if and only if

1. $\lambda_1\lambda_2 = -(n-1)$,
2. $X_{(1)}j = (\lambda_1 + \lambda_2)j$, that is, $Y_{(1)}$ is k-regular with $-n + 2k + 2 = \lambda_1 + \lambda_2$, and
3. $0 = J + X_{(1)}^2 - (\lambda_1 + \lambda_2)X_{(1)} + \lambda_1\lambda_2 I$.

Since $Y_{(1)}$ is k-regular, $A_{(1)}J = kJ$. Using the relations $X_{(1)} = 2A_{(1)} - J + I$ and $J = I + A_{(1)} + \overline{A}_{(1)}$, and substituting $-(n-1) = \lambda_1\lambda_2$ and $-n+2k+2 = \lambda_1 + \lambda_2$, the equation in item 3 above simplifies to

$$A_{(1)}^2 = kI + \frac{3k-n}{2}A_{(1)} + \frac{k}{2}\overline{A}_{(1)}$$

Thus item 3 is equivalent to $Y_{(1)}$ being an SRG$(n-1, k, \frac{3k-n}{2}, \frac{k}{2})$. This condition implies item 2 above (k-regularity). It also implies item 1. This is evident when Y is switching equivalent to the empty graph or to the complete graph, and otherwise, since $\{J, X_{(1)}, I\}$ is linearly independent, $\lambda_1\lambda_2$ is determined by $X_{(1)}^2$ via expression 3 above. Finally, note that the value $a = \frac{3k-n}{2}$ is determined by the other parameters n, k, and c by the SRG relation (3.3). This completes the proof. ■

Interestingly, Theorem (3.2.2) implies that if Z is an SRG on n vertices with $k = 2c$, then for any $v \in V(Z)$, $Z' = (Z \sqcup K_1)_{N(v)} \setminus \{v\}$ is again such. Waldron claims that it may be that Z and Z' are not isomorphic, and that this method has been used to discover new such SRGs [Wal09]. However, if n is even, then $Z \sqcup K_1$ is the unique even graph in its switching class, so in this case, $Z \cong Z'$.

We now show how the parameters n and d of an ETF relate to the parameters n and k of the associated SRGs, (i.e. the switching neighborhoods).

Corollary 3.2.2.1. Consider an ETF of n vectors in $\mathbb{R}^d$, and an associated SRG with parameters n and k. Then n, k, and d relate by the following formula..

$$k = \frac{1}{2}\left(n - 2 - \sqrt{\frac{d(n-1)}{n-d}} + \sqrt{\frac{(n-d)(n-1)}{d}}\right)$$

The integrality of k can be used with this formula to rule out certain configurations, though a stronger (and easy-to-prove) criterion will be discussed later.

Proof. Recall that the eigenvalues of an ETF of n vectors in $\mathbb{R}^d$ are

$$\lambda_1 = -\sqrt{\frac{d(n-1)}{n-d}} \qquad\qquad \lambda_2 = \sqrt{\frac{(n-d)(n-1)}{d}}$$

Plug these into $k = \frac{1}{2}(n - 2 + \lambda_1 + \lambda_2)$. ■

3.3 The pillars of Lemmens and Seidel

In this section we review Lemmens and Seidel's pillar method [LS73], especially their treatment of the $\alpha = 1/3$ case, in order to find the maximum number $N(S^{d-1})$ of real equiangular vectors (equivalently lines) in $\mathbb{R}^d$ up to dimension $d = 13$. A number theoretic fact available only in the real case implies that it is sufficient to find the maximum number $N(S^{d-1}, \pm\frac{1}{3})$ of equiangular lines at angle $\alpha = \frac{1}{3}$ in these dimensions. The basic idea is to start with a coclique in a fixed dimension d and consider what vectors can be added to the configuration. The fact that real equiangular vectors only have two possible inner products, $\pm\alpha$, implies that the projection onto the span of the coclique takes finitely many values h. The cosets $h + \Gamma$, up to multiplication by -1, called pillars, can be enumerated, and this leads to improved bounds on $N(S^{d-1}, \pm\alpha)$.

There have been many new developments since [LS73] that we do not have space to explain in detail. In the real case, an inequality, which we call the *coclique bound*, can be found by dragging a regular simplex of $d + 1$ vertices in $\mathbb{R}^d$ orthogonally into the $(d + 1)$th dimension. A newer strategy than that of Lemmens and Seidel is to use the coclique bound in conjunction with Ramsey's Theorem to build onto a clique, rather than onto coclique [BDKS16]. Intuitively it makes sense for Ramsey's Theorem to arise, because one can show that most of the brute force computation for finding $N(S^{d-1})$ for $d = 4, 5, 6, 10, 12$ (the nontrivial cases) can be diverted into the verification of the Ramsey numbers $R(3, 3) = 6$, $R(3, 4) = 9$, $R(4, 4) = 18$, and $R(4, 5) = 25$. I have omitted these examples to save space. In 2019, Jiang et. al. claimed to have solved the problem of determining $N(S^{d-1}, \pm\alpha)$ given d and α, though their result is asymptotic [JTY+19]. They also used Ramsey's theorem. Theorem (3.3.2) below shows why their result can be used to find $N(S^{d-1})$.

We first prove a number theoretic fact with profound consequences about the real case.

Theorem 3.3.1. For $n > 2d$ equiangular lines in $\mathbb{R}^d$ with angle α, $1/\alpha$ is an odd integer.

Proof. $\lambda = -1/\alpha$ is a root of the characteristic polynomial of the Seidel matrix X with multiplicity $n - d > 1$. Since X is an integer matrix, if λ is not an integer, then its conjugate is also a root with multiplicity $n - d$. This is impossible since the sum of these multiplicities is $2(n - d) > n$, so λ is an integer. Consider the (ordinary) adjacency matrix $A = \frac{1}{2}(X + J - I)$. The matrix J has eigenvalue 0 with multiplicity $n - 1$. Since λ has multiplicity > 1, X has a λ-eigenvector that is a 0-eigenvector of J, and so $\mu = \frac{1}{2}(\lambda - 1)$ is an eigenvalue of A. Since μ is rational and A is an integer matrix, μ is an integer. ■

Corollary 3.3.1.1. If $\mathcal{L}$ is an equiangular set in $\mathbb{R}^d$ attaining the absolute bound, then $d = 2$, $d = 3$, or d is two less than an odd square.

Theorem (3.3.1) and the relative bound imply that $N(S^{d-1})$ depends on $N(S^{d-1}, \pm\alpha)$ for finitely many odd integers $1/\alpha > 1$. Explicitly, Lemmens and Seidel write:

Theorem 3.3.2. Suppose $5 < d < (2m + 1)^2$ for some integer $m > 0$. Then

$$N(S^{d-1}) \leq \max\left\{N(S^{d-1}, \pm\tfrac{1}{3}), N(S^{d-1}, \pm\tfrac{1}{5}), \cdots, N(S^{d-1}, \pm\tfrac{1}{2m-1}), \frac{4dm(m+1)}{(2m+1)^2 - d}\right\}$$

Before we prove it, let us note that the maximum number $N(S^{13})$ of equiangular lines in $\mathbb{R}^{14}$ is known only to be either 28 or 29. We prove the theorem in either case. It is easy to show that there is no ETF in $\mathbb{R}^{14}$, and if 29 equiangular lines exist, they must have angle $1/5$. There are known configurations of 28 equiangular lines in $\mathbb{R}^{14}$ at angle $1/5$ [GKMS16, Tre08], and it is tractable to check via computer whether a 29th line can be added without increasing the dimension. I have verified this in Sage for the cited example.

Proof. If $N(S^{13}) = 28$, (which is an open question,) the assertion holds because there exist 28 equiangular lines (in $\mathbb{R}^7$, even) at angle $1/3$. Let $5 < d < (2m + 1)^2$. Constructions exist [LS73] showing that $N(S^{d-1}) > 2d$ if $d \neq 2, 3, 4, 14$, and we assume additionally that $N(S^{13}) > 28$. Thus by Theorem (3.3.1), $N(S^{d-1}) = N(S^{d-1}, \pm\alpha)$ for some odd integer $1/\alpha > 1$. Now suppose $1/\alpha \geq 2m+1$. Then by the relative bound,

$$N(S^{d-1}, \pm\alpha) \leq \frac{d(1-\alpha^2)}{1-d\alpha^2} \leq \frac{4dm(m+1)}{(2m+1)^2 - d}$$

$\blacksquare$

We already have the tools to calculate $N(S^{d-1})$ in dimensions $2, \cdots, 7$. The above theorem also holds in dimensions 3 and 4, which we will soon see. We have already dealt with dimensions $2, 3$, and 7. We will now show the tight upper bounds for dimensions 4,5 and 6. Lower bounds come from explicit constructions.

Theorem 3.3.3. The maximum number $N(S^3)$ of equiangular lines in $\mathbb{R}^4$ is 6.

Proof. Theorem (3.3.1) cannot be used (in conjunction with the relative bound) because if $d = 4$ and $n = 7$, then n is not greater than $2d$. Since the number of vertices is low enough, it is feasible to do an exhaustive search. The following Sage program iterates over one constituent of each switching class of graphs on 7 vertices, namely, the graph where the final vertex is isolated. The dimension d is then calculated for each graph. Caution: the graph switching equivalent to K_7 will register as 1-dimensional.

```
flag = true
for G in graphs(6):
    G = G.disjoint_union(graphs.CompleteGraph(1))
    spectrum = G.seidel_adjacency_matrix().eigenvalues()
    codim = spectrum.count(min(spectrum))
    d = 7 - codim
    if d < 5 and d != 1:
        flag = false
print flag
```

The program returns `True`, indicating that $N(S^3) = 6$. $\blacksquare$

The maximum number of real equiangular lines in dimensions 5 and 6 are a consequence of theorem (3.3.1) and the relative bound.

Theorem 3.3.4. $N(S^4) = 10$ and $N(S^5) = 16$.

Proof. Plugging the values $d = 5, 6$ and $n = 11, 17$, respectively, into the relative bound yields $\alpha > 1/3$, contradicting theorem (3.3.1). ■

In dimensions 8-13 the relative bound and theorem (3.3.1) imply that $\alpha = 1/3$. Thus to compute $N(S^{d-1})$ for these dimensions d it suffices to compute $N(S^{d-1}, \pm\frac{1}{3})$. Thus we introduce the pillars of Lemmens and Seidel.

Definition 3.3.1. Let $\mathcal{L}$ be a set of m unit vectors in $\mathbb{R}^n$ such that for any distinct $v, v' \in \mathcal{L}$, $\langle v, v' \rangle = -\alpha = -1/(m-1)$. That is, $\mathcal{L}$ is a regular unit simplex centered at the origin spanning an $m-1$ dimensional subspace. Let $P_{\mathcal{L}}$ be the orthogonal projector onto $\operatorname{span}\mathcal{L}$, and $P_{\mathcal{L}}^{\perp}$ be the orthogonal projector onto $(\operatorname{span}\mathcal{L})^{\perp}$. If $x \in \mathbb{R}^n$ is another vector such that $\mathcal{L} \cup \{x\}$ is equiangular, and $h = P_{\mathcal{L}}(x)$, then $h + P_{\mathcal{L}}^{\perp}$ is called the *pillar* of h.

It turns out any interesting configuration of real equiangular vectors at angle 1/3 can be obtained by building onto a unit regular tetrahedron spanning a 3 dimensional subspace.

Proposition 3.3.1. Let $\alpha = 1/3$ and $n \geq 4$. Consider n equiangular vectors at angle α with Seidel matrix X with $n > d(X)$. Then X has an induced subgraph switching equivalent to $\overline{K}_4$, or there is another graph X' with X as an induced subgraph such that $d(X') = d(X)$.

Proof. Assume X has no induced subgraph switching equivalent to $\overline{K}_4$. Since $n > d(X)$, X is not switching equivalent to K_n. So there exist three vertices u, v, w such that $Y = X[\{u, v, w\}]$ is switching equivalent to $\overline{K}_3$. Let u, v, w also denote the corresponding vectors, and let $x = u + v + w$. Then $\langle x, u \rangle = \langle x, v \rangle = \langle x, w \rangle = \frac{1}{3}$. If y is another vertex of the graph X (and y also denotes the corresponding vector), then $\langle x, y \rangle = \langle u, y \rangle + \langle v, y \rangle + \langle w, y \rangle = \pm\frac{1}{3}$ or ± 1. If $\langle x, y \rangle = \pm 1$ then y is adjacent or nonadjacent to all three of u, v, w, which contradicts our assumption that X contains no induced subgraph switching equivalent to $\overline{K}_4$. Therefore the set of vectors obtained by adding x is equiangular, and of rank no higher than $d(X)$. ■

Consider a set $\mathcal{L}$ of unit vectors in $\mathbb{R}^n$ with inner products of distinct vectors equal to $-1/(m-1)$ and a vector x such that $\mathcal{L} \cup \{x\}$ is equiangular (with angle $\alpha = 1/(m-1)$). Let $h = P_{\mathcal{L}}(x)$ and $c = P_{\mathcal{L}}^{\perp}(x)$. Then $\langle x, v_i \rangle = \epsilon_i \alpha$ with $\epsilon_i = \pm 1$ for each i. Recall that $\sum v_i = 0$. It follows that $\epsilon_1 + \cdots + \epsilon_m = 0$, so m is even and half of the ϵ_i are $+1$ and half are -1. (When m is odd, $\mathcal{L}$ is a maximally equiangular set of vectors! If a graph has a proper induced subgraph switching equivalent to $\overline{K}_t$ with t odd, then for any odd $s \leq t$, the angle is not $1/(t-1)$.)

We may express h as $\sum c_i v_i$. To find c_i, take the inner product with v_i.

$$\frac{\epsilon_i}{m-1} = \langle x, v_i \rangle = \langle h, v_i \rangle = \frac{m}{m-1} c_i - \frac{1}{m-1} \sum_{j=1}^{m} c_j$$

Now clearing out the denominator and summing over i on the leftmost and rightmost expressions yields

$$0 = \sum \epsilon_i = \sum c_i + m \sum c_i = (1+m) \sum c_i$$

from which it follows that $\sum c_i = 0$ and so

$$h = \frac{1}{m} \sum \epsilon_i v_i \tag{3.7}$$

Since the ϵ_i determine h, there are $\binom{m}{\frac{1}{2}m}$ pillars. We are only concerned with finding equiangular lines, as opposed to vectors, so we work modulo the sign of $\epsilon = (\epsilon_1, \cdots, \epsilon_m)$, leaving half of the pillars in consideration. Now suppose instead we add vectors $x_1, \cdots, x_s$ such that $\mathcal{L} \cup \{x_1, \cdots, x_s\}$ is equiangular. Let $h_a = P_{\mathcal{L}}(x_a)$ and $\langle x_a, v_i \rangle = \epsilon_i^{(a)} \alpha$ for each $a = 1, \cdots, s$ and $i = 1, \cdots, m$. Let $E = [\epsilon^{(1)} \cdots \epsilon^{(s)}]$ where $\epsilon^{(a)} = (\epsilon_1^{(a)}, \cdots, \epsilon_m^{(a)})$ for each a. It turns out that $\frac{1}{m(m-1)} E^T E$ is the Gram matrix of $[h_1 \cdots h_s]$, which we now prove.

$$\langle h_a, h_b \rangle = \frac{1}{m^2} \sum_{ij} \epsilon_i^{(a)} \epsilon_j^{(b)} \langle v_i, v_j \rangle = \frac{1}{m^2} \left(\frac{m}{m-1} \sum_i \epsilon_i^{(a)} \epsilon_i^{(b)} - \frac{1}{m-1} \sum_{ij} \epsilon_i^{(a)} \epsilon_j^{(b)} \right)$$

Please note that the sum over i and j on the right-hand-side includes the diagonal terms. Now recall that $\sum_i \epsilon_i^{(a)} = 0$ (or $\sum_i \epsilon_i^{(b)} = 0$), so $\sum_{ij} \epsilon_i^{(a)} \epsilon_j^{(b)} = 0$, and therefore

$$\langle h_a, h_b \rangle = \frac{1}{m(m-1)} \sum_i \epsilon_i^{(a)} \epsilon_i^{(b)} = \frac{1}{m(m-1)} \langle \epsilon^{(a)}, \epsilon^{(b)} \rangle = \frac{1}{m(m-1)} (E^T E)_{ab}$$

Lemmens and Seidel next prove the following theorem.

Theorem 3.3.5. Let $L = [v_1 \cdots v_m]$ be a sequence of unit vectors in $\mathbb{R}^n$ with common inner products $-\alpha = -1/(m-1)$. There exist $m+s$ equiangular vectors $v_1, \cdots, v_m, x_1, \cdots, x_s$ spanning $\mathbb{R}^n$ if and only if there exist an $s \times s$ Seidel matrix A and an $m \times s$ matrix E of

± 1s with vertical sum 0 such that the matrix

$$C = I + \frac{1}{m-1}A - \frac{1}{m(m-1)}E^T E \tag{3.8}$$

is positive semidefinite of rank $n - m + 1$.

Proof. $(\implies)$: Let $E = [\epsilon^{(1)} \cdots \epsilon^{(s)}]$ be the matrix whose columns $\epsilon^{(j)}$ are determined by $\langle v_i, x_j \rangle = \epsilon_i^{(j)} \alpha$. The row sum of E is 0. Let S be the Seidel matrix of L and let A be the Seidel matrix of $X = [x_1 \cdots x_s]$. Then the Seidel matrix of $[v_1 \cdots v_n \, x_1 \cdots x_s]$ is

$$S' = \begin{pmatrix} S & E \\ E^T & A \end{pmatrix}$$

S' has rank n. The Gram matrix for X, $I + \frac{1}{m-1}A$, is the sum of the Gram matrix C for $[c_1 \cdots c_s]$ and the gram matrix H for $[h_1 \cdots h_s]$. As noted above, $H = \frac{1}{m(m-1)}E^T E$, so equation (3.8) holds. Since $\operatorname{span}\{v_1, \cdots, v_m, x_1, \cdots, x_s\} = \mathbb{R}^n$ and $\dim \operatorname{span}\{v_1, \cdots, v_m\} = m - 1$, C has rank $n - (m-1)$.

$(\impliedby)$: By proposition (2.2.2), C is the the Gram matrix of a sequence $[c_1 \cdots c_s]$ of vectors spanning an $n - (m-1)$ dimensional subspace V of $\mathbb{R}^n$. Let $\{v_1, \cdots, v_m\}$ be a regular unit simplex centered at the origin of $V^\perp$. Let $[h_1 \cdots h_s]$ be the vectors in $V^\perp$ defined by (3.7). Since equation (3.8) holds, A is the Seidel matrix of $[x_1 \cdots x_s]$ where $x_i = c_i + h_i$ for each i. Therefore S' (as above) is the Seidel matrix of $[v_1 \cdots v_m \, x_1 \cdots x_s]$. If S' has rank n' then by $(\implies)$, C has rank $n' - m + 1$, so S' has rank n. ■

To see how many vectors x can be in one pillar, we let $E = [\epsilon \cdots \epsilon]$ have all columns equal to a vector ϵ of ± 1 with vertical sum 1. Then $E^T E = mJ$, and the Gram matrix for $[c_1 \cdots c_s]$ is

$$C = \frac{m-2}{m-1}I - \frac{1}{m-1}(J - I - A)$$

The vectors c_i have length $\sqrt{(m-2)/(m-1)}$ and mutual inner products 0 and $-2/(m-1)$. After renormalizing, their mutual inner products are 0 and $-2/(m-2)$. Since m is even, to bound the maximum number of vectors in the pillar we bound the maximum number of unit vectors in $\mathbb{R}^n$ with mutual inner products 0 and $-1/r$ for a positive integer r.

Theorem 3.3.6. $N(S^{n-1}, \{0, -\frac{1}{r}\}) = \lfloor \frac{n}{r} \rfloor + n$.

The upper bound follows from the linear programming bounds discussed in the next chapter. Lemmens and Seidel prove this bound in a different way, using Perron-Frobenius theory, and at the same time prove that this is also a lower bound.

A better bound (constant with respect to the dimension) can be found when multiple pillars are occupied. Consider the case of two occupied pillars:

$$E = \Big[\underbrace{\epsilon^{(1)} \cdots \epsilon^{(1)}}_{s} \underbrace{\epsilon^{(2)} \cdots \epsilon^{(2)}}_{t}\Big] \Big\}m$$

Let l be the Hamming distance of $\epsilon^{(1)}$ and $\epsilon^{(2)}$, (the number of unequal coordinates). Without loss of generality assume $m \geq 2l$; (otherwise replace $x_1, \cdots, x_s$ with $-x_1, \cdots, -x_s$, which has the effect of replacing $\epsilon^{(1)}$ with $-\epsilon^{(1)}$ and replacing l with $m - l$). Then

$$E^T E = \begin{pmatrix} mJ & (m-2l)J \\ (m-2l)J^T & mJ \end{pmatrix}$$

The (ordinary) adjacency matrix for the complement graph determined by the sequence of equiangular vectors $[x_1, \cdots, x_s, y_1, \cdots, y_t]$ is

$$\frac{1}{2}(J - I - A) = \begin{pmatrix} B_1 & B_2 \\ B_2^T & B_3 \end{pmatrix}$$

where B_1 is $s \times s$ and B_3 is $t \times t$. This simplifies to

$$\frac{m-1}{2} C = \begin{pmatrix} (\frac{1}{2}m - 1)I - B_1 & \frac{l}{m}J - B_2 \\ \frac{l}{m}J^T - B_2^T & (\frac{1}{2}m - 1)I - B_3 \end{pmatrix} \tag{3.9}$$

We now consider the case that multiple pillars are occupied. Using theorem (3.3.5) and the above decomposition, we rederive a bound on the number of vectors x in each pillar. This contrasts to the case in which all of the added vectors belong to the same pillar.

Theorem 3.3.7. Let m be even, and let $v_1, \cdots, v_m$ be unit vectors where for any distinct v_i, v_j, $\langle v_i, v_j \rangle = -1/(m-1)$. Further assume that $v_1, \cdots, v_m, x_1, \cdots, x_s, y$ are equiangular, and $x_1, \cdots, x_s$ belong to one pillar $h_1 + \Gamma$, and y belongs to a different pillar $h_2 + \Gamma$. Let the Hamming distance of the corresponding vectors $e^{(1)}$ and $e^{(2)}$ be l. Without loss of generality assume that $0 < l \leq \frac{1}{2}m$. Then

$$s \leq \frac{m^2(m-2)^2}{4l^2}$$

Please be careful when counting to consider the vectors in $h + \Gamma$ as well as $-h + \Gamma$.

Proof. We obtain an expression for C from (3.9), substituting in $B_1 = 0$, $B_2 = b$ (an $s \times 1$ column of 0s and 1s), and $B_3 = 0$.

$$\frac{m-1}{2}C = \begin{pmatrix} (\frac{1}{2}m-1)I & \frac{l}{m}j - b \\ \frac{l}{m}j^T - b^T & \frac{1}{2}m - 1 \end{pmatrix}$$

where j is the $s \times 1$ column of all 1s. Schur's determinant identity reads

$$\det \begin{pmatrix} A_1 & A_2 \\ A_3 & A_4 \end{pmatrix} = \det(A_4)\det\bigl(A_1 - A_2A_4^{-1}A_3\bigr)$$

In the case that A_4 is 1×1 this simplifies to $\det(A_1)\det(A_4) - \det(A_2)\det(A_3)$. Since C is positive semidefinite, its determinant is non-negative. So

$$\begin{aligned} 0 &\leq \left(\frac{1}{2}m - 1\right)^2 - \left(\frac{l}{m}j^T - b^T\right)\left(\frac{l}{m}j - b\right) \\ &= \left(\frac{1}{2}m - 1\right)^2 - \frac{l^2}{m^2}j^Tj + \frac{l}{m}j^Tb + \frac{l}{m}b^Tj - b^Tb \end{aligned}$$

Note that $j^Tj = s$ and $j^Tb = b^Tj = b^Tb$ is the number of 1s in b, (which is non-negative). So by our assumption that $l \leq \frac{1}{2}m$,

$$0 \leq \left(\frac{1}{2}m - 1\right)^2 - \frac{l^2}{m^2}s - \frac{m-2l}{m}b^Tb \leq \left(\frac{1}{2}m - 1\right)^2 - \frac{l^2}{m^2}s$$

■

Theorem 3.3.8. (The reward of all of these calculations is the following.)

d	$N(S^{d-1})$	$N(S^{d-1}, \pm\frac{1}{3})$
7	28	28
8	28	28
9	28	28
10	28	28
11	28	28
12	28	28
13	28	28
14	?	28
$d > 14$	$*$	$2d - 2$

The ? indicates that $N(S^{13})$ is unknown, and the $*$ indicates that various other, but not all values of $N(S^{d-1})$ are known for $d > 14$.

Proof. We have already seen that the relative bound implies that $N(S^{d-1}) = N(S^{d-1}, \pm\frac{1}{3})$ for $d = 7, \cdots, 13$. We have also aluded to constructions of the indicated number of equiangular lines. So it suffices to upper bound $N(S^{d-1}, \pm\frac{1}{3})$ for each d by the values indicated.

By Proposition (3.3.1), we consider pillars for a coclique of size $m = 4$. There are three pillars, each at hamming distance 2 from the others, corresponding to the column vectors

$$\epsilon^{(1)} = (+, +, -, -), \epsilon^{(2)} = (+, -, +, -), \epsilon^{(3)} = (+, -, -, +)$$

By Theorem (3.3.5), the remaining s vectors correspond to an $s \times s$ Seidel matrix A, and an $m \times s$ matrix E of 0s and 1s with vertical sum 0, such that $C = I + \frac{1}{3}A - \frac{1}{12}E^T E$ is positive semidefinite of rank $d-3$. If more than one of the pillars are occupied, by Theorem (3.3.7), each pillar contains at most 8 vectors, for a total of at most $4 + 3 \times 8 = 28$ vectors. If only one pillar is used, Theorem (3.3.6) implies that the pillar contains at most $2d - 6$ vectors. Adding to the four original vectors, this gives an upper bound of $2d - 2$ vectors. ■

Lemmens and Seidel also use pillars to analyze the $\alpha = 1/5$ case [LS73], though it becomes more complicated. This case was also studied by Neumaier [Neu89].

3.4 Co-chains of graph-like structures and Boolean representations of $\mathfrak{sl}(2, \mathbb{C})$

We have seen that rigid configurations of equiangular lines correspond to switching classes of graphs on $[n]$. These structures also correspond to two-graphs on $[n]$. A *two-graph* on $[n]$ is a set of 3-subsets of $[n]$ such that any 4-subset of $[n]$ contains an even number of these 3-sets. Let Y be a graph on $[n]$. The *two-graph generated by* Y, $\delta(Y)$, is the set of 3-sets $S \subseteq [n]$ such that the induced subgraph $Y[S]$ has an even number of edges. The proof that $\delta(Y)$ is a two-graph is done locally by checking that all graphs on vertex set $[4]$ contain an even number of 3-sets that contain an odd number of edges. (It is easy to check the 11 isomorphism classes of graphs.) It can also be checked locally that graphs Y_1 and Y_2 are (ordered) switching equivalent if and only if $\delta(Y_1) = \delta(Y_2)$. It turns out that every 2-graph has the form $\delta(Y)$ for some Y, and this construction generalizes quite a lot. We consider two co-chains that generalize this concept, and show that there is an $\mathfrak{sl}(2, \mathbb{C})$ action that reveals a certain enumerative property about the unordered structures.

For a fixed $n \in \mathbb{N}$, and for each $k = 0, \cdots, n$, let $V_k^n = V_k$ be the $\mathbb{F}_2$-vector space

$$V_k = \left\{ f : \binom{[n]}{k} \to \mathbb{F}_2 \right\}$$

where $\binom{[n]}{k}$ is the set of k-subsets of $[n]$. The coboundary maps are defined

$$\begin{aligned} &\delta_k : V_k \to V_{k+1} \\ &\delta_k(f)(T) = \sum_{t \in T} f(T \smallsetminus t) \end{aligned}$$

The first few vector spaces can be viewed as

$$\begin{aligned} V_1 &= \{\text{subsets of } [n]\} \\ V_2 &= \{\text{graphs on } [n]\} \\ V_3 &= \{\text{sets of 3-sets on } [n]\} \\ &\vdots \end{aligned}$$

Theorem 3.4.1. The sequence

$$V_1 \xrightarrow{\delta_1} V_2 \xrightarrow{\delta_2} V_3 \xrightarrow{\delta_3} \cdots$$

is exact.

This generalizes the characterization of complete bipartite graphs as graphs with the property that any 3-set of vertices has either 0 or 2 edges, which is the statement $\ker \delta_2 = \operatorname{im} \delta_1$. It also generalizes the statement that two-graphs are precisely the sets of 3-sets on $[n]$ of the form $\delta(Y)$ for some graph Y, which is the statement $\ker \delta_3 = \operatorname{im} \delta_2$.

Proof. To see that $\operatorname{im} \delta_{k-1} \subseteq \ker \delta_k$, observe that

$$\begin{aligned} \delta_k \delta_{k-1}(f)(T) &= \sum_{t \in T} \delta_{k-1}(f)(T \setminus t) \\ &= \sum_{t \in T} \sum_{s \in T \setminus t} f(T \setminus t \setminus s) \\ &= \sum_{\{s,t\} \subseteq T,\ s \neq t} (1+1) f(T \setminus t \setminus s) \\ &= 0 \end{aligned}$$

We establish $\operatorname{im} \delta_{k-1} \supseteq \ker \delta_k$ by dimension counting. View δ_k as the $\binom{[n]}{k+1} \times \binom{[n]}{k}$ matrix with entry $(\text{row}, \text{col}) = (T, S)$ equal to 1 if $S \subseteq T$ and 0 otherwise. These matrices, (technically their transposes,) and their generalizations have been investigated, by P. Frankl, Linial and Rothschild, and Wilson [LR81, Wil90]. Frankl proved that the rank of δ_k is

$$\binom{n-1}{k}$$

Therefore

$$\operatorname{rk} \delta_{k-1} + \operatorname{rk} \delta_k = \binom{n-1}{k-1} + \binom{n-1}{k} = \binom{n}{k} = \dim V_k$$

So the result follows by the rank-nullity theorem. ■

Other variants of this construction exist, and it would be nice to formulate it in such a way that the well-established machinery of algebraic topology could be applied. D. Wagner suggested the following variation. For each $k = 0, 1, 2, \cdots$ let $W_k^n = W_k$ be the $\mathbb{Z}$-module

$$W_k = \left\{ f : \binom{[n]}{k} \to \mathbb{Z} \right\}$$

with coboundary maps

$$\begin{aligned} &\delta_k : W_k \to W_{k+1} \\ &\delta_k(f)(T) = \sum_{t \in T} (-1)^{j(t,T)} f(T \setminus t) \end{aligned}$$

where t is the $j(t,T)$th element of T. By a similar proof we again have $\operatorname{im} \delta_{k-1} \subseteq \ker \delta_k$.

These constructions generalize some notions used in the study of real equiangular lines. For instance, we have considered at various points the switching class of the graph Y^+ obtained from a graph Y by adding an isolated vertex. This generalizes to the operation

$$\begin{aligned} V_k^n &\to V_{k+1}^{n+1} \\ f &\mapsto \delta_k(f^+) \end{aligned}$$

The same operation can be defined for $W_k^n \to W_{k+1}^{n+1}$. We can also consider the unordered structures, *i.e.* the orbits of V_k^n under the symmetric group S_n. Let $\mathcal{V}_k^n = \mathcal{V}_k$ and $\mathcal{W}_k^n = \mathcal{W}_k$ denote the sets of *isomorphism classes* of elements of V_k^n and W_k^n respectively. There is no known "closed formula" for the number $g_{n,i} = g_{n,2,i}$ of nonisomorphic graphs on n vertices with i edges, however, it is known that these numbers are in a sense (described precisely in the following theorem) palindromic and unimodal. This generalizes to elements of $\mathcal{V}_k^n$:

Theorem 3.4.2. Let $g_{n,k,i}$ denote the number of elements $[f] \in \mathcal{V}_k^n$ such that $f(x) = 1$ for exactly i vertices $x \in [n]$. (Note that $g_{n,k,i}$ is well-defined.) Then for $n, k \in \mathbb{N}$, the sequence

$$\left(g_{n,k,i} \ : \ i = 0, \cdots, \tbinom{n}{k} \right) \tag{3.10}$$

is palindromic and unimodal, *i.e.* $g_{n,k,i} = g_{n,k,\binom{n}{k}-i}$ for each i and $g_{n,k,0} \leq \cdots \leq g_{n,k,\lceil \binom{n}{k}/2 \rceil}$.

The proof uses the representation theory of $\mathfrak{sl}(2, \mathbb{C})$, especially the so-called *Boolean representations.* We describe the finite dimensional irreducible representations of $\mathfrak{sl}(2, \mathbb{C})$ for reference, and state the theorems we need, but omit the proofs as they can be found in any textbook on the subject.

We begin by defining an $m+1$ dimensional irreducible representation of $\mathrm{SL}(2,\mathbb{C})$ for each integer $m \geq 0$. Let V_m be the complex vector space of complex homogeneous polynomials of degree m in 2 variables x_1 and x_2. V_m has basis $E_m = \{x_1^k x_2^{m-k} : k = 0, \cdots, m\}$. Let G be a 2×2 matrix Lie group. Consider the *left-translation representation*, given by

$$\begin{aligned} \Pi_m : G &\to \mathrm{GL}(V_m) \\ A(f(x)) &\mapsto f(A^{-1}x) \end{aligned}$$

It is straight forward to check that this defines a representation, and that the corresponding Lie algebra representation $\pi_m : \mathfrak{g} \to \mathfrak{gl}(V_m)$ is given by

$$\pi_m(X)(f(x)) = -\frac{\partial f}{\partial x_1}(X_{11}x_1 + X_{12}x_2) - \frac{\partial f}{\partial x_2}(X_{21}x_1 + X_{22}x_2) \tag{3.11}$$

Specializing to $G = \mathrm{SL}(2,\mathbb{C})$, a nice basis for $\mathfrak{sl}(2,\mathbb{C})$ is the following.

$$X = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \qquad Y = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} \qquad H = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

The commutation relations are

$$[H,X] = 2X \qquad [H,Y] = -2Y \qquad [X,Y] = H \tag{3.12}$$

π_m is determined by linearly extending its action on a basis. For instance, by (3.11),

$$\pi_m(X) = -x_2\frac{\partial}{\partial x_1} \qquad \pi_m(Y) = -x_1\frac{\partial}{\partial x_2} \qquad \pi_m(H) = -x_1\frac{\partial}{\partial x_1} + x_2\frac{\partial}{\partial x_2}$$

Now let us consider the action of these operators on the basis E_m for V_m. We have

$$\begin{aligned} \pi_m(X)(x_1^k x_2^{m-k}) &= -kx_1^{k-1}x_2^{m-k+1} \\ \pi_m(Y)(x_1^k x_2^{m-k}) &= -(m-k)x_1^{k+1}x_2^{m-k-1} \\ \pi_m(H)(x_1^k x_2^{m-k}) &= (m-2k)x_1^k x_2^{m-k} \end{aligned}$$

Note that for each $k = 0, \cdots, m$, the monomial $x_1^k x_2^{m-k}$ is an eigenvector of $\pi_m(H)$ with eigenvalue $m-2k$. These $m+1$ distinct eigenvalues and eigenvectors can be used to diagonalize $\pi_m(H)$. It is clear that the V_m are irreducible, because if $f(x) \in V_m$ is nonzero, and $x_1^k x_2^{m-k}$ is the monomial of $f(x)$ with nonzero coefficient of highest k, then $\pi_m(X)^k(f(x))$ is a multiple of x_2^m, and then for each j, the monomial $x_1^j x_2^{m-j}$ is a multiple of $\pi_m(Y)^j(x_2^m)$. Of course, the V_m are nonisomorphic representations since they have different dimensions. It turns out that these are the only finite dimensional irreducible representations of $\mathfrak{sl}(2,\mathbb{C})$.

Theorem 3.4.3. Let $\pi : \mathfrak{sl}(2,\mathbb{C}) \to \mathfrak{gl}(V)$ be a finite dimensional complex irreducible representation of $\mathfrak{sl}(2,\mathbb{C})$. Then π is isomorphic to one of the aforementioned π_m.

Theorem 3.4.4. Every finite dimensional representation of $\mathfrak{sl}(2,\mathbb{C})$ is completely reducible, *i.e.* isomorphic to a direct sum of irreducible representations.

Theorem 3.4.5. Let $\pi : \mathfrak{sl}(2,\mathbb{C}) \to \mathfrak{gl}(V)$ be a finite dimensional (complex) representation. Let $V = \bigoplus_{\lambda\in\mathbb{Z}} V_\lambda$ be the decomposition of V into the eigenspaces of $\pi(H)$, indexed by eigenvalue λ. (These are called *weight spaces.*) Then these eigenvalues are integers, and then the restricted operator $\pi(X) : V_\lambda \to V_{\lambda+2}$ is injective if $\lambda \leq -1$ and surjective if $\lambda \geq -1$, and for $\lambda \geq 0$, the restricted operator $\pi(X)^\lambda : V_{-\lambda} \to V_\lambda$ is an isomorphism. Similarly, the restricted operator $\pi(Y) : V_\lambda \to V_{\lambda-2}$ is injective if $\lambda \geq 1$ and surjective if $\lambda \leq 1$, and for $\lambda \geq 0$, the restricted operator $\pi(Y)^\lambda : V_\lambda \to V_{-\lambda}$ is an isomorphism.

Corollary 3.4.5.1. Let $V = \bigoplus_{\lambda\in\mathbb{Z}} V_\lambda$ be as above, and let $d_\lambda = \dim V_\lambda$. Then the sequences $(\cdots, d_{-4}, d_{-2}, d_0, d_2, d_4, \cdots)$ and $(\cdots, d_{-5}, d_{-3}, d_{-1} = d_1, d_3, d_5, \cdots)$ are palindromic and unimodal with maximums d_0 and $d_{-1} = d_1$, respectively. That is, $d_0 \geq d_2 \geq d_4 \geq \cdots$, $d_1 \geq d_3 \geq d_5 \geq \cdots$, and $d_\lambda = d_{-\lambda}$ for each λ.

Thus, to show a sequence is palindromic and unimodal, it suffices to show that it is the sequence of dimensions, counting by 2s, of weight spaces for some representation of $\mathfrak{sl}(2,\mathbb{C})$.

A useful tool for working with finite dimensional $\mathfrak{sl}(2,\mathbb{C})$ representations is the Poincaré polynomial. Like the character of a finite dimensional representation of a finite group, the Poincaré polynomial is a total invariant of the equivalence class of representations.

Definition 3.4.1. Let $V = \bigoplus_{\lambda\in\mathbb{Z}} V_\lambda$ be a finite dimensional representation of $\mathfrak{sl}(2,\mathbb{C})$, where V_λ is the λ-weight space. The *Poincare polynomial* for V is the following Laurent polynomial.

$$P(V;t) = \sum_{\lambda\in\mathbb{Z}} (\mathrm{dim} V_\lambda) t^\lambda$$

Thus, if V is the unique irreducible representation $V(m)$ of dimension $m+1$, then

$$P(V;t) = t^{-k} + t^{-k+2} + \cdots + t^{k-2} + t^{k} = \frac{t^{-k-1} - t^{k+1}}{t^{-1} - t} \tag{3.13}$$

It is straight forward to show that the Poincare polynomial has the following properties:

1. $P(V_1;t) = P(V_2;t)$ if and only if $V_1 \cong V_2$ (isomorphism of representations).
2. $P(V_1 \oplus V_2;t) = P(V_1;t) + P(V_2;t)$.
3. $P(V_1 \otimes V_2;t) = P(V_1;t)P(V_2;t)$.

By complete reducibility, the Poincare polynomial for any finite dimensional representation of $\mathfrak{sl}(2,\mathbb{C})$ is determined by (3.13) and property 2. Properties 2 and 3 allow us to decompose tensor product representations into direct sums.

We now introduce Boolean representations.

Definition 3.4.2. Let A be a finite set of size m and let $P(A) = \bigcup_k P_k(A)$ be its powerset, (where $P_k(A)$ is the set of k element subsets of A). The *Boolean representation* of $\mathfrak{sl}(2,\mathbb{C})$ for the set A is the pair (ρ, V) where $V = \mathbb{C}P(A)$ is the space of formal linear combinations of $P(A)$, and $\rho : \mathfrak{sl}(2,\mathbb{C}) \to V$ is defined as follows. For $S \in P(A)$, define

$$\rho(X)(S) = \sum_{a \in A \smallsetminus S} S \cup \{a\} \qquad\qquad \rho(Y)(S) = \sum_{a \in S} S \smallsetminus \{a\}$$

and $\rho(H) = \rho(X)\rho(Y) - \rho(Y)\rho(X)$. This is clearly an $\mathfrak{sl}(2,\mathbb{C})$-representation, and it can easily be shown that

$$\rho(H)(S) = (2k - m)S$$

so $P(A)$ is an eigenbasis. Thus if $\lambda = 2k - m$ and V_λ is the λ-weight space, then $\dim V_\lambda = \binom{m}{k}$. By the binomial theorem, the Poincare polynomial is

$$P(V;t) = \sum_{k=0}^{m} \binom{m}{k} t^{2k-m} = (t^{-1} + t)^m$$

so $V \cong V(1)^{\otimes m}$, where $V(1)$ is the unique 2 dimensional irreducible representation of $\mathfrak{sl}(2,\mathbb{C})$ (up to isomorphism), hence the name *Boolean representation.*

Theorem 3.4.6. Suppose a vector space V carries a representation π of a Lie algebra and $H \subseteq \mathrm{GL}(V)$ is a group of automorphisms of π. Then the H*-equivariant subspace*

$$V^H = \{v \in V \mid \phi(v) = v \text{ for all } \phi \in H\}$$

is a subrepresentation of V.

We are now ready to prove Theorem (3.4.2).

Proof of Theorem (3.4.2). Let A be set of all k-sets on $[n]$. Then A corresponds to the constant function $f \equiv 1$ in V_k^n. Consider the Boolean representation ρ of $\mathfrak{sl}(2, \mathbb{C})$ for the set A. Then S_n acts in the obvious way on $\mathbb{C}P(A)$, where $P(A)$ is the powerset of A and $\mathbb{C}P(A)$ is the complex vector space with basis $P(A)$. Provided that the elements of S_n commute with $\rho(X)$ and $\rho(Y)$, we obtain the S_n-equivariant subrepresentation $\mathbb{C}P(A)^{S_n}$. The $\rho(H)$-eigenspaces of $\mathbb{C}P(A)^{S_n}$ are spanned by uniform superpositions $\sum_{g \in [f]} g$ for each isomorphism class $[f] \in \mathcal{V}_k^n$. It follows that the sequence (3.10) is palindromic and unimodal. It is simple to show that the elements of S_n commute with $\rho(X)$ and $\rho(Y)$:

Let $\Sigma \in S_n$ and let $S \in P_i(A)$. To see that Σ commutes with $\rho(X)$, observe that

$$\begin{aligned}
(\Sigma \circ \rho(X))(S) &= \sum_{a \in A \smallsetminus S} \Sigma(S \cup \{a\}) \\
&= \sum_{a \in A \smallsetminus S} \Sigma(S) \cup \{\Sigma(a)\} \\
&= \sum_{a' \in A \smallsetminus \Sigma(S)} \Sigma(S) \cup \{a'\} \\
&= (\rho(X) \circ \Sigma)(S)
\end{aligned}$$

We omit the proof that Σ commutes with $\rho(Y)$ as it is very similar. ■

Chapter 4

Linear programming bounds

The linear programming bounds (LP bounds) for codes are a vast generalization of various bounds we have seen so far:

1. The relative bound: If $v_1, \cdots, v_n$ are equiangular unit vectors in K^d at angle α, then

   $$n \leq \frac{d(1-\alpha^2)}{1-d\alpha^2} \tag{4.1}$$

   The relative bound is a special case of the first Welch bound.

2. The first Welch bound: If $v_1, \cdots, v_n$ are unit vectors in K^d, then

   $$\max_{i \neq j} |\langle v_i, v_j \rangle|^2 \geq \frac{n-d}{d(n-1)}$$

   We proved the first Welch bound in section 2.4. As the name implies, there is an entire sequence of Welch bounds. These are also an instance of the LP bounds.

3. The general Welch bounds: If $v_1, \cdots, v_n$ are unit vectors in K^d, then

   $$\max_{i \neq j} |\langle v_i, v_j \rangle|^{2k} \geq \frac{1}{n-1} \left(\frac{n}{\binom{d+k-1}{k}} - 1 \right)$$

   for $k = 1, 2, 3, \cdots$.

4. The MUB bound: The maximum cardinality of a set of mutually unbiased bases for $\mathbb{C}^d$ is less than or equal to $d+1$.

5. The maximum cardinality of a set of unit vectors in $\mathbb{R}^d$ with mutual inner products $-1/t$ and 0 is less than or equal to $d+\frac{d}{t}$. (Since the cardinality of such a set is an integer, it must be less than or equal to $d+\lfloor\frac{d}{t}\rfloor$, and it this can be shown to be attainable for all $d \in \mathbb{N}$ and $t > 0$ using the Perron-Frobenius theorem.) This upper bound, without the floor signs, is an instance of the LP bounds. We used it when reviewing Lemmens and Seidel's method of pillars.

These are all generalized by the LP bounds. The absolute bounds for equiangular lines also have a more general form, which we describe in the next section. Our first step towards proving the LP bounds will be to define the kind of space in which the LP bounds hold, and along the way, we prove a generalized absolute bound. Godsil's definition of *polynomial space* amounts to a set X equipped with a separation function and an inner product on the space $\mathrm{Pol}(X)$ of polynomial functions on X satisfying four axioms. The LP bounds do not use Godsil's 2nd, 3rd, or 4th axioms, but we will need another axiom, essentially the *identity of indiscernibles* axiom of metric spaces. Godsil then defines Q-polynomial spaces as polynomial spaces satisfying two additional axioms. Q-polynomial spaces enlarge the class of Q-polynomial symmetric association schemes, and this is where they get their name [God93]. We prove the LP bounds for spaces with this Q-polynomial property, which we call *Q-polynomial code space*.

In [CS98], the LP bounds are derived for codes in compact 2-point homogeneous spaces, but in order to evaluate the bound, presumably, one has to derive some facts about the irreducible representations of the space's automorphism group. If such a space is Q-polynomial (which is true in all of our examples, and we expect might always be true), then the bound can be calculated by iteratively finding the zonal orthogonal polynomials. We prove the LP bounds in the more general setting of a Q-polynomial code space. These are defined combinatorially, without assuming the existence of an automorphism group, which turns out to be an externality.

Thus, our proof of the LP bounds is almost exactly the same as the one presented in [CS98], but we distill out the combinatorial parts. The rest, (the parts involving representation theory,) is left for sections 4.4 and 4.5, where we describe how compact 2-point homogeneous spaces come into the picture. Lastly we will do some explicit calculations involving the zonal orthogonal polynomials in section 4.6.

4.1 Polynomial functions and absolute bounds

We first state the generalized absolute bound, which is an upper bound on the size of a code of a given degree in a "space". There are few requirements on the space, so we state the theorem in generality, though our spaces of interest are all "Q-polynomial code spaces".

Let X be a set. A *separation function* on X is a function $\tau : X \times X \to \mathbb{R}$ satisfying $\tau(x, y) = \tau(y, x)$ for all $x, y \in X$, (symmetry). For each $x_0 \in X$, we define $\tau_{x_0} : X \to \mathbb{R}$ by $\tau_{x_0}(x) = \tau(x_0, x)$. Given $x_0 \in X$ and $f \in K[t]$ we denote $f_{x_0} = f \circ \tau_{x_0}$. A function on X of the form f_{x_0} is called a *zonal polynomial function* centered at x_0 on X (with respect to τ). We denote the real vector space of all such functions by $\mathrm{Zon}_{x_0}(X)$, and the space spanned by all zonal polynomial functions by $\mathrm{Zon}(X) = \sum_{x_0 \in X} \mathrm{Zon}_{x_0}(X)$. More generally, *polynomial functions* on X (with respect to τ) are elements of the degree-filtered algebra $\mathrm{Pol}(X)$ generated under pointwise multiplication by the zonal polynomial functions at varying points $x_0 \in X$. In other words,

$$\begin{aligned}
&\mathrm{Pol}(X, 0) = \{\text{constant functions } X \to \mathbb{R}\}\\
&\mathrm{Pol}(X, 1) = \{f_{x_0} \,|\, x_0 \in X\ ,\ f \in \mathbb{R}[t] \text{ linear or constant}\}\\
&\mathrm{Pol}(X, r + 1) = \mathrm{span}_{\mathbb{R}}\{fg \,|\, f \in \mathrm{Pol}(X, 1)\ ,\ g \in \mathrm{Pol}(X, r)\} \text{ for } r \geq 1\\
&\mathrm{Pol}(X) = \bigcup_{r=0}^{\infty} \mathrm{Pol}(X, r)
\end{aligned}$$

Polynomial functions in $\mathrm{Pol}(X, r) \setminus \mathrm{Pol}(X, r - 1)$ are said to have *degree* r. For a zonal polynomial function f_{x_0}, note that $\deg f_{x_0} \leq \deg f$. Depending on X and τ, it might be the case that not every polynomial function is a linear combination of zonal polynomial functions, *i.e.* $\mathrm{Zon}(X) \subsetneq \mathrm{Pol}(X)$.

Let X be a set equipped with a separation function τ. We call a finite subset $\mathcal{C}$ of X a *code*. The *degree* of $\mathcal{C}$ is the number $|\{\tau(x, y) \,|\, x, y \in \mathcal{C},\ x \neq y\}|$. $\mathcal{C}$ is an *orderable code* if there exists a total order $<$ on $\mathcal{C}$ such that for all $x, y \in \mathcal{C}$ with $y < x$, $\tau(x, x) \neq \tau(x, y)$. We call (X, τ) *discernible* if there exists $\alpha_0 \in \mathbb{R}$ such that for all $x, y \in X$, $\tau(x, y) = \alpha_0$ if and only if $x = y$. If (X, τ) is discernible, then every code $\mathcal{C} \subseteq X$ is orderable. (All of our spaces of interest are discernible.)

Theorem 4.1.1 (Absolute bound)**.** Let X be a set with separation function τ, and let $\mathcal{C}$

be an orderable code of degree at most r in X. Then

$$|\mathcal{C}| \leq \dim \mathrm{Pol}(X, r)$$

For the proof, see [God93]. Given how general the absolute bounds are, and thus how little information about the space X is taken into account, it should be no surprise that the absolute bounds are not always attained. For instance, if $\mathrm{Pol}(X, r)$ is infinite dimensional, then the absolute bound for codes of degree r is vacuous. Thus Godsil requires that $\mathrm{Pol}(X, 1)$ is finite dimensional, which implies that $\mathrm{Pol}(X, r)$ is finite dimensional for all r [God93]. Even for a well-chosen separation function, however, the absolute bound is rarely attained.

Let us introduce an example that we will return to throughout the chapter: the sphere. $X = S^{d-1}$ and $\tau(x, y) = x \cdot y$. Then $\mathrm{Pol}(X, 1) = (\mathbb{R}^d)^* \cong \mathbb{R}^d$, and similarly, $\mathrm{Pol}(X)$ can be identified with the coordinate ring

$$\mathbb{R}[X] = \mathbb{R}[x_1, \cdots, x_d]/(x_1^2 + \cdots + x_d^2)$$

Traditionally, one studies this example by first describing certain polynomial functions called the *spherical harmonics*. We will soon define the harmonic subspaces, which generalize this to other spaces X. The spherical harmonics can be calculated explicitly by solving Laplace's equation using the method of *separation of variables*, [AAR99, Gal]. In cartesian coordinates, the harmonics are multivariable polynomials, and in spherical coordinates, they are more complicated, and in a sense generalize the sine and cosine functions. Harish-Chandra generalized this idea further to certain symmetric spaces by introducing an algebra of invariant differential operators. We will not use any differential operators in our approach, but we note that in the case of the sphere, this method quickly leads one to the formula for $\dim \mathrm{Pol}(X, r)$. Then, using the absolute bound, one can easily see, for instance, that the maximum cardinality of a spherical code of degree 2 cannot exceed

$$\binom{d+3}{2}$$

This example was shown in [DGS77], where it also noted that the bound is occasionally saturated as follows: Consider a set of $\binom{d+3}{2}$ equiangular vectors in $\mathbb{R}^{d+1}$. (Four such configurations are known to exist.) Apply Seidel switches such that one vertex v is isolated, and then apply the affine shift to the entire configuration that sends $v \mapsto 0$. After renormalizing, the result is a degree 2 spherical code of cardinality $\binom{d+3}{2}$ in a d dimensional subspace.

Turning things around, this serves as a proof of the absolute bound for equiangular lines in $\mathbb{R}^d$. The absolute bound for real equiangular lines can also be found by considering the space $K\mathbb{P}^{d-1}$ with separation function $\tau(x, y) = |\langle x, y\rangle|^2$ and $K = \mathbb{R}$. Similarly, the absolute bound on complex equiangular lines is found for $K = \mathbb{C}$.

If $\mathrm{Zon}(X) = \mathrm{Pol}(X)$, we can also evaluate the absolute bound if we know the reproducing kernels of the harmonic subspaces. To define the harmonic subspaces, we must specify an inner product $\langle \cdot, \cdot \rangle$ on $\mathrm{Zon}(X)$. This technique for computing the absolute bound does not require one to solve any differential equations, but we need a method to calculate the reproducing kernels. In section 4.3 we introduce the zonal orthogonal polynomials, which by their very definition can be calculated (with some difficulty), and sometimes coincide with the reproducing kernels. We define the *harmonic subspaces* for $(X, \tau, \langle \cdot, \cdot \rangle)$ as follows:

$$\mathrm{Harm}(X, r) = \mathrm{Pol}(X, r) \cap \mathrm{Pol}(X, r-1)^{\perp} \cap \mathrm{Zon}(X)$$

If W is a finite dimensional real or complex inner product space of functions on a set X, then the *reproducing kernel* for W at a point $x \in X$ is the unique function $k_{(x)} \in W$ such that for all $f \in W$, $\langle f, k_{(x)} \rangle = f(x)$. Given any orthonormal basis $\{f_1, \cdots, f_d\}$ for W, the reproducing kernel satisfies $k_{(x)} = \overline{f_1(x)} f_1 + \cdots + \overline{f_d(x)} f_d$. However, note that in our application, *a priori* we do not know the dimension, let alone any basis for $\mathrm{Harm}(X, r)$.

In general, $k_{(x)}$ is not a zonal polynomial function centered at x, but the notation $k_{(x)}$ is chosen to suggest a similar behavior. For instance, suppose a group G acts on X, and thereby defines a representation on W by left translation ($L(g)(f)(x) = f(g^{-1}x)$). Assuming this representation is unitary, then for any $g \in G$ and $x \in X$, $gk_{(x)} = k_{(gx)}$, which looks just like the property that for any $g \in G$, $f \in \mathbb{R}[t]$, and $x \in X$, $gf_x = f_{gx}$.

Proposition 4.1.1. Let W be a finite dimensional real or complex inner product space of functions (including constant functions) on a set X, with inner product satisfying $\langle f, g \rangle = \langle 1, fg \rangle$ for all $f, g \in W$ and $\langle 1, 1 \rangle = 1$. Then $k_{(x)}(x) = d$ where $d = \dim W$.

Proof. Choose an orthonormal basis $f_1, \cdots, f_d$ for W. Then

$$k_{(x)}(x) = \sum_{i=1}^{d} \langle k_{(x)}, f_i \rangle f_i(x) = \sum_{i=1}^{d} |f_i(x)|^2$$

The inner product of the above constant function with the constant function 1 is

$$k_{(x)}(x) = \langle 1, k_{(x)}(x) \rangle = \sum_{i=1}^{d} \langle f_i, f_i \rangle = d$$

as claimed. ■

If $\mathrm{Pol}(X) = \mathrm{Zon}(X)$, and $\mathrm{Zon}(X)$ has an inner product as required above, then access to the reproducing kernels $k_{i,(x)}$ for the harmonic subspaces $\mathrm{Harm}(X, i)$ lets one compute

$$\dim \mathrm{Pol}(X, r) = \sum_{i=1}^{r} k_{i,(x)}(x)$$

for a point $x \in X$, and thereby lets one compute the absolute bound.

4.2 Code spaces

Our goal is to state and prove the LP bounds. We saw that it is convenient to place some restrictions on the set X in order to be able to evaluate the absolute bounds in practice. We will need these restrictions and more for the LP bounds, so let us define some terminology.

Definition 4.2.1. We define a *code space* to be an ordered quadruple $(X, \tau, \langle \cdot, \cdot \rangle, \alpha_0)$ where X is a set, τ is a separation function on X, $\langle \cdot, \cdot \rangle$ is an inner product on $\mathrm{Zon}(X)$, $\alpha_0 \in \mathbb{R}$, and the following axioms hold.

1. $\tau(x, y) = \alpha_0$ if and only if $x = y$ for all $x, y \in X$ (discernibility),
2. $\tau(x, y) = \tau(y, x)$ for all $x, y \in X$ (symmetry)

We usually abbreviate $(X, \tau, \langle \cdot, \cdot \rangle, \alpha_0)$ as X with τ, $\langle \cdot, \cdot \rangle$, and α_0 being implicit. (This sort of remark will be henceforth omitted from definitions.)

Our definition is a relaxation of Godsil's definition of *polynomial space*, though Godsil does not include axiom 1 [God93]. We acknowledge that the pursuit of generality is not always a good use of one's time, nor is it necessarily beneficial unless one has an example that *requires* the generality. Our reason for relaxing Godsil's axioms is that Godsil proves different theorems than the ones we prove here, three of his four axioms are simply never used in the proof of the LP bound.

A code space is very similar to a *semimetric space*, which is a pair (X, d) satisfying discernibility with $d(x, x) = 0$ and symmetry. Given a semimetric space (X, d), let $\phi : \operatorname{im} d \to \mathbb{R}$ be an injective function, and give X the separation function $\tau = \phi \circ d$. Suppose there exists a measure μ on X allowing for the integration of polynomial functions (with respect to τ), and $\mu(X) = 1$. Then we can define an inner product on $\mathrm{Zon}(X)$ by

$$\langle f_x, g_y \rangle = \int_X f_x(z) g_y(z) d\mu(z)$$

Then $(X, \tau, \langle \cdot, \cdot \rangle, \phi(0))$ is a code space. Many interesting code spaces arise in this way. For instance, if X is finite, then μ can be taken to be the uniform measure.

We will consider three infinite spaces that arise in this way. The first is the sphere $X = S^{d-1}$ with separation function $\tau(x, y) = x \cdot y$ spherical measure μ. The other two are the real and complex projective spaces. Recall that the distance between two points $[u], [v] \in K\mathbb{P}^{d-1}$ is $\sqrt{2 - 2|\langle u, v\rangle|^2}$, which is uniquely determined by $|\langle u, v\rangle|^2$. Thus in these spaces we use $\tau([u], [v]) = |\langle u, v\rangle|^2$ as our separation function. We will describe a method to obtain measures on these spaces when we discuss two-point homogeneous spaces, and we will give an explicit formula in section 4.6. It will become clear in the examples section why the squared overlap is a more natural choice than the unsquared version. To avoid possible confusion caused by using the squared overlap, we use letters like α, α_0, A for angles and such when working with the space S^{d-1} and β, β_0, B with $K\mathbb{P}^{d-1}$.

Another infinite space that would be nice to consider is the Euclidean space $X = \mathbb{R}^d$. Unfortunately, we are at a loss to come up with a suitable inner product, and this seems to arise from the fact that $\mathbb{R}^d$ is not compact.

One may wonder why we don't let $\tau(x, y)$ in the aforementioned examples equal the geodesic between x and y. The reason is that the expressions for the zonal orthogonal polynomials will be cleaner, and this is the main reason for allowing flexibility in our definitions.

A code space X *embeds into* another code space Y if there exists an injection $f : X \to Y$

such that for all $x, y \in X$,

$$\tau_X(x, y) = \tau_Y(f(x), f(y))$$

It turns out that if a code space X embeds into a sphere, real projective space, or complex projective space, then $\mathrm{Pol}(X) = \mathrm{Zon}(X)$. As noted in the previous section, this fact can be used to justify an evaluation of the absolute bounds. We omit the proof of this fact. Godsil proves it for spherical embeddings [God93]. We note this only because most of our spaces of interest have this property.

Let us take note of three commonly encountered finite code spaces.

1. The Hamming scheme $H(d, q)$, $(d, q \in \mathbb{N})$. The underlying space consists of the strings of length n over an alphabet Σ of size q. Without loss of generality, let $\Sigma = \{0, \cdots, q-1\}$. (We will use arithmetical properties of Σ below for the sole purpose of showing that the Hamming scheme embeds into the sphere.) The *Hamming distance* $h(x, y)$ between two strings x, y is the number of letters in x and y that differ. When $h(x, y) = 1$, the strings x and y have an edge in the *Hamming graph*, and the distance between x and y in the Hamming graph is $h(x, y)$. This graph is distance regular, so $H(d, q)$ is an association scheme where x and y are ith associates precisely when $h(x, y) = i$. Let the separation function be $\tau(x, y) = d - h(x, y)$. The Hamming scheme embeds into the sphere. For $q = 2$, we can send each string $x = x_1 \cdots x_d$ to $((-1)^{x_1}, \cdots, (-1)^{x_d}) \in S^{d-1}$. For $q \geq 2$, send each string $x = x_1 \cdots x_d$ to

   $$\frac{1}{\sqrt{qd}} \left((\omega^{(0)x_1}, \cdots, \omega^{(q-1)x_1}), \cdots, (\omega^{(0)x_d}, \cdots, \omega^{(q-1)x_d}) \right) \in S^{2qd-1}$$

   where $\omega^{x_i t}$ stands for $(\cos \frac{2\pi i x_i t}{q}, \sin \frac{2\pi i x_i t}{q})$.

2. The Johnson scheme $J(v, k)$, $v \in \mathbb{N}$, $k \in [v]$. The underlying set consists of the set of k-subsets of the set $[v]$ (or binary strings of length v with exactly k 1s). The *Johnson distance* $j(x, y)$ between two k-subsets x, y is $k - |x \cap y|$, (the same as the Hamming distance, if we think of the elements as strings). When $j(x, y) = 1$, the sets x and y have an edge in the *Johnson graph*, and the distance between x and y in the Johnson graph is $j(x, y)$. This graph is again distance regular, so $J(v, k)$ is an association scheme where x and y are ith associates precisely when $j(x, y) = i$. We use separation function $\tau(x, y) = v - j(x, y)$. The Johnson scheme also embeds

into the sphere. If v_x is the indicator vector for the k-set x, then we send x to $\frac{1}{\sqrt{k}}v_x \in S^{v-1}$.

3. Symmetric association schemes. This includes all distance regular graphs. We restrict our attention to Q-polynomial (symmetric) association schemes, which are Q-polynomial code spaces. We make this restriction because we can only prove the LP bounds in a Q-polynomial code space. The class of Q-polynomial association schemes includes many distance regular graphs, such as the Hamming graphs, the Johnson graphs, and perhaps all distance transitive graphs (see remarks at the end of Section 4.5). As in the examples above, we use separation function $\tau(x,y) = n - d(x,y)$, where $n = |V(G)|$ and $i = d(x,y)$ means x and y are ith associates.

Definition 4.2.2. A *code* in a code space X is a finite subset $\mathcal{C}$ of X. The elements of a code are called *codewords*. Let $A \subseteq \operatorname{im}\tau \smallsetminus \{\alpha_0\}$. An *$A$-code* in X is a code $\mathcal{C}$ in X such that for all distinct $x, y \in \mathcal{C}$, $\tau(x,y) \in A$. We denote the maximum cardinality of an A-code in X by $N(X,A)$ and we denote $N(X) = \sup\{N(X,\alpha) : \alpha \in \mathbb{R} \smallsetminus \{\alpha_0\}\}$.

We have used this notation previously when discussing equiangular lines.

While a code is required to be finite, the set A is not. This flexibility is added to the definition because we are generally interested in finding or bounding $N(X,A)$ for a given set A. This is what the LP bounds are designed to do.

A $\{\pm\alpha\}$-code in S^{d-1} is a set of equiangular unit vectors in $\mathbb{R}^d$. A $\{\beta\}$-code in $K\mathbb{P}^{d-1}$ is a set of equiangular lines in K^d. Similarly, a B-code in $\mathbb{RP}^{d-1}$ corresponds to an A-code in S^{d-1} up to Seidel switching, where $A = \{\pm\sqrt{\beta} \mid \beta \in B\}$. (Note that $-1 \notin A$). An $\{r,R\}$-code (or an $[r,R]$-code) in $\mathbb{R}^d$ is a packing of balls of radius (at least) r into a ball of radius $R + 2r$. As mentioned before, we cannot apply the LP bounds here for lack of a separation function, but the number $N(\mathbb{R}^d, [r,R])$ could be bounded by packing tiny balls into a small ball in a sphere or projective space of real dimension d.

From this, one can bound the sphere packing density in $\mathbb{R}^d$:

$$\lim_{R/r\to\infty} \frac{N(\mathbb{R}^d, [r,R])}{V_d(R)}$$

where $V_d(R)$ is the volume of the ball enclosed by the $d-1$ sphere of radius R.

Codes in the Hamming scheme or Johnson scheme are used for error correction, where they represent messages somehow encoded by adding redundancy that makes them more resilient against noise. When a codeword is sent across a noisy channel, it accrues errors in the individual letters, but if there aren't too many errors, it can be recovered by decoding a nearest codeword (or *the* nearest codeword in some cases).

When q is a prime power, it is convenient to regard the underlying alphabet Σ of $H(d,q)$ as a finite field $\mathbb{F}_p^r$ of order q. Here, p can be any prime power such that $p^r = q$. For such a choice of p, a code $\mathcal{C}$ in $H(d,q)$ is called *linear* if it is a subspace of the $\mathbb{F}_p$-vector space $\mathbb{F}_q^n$. For a fixed codeword $x \in \mathcal{C}$, $\mathcal{C}$ is a coset $\mathcal{C}' + x$ of a linear code $\mathcal{C}'$, and the bijection $\mathcal{C} \to \mathcal{C}'$ given by $y \mapsto y - x$ preserves the separation function. Thus we may as well only consider linear codes. Every linear code has a *dual code* $\mathcal{C}^\perp$, which is just the orthogonal subspace under the dot product. It is not generally the case (contrasting with *inner product spaces* over $\mathbb{R}$ or $\mathbb{C}$) that $\mathcal{C} \cup \mathcal{C}^\perp$ is the entire space. Indeed, there even exist *self-dual codes* (codes $\mathcal{C}$ that are proper subsets of $H(d,q)$ for which $\mathcal{C} = \mathcal{C}^\perp$), such as the extended Golay code. Nevertheless, it still holds that $\dim \mathcal{C} + \dim \mathcal{C}^\perp = rn$. The minimum Hamming distance of a linear code is related to the *strength* of its dual. We will not discuss strength and t-designs in this thesis so we leave it at that.

4.3 Linear programming bounds in a Q-polynomial code space

In this section we derive an upper bound on $N(X, A)$, which we call the *linear programming bound* or *LP bound*. A related LP bound exists, which gives a lower bound on the size of a code of a given strength. We reserve the term "LP bound" (without modifier) to refer to the upper bound to be derived for $N(X, A)$. Still, this lower LP bound is related to the one we derive. For example, in the case of a linear Hamming code, the lower LP bound can be applied to the dual code, which gives a bound on the original code in terms of its largest minimum Hamming distance. See [God93] for details.

Definition 4.3.1. A code space X is *regular* if for each $i, j \in \mathbb{N}$, the *moments* $\langle \tau_x^i, \tau_x^j \rangle$ depend only on $i + j$ (and not on x). X is *distance regular* if it is regular, and the numbers $\langle \tau_x^i, \tau_y^j \rangle$ depend only on i, j, and $\tau(x, y)$, (but not explicitly on x or y).

Godsil calls such spaces *1-homogeneous* and *2-homogeneous*, respectively. We have opted for a different terminology because we will later discuss homogeneous and 2-point homogeneous spaces, which involve a group action. In our usage, regular graphs are regular code spaces, and distance regular graphs are distance regular code spaces.

Suppose a group G acts transitively on a code space X. Then it is clear that X is regular. If the left translation representation of G on $\mathrm{Zon}(X)$ is unitary and G acts distance transitively on X (*i.e.* for all $x_1, y_1, x_2, y_2 \in X$, there exists $g \in G$ such that $x_2 = gx_1$ and $y_2 = gy_1$ if and hence only if $\tau(x_1, y_1) = \tau(x_2, y_2)$,) then X is distance regular. The spaces $H(d, q), J(v, k), S^{d-1}, \mathbb{RP}^{d-1}$, and $\mathbb{CP}^{d-1}$ each have a distance transitive automorphism group, and this is also true of any distance transitive graph.

Note, however, that regularity and distance regularity are defined purely combinatorially, making no reference to any group action. Just as distance regular graphs were conceived as a combinatorial generalization of distance transitive graphs, these regularity properties generalize transitivity and distance transitivity of a group action. For the LP bound, we will need a property even stronger than distance regularity, (which not all distance regular graphs satisfy,) called the *Q-polynomial* property or *addition rule* [God93].

Let $x_0 \in X$ and consider the linear map

$$\begin{aligned} \mathbb{R}[t] &\to \mathrm{Zon}_{x_0}(X) \\ f &\mapsto f_{x_0} \end{aligned}$$

This lets us define a (possibly degenerate) inner product on $\mathbb{R}[t]$ by

$$(f, g)_{x_0} = \langle f_{x_0}, g_{x_0} \rangle$$

In a regular code space X, $(\cdot, \cdot)_{x_0}$ does not depend on the base point x_0, so we denote it by $(\cdot, \cdot)$. While it is possible to define the zonal orthogonal polynomials locally, we only consider the regular case because the LP bounds require globally defined zonal orthogonal polynomials.

The *zonal orthogonal polynomials* for X are the elements of $\mathbb{R}[t]$ defined recursively as follows: $z_0 = 1$, and having defined $z_0, \cdots, z_{r-1} \in \mathbb{R}[t]$, we let z_r be the unique polynomial of degree r that is orthogonal to $z_0, \cdots, z_{r-1}$ and satisfies

$$(z_r, z_r) = z_r(\alpha_0) \neq 0 \tag{4.2}$$

if such a polynomial exists, and terminate otherwise. These polynomials are well-defined. First, if $\{z_0, \cdots, z_{r-1}\} \in \mathbb{R}[t]$ is an independent set of polynomials of degrees $< r$, then $\{z_0, \cdots, z_r\}$ is independent for any polynomial z_r of degree r, so in particular, there must be such a z_r that is orthogonal to $z_0, \cdots, z_{r-1}$. Any other polynomial of degree r is then spanned by $\{z_0, \cdots, z_r\}$, so if $(z_r, z_r) \neq 0$ then z_r is determined up to nonzero real scalar. Then, if $(z_r, z_r) = cz_r(\alpha_0)$ for some nonzero $c \in \mathbb{R}$, then the unique nonzero $k \in \mathbb{R}$ such that $z_r' = kz_r$ satisfies (4.2) is $k = \frac{1}{c}$, so replace z_r with $\frac{1}{c}z_r$. Note that by discernibility and regularity, z_r does not depend on x_0.

Definition 4.3.2. We call a regular code space a *Q-polynomial code space* if it satisfies a property Godsil calls the *addition rule*: for all $x, y \in X$ and $r, s \in \mathbb{N}$,

$$\langle z_{r,x}, z_{s,z} \rangle = \begin{cases} z_{r,x}(y) = z_{r,y}(x) & \text{if } r = s \\ 0 & \text{if } r \neq s \end{cases}$$

The LP bounds hold in all Q-polynomial code spaces. All of our spaces of interest are Q-polynomial, though this must be verified in each case. It follows from our definition that z_{r,x_0} has degree r and moreover is an element of $\mathrm{Harm}(X, r)$. The following proposition gives two other characterizations of a Q-polynomial code space.

Proposition 4.3.1. Let X be a regular code space. Suppose that for all distinct $s_1, s_2 \in \mathbb{N}$, $\langle z_{x,s_1}, z_{y,s_2} \rangle = 0$. Then for each $r \in \mathbb{N}$, the following are equivalent:

1. The reproducing kernel for $\mathrm{Harm}(X, r)$ at x_0 is zonal at x_0 for one and hence all $x_0 \in X$.
2. z_{r,x_0} is the reproducing kernel for $\mathrm{Harm}(X, r)$ at x_0 for one and hence all $x_0 \in X$.
3. $\langle z_{r,x}, z_{r,y} \rangle = z_{y,r}(x) = z_{x,r}(y)$ for all $x, y \in X$. (X is Q-polynomial.)

In order to prove that 3 implies 1 or 3 implies 2 in Proposition (4.3.1), we need the following fact. The remainder of the proof is easy.

Proposition 4.3.2. Every Q-polynomial code space is distance regular.

Proof. Since the zonal orthogonal polynomials $z_0, \cdots, z_i$ are linearly independent of degrees $0, \cdots, i$, respectively, we may write

$$t^i = c_{i,0}z_0 + \cdots + c_{i,i}z_i$$

Solving for the coefficients $c_{i,k}$ we find

$$t^i = \frac{(t^i, z_0)}{z_0(\alpha_0)} z_0 + \cdots + \frac{(t^i, z_i)}{z_i(\alpha_0)} z_i$$

And so

$$\tau_x^i = \frac{(t^i, z_0)}{z_0(\alpha_0)} z_{x,0} + \cdots + \frac{(t^i, z_i)}{z_i(\alpha_0)} z_{x,i}$$

Now, by the addition rule we get

$$\langle \tau_x^i, \tau_y^j \rangle = \sum_{k=0}^{\min\{i,j\}} \frac{(t^i, z_k)(t^j, z_k)}{z_k(\alpha_0)} z_k(\tau(x,y))$$

which depends on only i, j, and $\tau(x, y)$.

■

For convenience, we also introduce the following notation. If X is a Q-polynomial code space, let $J_r : X \times X \to K$ be the function

$$J_r(x,y) := \langle z_{x,r}, z_{y,r} \rangle \overset{\star}{=} z_{x,r}(y) = z_r(\tau(x,y))$$

The $\star$ over the middle equality indicates that the addition rule is used. The restriction $J_r|_{\mathcal{C}\times\mathcal{C}}$ of J_r to a code $\mathcal{C}$ in X can be viewed as a positive semidefinite $\mathcal{C} \times \mathcal{C}$ matrix (since it is a Gram matrix). In some applications, J_r is called the *feature map*.

Definition 4.3.3. The *distance distribution* $\{p_\alpha \,|\, \alpha \in A\}$ of an A-code $\mathcal{C}$ in a Q-polynomial code space X consists of the data

$$p_\alpha = \frac{1}{|\mathcal{C}|} \Big| \{(x,y) \in \mathcal{C} \times \mathcal{C} \,|\, \tau(x,y) = \alpha\} \Big|$$

Denote $\overline{A} = A \cup \{\alpha_0\}$. Then

$$\begin{aligned} p_{\alpha_0} &= 1 \\ p_\alpha &\geq 0 \text{ for all } \alpha \in A \\ \sum_{\alpha \in \overline{A}} p_\alpha &= |\mathcal{C}| \end{aligned}$$

The *transformed distribution* of $\{p_\alpha \,|\, \alpha \in A\}$ is the set $\{q_r \,|\, r \in \mathbb{N}\}$ where

$$q_r = \frac{1}{|\mathcal{C}|} \sum_{\alpha \in \overline{A}} p_\alpha z_r(\alpha) \overset{\star}{=} \frac{1}{|\mathcal{C}|^2} \sum_{x,y \in \mathcal{C}} J_r(x, y) \geq 0$$

The inequality above holds because J_r is positive semidefinite. The $\star$ indicates that we have just made (crucial) use of the addition rule. Consider now a finite subset $\{\alpha_1, \cdots, \alpha_N\}$ of A. It is clear that an upper bound for $N(X, A) - 1$ is the optimum value of the following *primal LP*. (The -1 is because $\alpha_0 \notin \{\alpha_1, \cdots, \alpha_N\}$.)

$$\begin{array}{lll} \text{Maximize} & \displaystyle\sum_{i=1}^{N} p_{\alpha_i} & \\ \text{Subject to} & p_{\alpha_i} \geq 0 & \text{for } i = 1, \cdots, N \\ & \displaystyle\sum_{i=1}^{N} p_{\alpha_i} z_r(\alpha_i) \geq -1 & \text{for } r = 0, 1, 2, \cdots \end{array}$$

This LP only takes into account the distance distribution of the code and no other combinatorial or geometric information, so we should not expect the optimum value to be necessarily attainable by a code. Perhaps surprisingly then, in certain Q-polynomial code spaces, codes attaining this optimum value do exist. They are called *tight codes*. The *dual LP* is more useful in practice. For real variables $f_1, \cdots, f_N$, the dual LP reads:

$$\begin{array}{lll} \text{Minimize} & \displaystyle\sum_{r=1}^{N} f_r & \\ \text{Subject to} & f_r \geq 0 & \text{for } r = 1, \cdots, N \\ & \displaystyle\sum_{r=1}^{N} f_r z_r(\alpha) \leq -1 & \text{for all } \alpha \in A \end{array}$$

The Weak Duality Theorem can be molded into the following statement.

Theorem 4.3.1 (LP bound)**.** Let $\mathcal{C}$ be an A-code in a Q-polynomial code space X, $f(t) \in \mathbb{R}[t]$ a real polynomial of degree N, and $f_0, \cdots, f_N \in \mathbb{R}$ its zonal coefficients, that is,

$$f(t) = \sum_{r=0}^{N} f_r z_r(t)$$

Assume $f_0 > 0$ and $f_r \geq 0$ for all $r = 1, \cdots, N$. Further assume that $f(\alpha) \leq 0$ for all $\alpha \in A$. Then

$$|\mathcal{C}| \leq \frac{f(\alpha_0)}{f_0}$$

We will apply these bounds in our spaces of interest in Section 4.6. It may be convenient to compute the zonal coefficients as follows using the orthogonality of the z_r.

$$f_r = \frac{(f, z_r)}{z_r(\alpha_0)}$$

All of these must be computed to verify that $f(x)$ satisfies the hypotheses of the theorem, but to compute the bound one only needs to compute $f_0 = (f, 1)$.

4.4 An algebraic Peter-Weyl Theorem

In this section we introduce 2-point homogeneous spaces. By the end of the next section we will derive a criterion for showing that such a space is a Q-polynomial. It might be the case that this criterion is always satisfied, and that its proof involves the Laplacian (which is why I was unable to prove it now). The main result of this section is that when this criterion is satisfied (and in a slightly more general case), the zonal orthogonal polynomials are equal to (sums of) some functions we introduce here using representation theory. It is only in light of this observation that one is justified in calculating the LP bounds as described in section 4.3. Our main motivation for studying LP bounds is to better understand spherical and projective codes, and indeed, S^{d-1}, $\mathbb{RP}^{d-1}$, and $\mathbb{CP}^{d-1}$ are all 2-point homogeneous spaces. Two other examples we have considered, the Hamming graph and the Johnson graph, are distance transitive graphs, and all distance transitive graphs will be covered by this case.

General Q-polynomial association schemes or Q-polynomial distance regular graphs are not accounted for in the theory presented here.

Definition 4.4.1. A *topological group* is a group equipped with a topology in which multiplication and inversion are continuous. Suppose a compact group G acts transitively on a code space X. Then we call X a *homogeneous space*. If G acts distance transitively, we call X a *2-point homogeneous space*.

When G acts transitively on X, the Orbit-Stabilizer Theorem furnishes a canonical identification $X = G/H$ where H is the stabilizer of x_0. Then X is compact under the quotient topology, and this topology should always equal the topology of interest on X.

There are varying conventions about the definition of 2-point homogeneous space. It is also commonly assumed that X is a Riemannian manifold or that G is a Lie group. We will not make these assumptions, but incidentally, all of our spaces of interest are Riemannian manifolds whose automorphism groups are Lie groups.

It has been proven that if X is an infinite compact connected 2-point homogeneous Riemannian manifold, then X must be one of the spaces listed in the following table [CS98]. The last three columns are included here for reference and will be explained in section 4.6. Caution: erroneous values for a and b have propagated through some of the literature. The values here are taken from [TZ12] and partially verified by me in the first three cases only.

X	G	H	a	b	$z_r(t)$
S^{d-1}	$\mathrm{SO}(d)$	$\mathrm{SO}(d-1)$	$(d-3)/2$	$(d-3)/2$	$P_r^{(a,b)}(t)$
$\mathbb{RP}^{d-1}$	$\mathrm{SO}(d)$	$\mathrm{O}(d-1)$	$(d-3)/2$	$-1/2$	$P_r^{(a,b)}(2t-1)$
$\mathbb{CP}^{d-1}$	$\mathrm{SU}(d)$	$\mathrm{U}(d-1)$	$d-2$	0	$P_r^{(a,b)}(2t-1)$
$\mathbb{HP}^{d-1}$	$\mathrm{Sp}(d)$	$\mathrm{Sp}(d-1) \times \mathrm{Sp}(1)$	$2d-3$	1	$P_r^{(a,b)}(2t-1)$
$\mathbb{OP}^2$	F_4	$\mathrm{Spin}(9)$	7	3	$P_r^{(a,b)}(2t-1)$

The *automorphism group* of a code space X is the group of all bijections $X \to X$ preserving τ. For instance, the automorphism group of S^{d-1} is $\mathrm{O}(d)$. However, it is sufficient to use any distance transitive subgroup, so we will imprecisely refer to each of the preferred groups G listed in the table as the *automorphism group* of X.

In our spaces of interest (the first three), it is easy to show that G acts distance transitively on X and that $\mathrm{Stab}(x_0) \cong H$. Slight variations are available for the automorphism groups, but we use the ones mentioned here. We will not discuss the quaternionic projective spaces or the octonionic plane, though they are often included in treatments of the subject. We leave the work of computing the zonal orthogonal polynomials for these spaces to Section 4.6. As mentioned in Section 4.2, it would be interesting to consider codes in Euclidean space $\mathbb{R}^d$. The automorphism group can be taken to be the special Euclidean group $\mathrm{SE}(d) = \mathrm{T}(d) \rtimes \mathrm{SO}(d)$. Again, the orbit stabilizer theorem gives the identification $\mathbb{R}^d = \mathrm{SE}(d)/\mathrm{SO}(d)$. The quotient topology on $\mathbb{R}^d$ inherited from $\mathrm{SE}(d)$ equals the topology induced by the Euclidean metric on $\mathbb{R}^d$, so since $\mathbb{R}^d$ is not compact, neither is $\mathrm{SE}(d)$. Thus, Euclidean space does not fit into our framework.

If G is a locally compact group, then there exists a unique measure μ on G that is invariant under left resp. right translation and satisfies $\mu(G) = 1$, called the *left* resp. *right Haar measure*. When G is compact, left and right Haar measure are equal. Compact groups generalize finite groups, which are compact groups under the discrete topology. Thus, for finite groups, the integrals shown below can be replaced with summations normalized by the order of the group.

The Haar measure allows us to define the set $L^2(G)$ of square integrable complex-valued functions (modulo almost-everywhere zero functions) on a locally compact group G. The reason for modding out the almost-everywhere zero functions so that the following inner product is non-degenerate, and so $L^2(G)$ becomes a Hilbert space.

$$\langle f, g\rangle = \int_G f(x)\overline{g(x)}d\mu(x) \tag{4.3}$$

In the finite case, $L^2(G)$ is the same as $\mathbb{C}G$, the space of $\mathbb{C}$-formal sums over G. There is much to be said here but it is not especially relevant, so let us only give a brief summary: $L^2(G)$ is an algebra under convolution, and this leads to a very beautiful and general theory of the Fourier transform defined over compact groups. For the most part, this is tangential to our study of the LP bounds, and the algebra structure of $L^2(G)$ can be ignored. $\mathbb{C}G$ is called the *group algebra*, so that is what we shall call it, even though we usually won't need to think about its algebra structure. That being said, in the next section we will derive a small fragment of Fourier theory in order to point out that we can, without loss of generality, assume that (G, H) is a Gelfand pair.

Since $L^2(G)$ generalizes the group algebra from the finite case, one might expect it to carry the various kinds of regular representations, and indeed it does. We will only ever consider left, complex representations in this thesis. There are two such regular representations: the left translation and right translation representations. They are defined by

$$L(g)(f(x)) = f(g^{-1}x)$$
$$R(g)(f(x)) = f(xg)$$

By associativity of the group law, these actions commute, so we obtain a (more general) *two-sided (left) regular representation* S of $G \times G$ given by

$$S(g_1, g_2)(f(x)) = f(g_1^{-1}xg_2)$$

All of these representations are unitary with respect to the inner product (4.3).

The left (or right) translation representation of a finite group is isomorphic to the direct sum over a complete set of non-isomorphic irreducible representations, each with multiplicity equal to their degree, and a similar (more general) isomorphism holds for the two-sided regular representation, (Maschke's Theorem). This theorem generalizes to compact groups. The Peter-Weyl Theorem asserts that this (algebraic) truth generalizes to all compact groups, and further addresses some analytic nuances that arise in the infinite case. For instance, what should play the role of the group algebra? At one extreme, we could consider the space of all functions on G, and at another, we could consider the space of functions on G with finite support, or perhaps, if G has some nice algebraic properties, polynomial functions on G. The space $L^2(G)$ lies somewhere in the middle and appears in the usual statement of the Peter-Weyl Theorem. For our purposes, a more manageable space $C_{\text{alg}}(G)$ will do the job. We will not need the analytic implications of the Peter-Weyl theorem, which are harder to prove anyway, and this is our reason for preferring $C_{\text{alg}}(G)$.

Recall that if $\rho : G \to GL(V)$ is a finite dimensional representation, then the *dual* of ρ is the representation

$$\begin{aligned} \rho^t : G &\to GL(V^*) \\ g &\mapsto \rho(g^{-1})^t \end{aligned}$$

where V^* is the dual space of V (the vector space of linear functionals $V \to \mathbb{C}$), and if

$T \in \mathrm{Hom}(V, W)$, then its *transpose* is the map

$$\begin{aligned} T^t : W^* &\to V^* \\ \phi &\mapsto \phi \circ T \end{aligned}$$

If V is a finite dimensional complex inner product space, then every $\phi \in V^*$ has the form u^* for some $u \in V$, given by $u^*(v) = \langle v, u \rangle$.

We need to introduce some notations and state a few facts before stating the weakened version of the Peter-Weyl theorem, which we call the *Algebraic Peter-Weyl theorem*. First, if G is a compact group, then we let $\widehat{G}$ denote any choice of a complete set of pairwise non-isomorphic finite dimensional irreducible unitary representations of G. It does no harm to remove the requirement of finite dimensionality in our definition of $\widehat{G}$, due to the following theorem.

Theorem 4.4.1. Every irreducible representation of a compact group is finite dimensional.

It also turns out that every finite dimensional representation of a locally compact group is unitarizable. This fact, unlike Theorem (4.4.1), is easy to prove (believing that left Haar measure exists). It has essentially the same proof as in the finite case, using the quintessential "averaging trick" of representation theory. Thus, the representations in $\widehat{G}$ would automatically all be finite dimensional.

On the other hand, Theorem (4.4.1) turns out to be quite nontrivial, and luckily, we don't really need it. It is perfectly fine for our purpose to simply stipulate that $\widehat{G}$ only contain finite dimensional representations. Sometimes Theorem (4.4.1) is included in the statement of the Peter-Weyl theorem.

Next, for each $\rho \in \widehat{G}$ and $v \otimes u^* \in V_\rho \otimes V_\rho^*$, define

$$f_{\rho; v \otimes u^*}(x) = u^*(\rho(x)^{-1} v) = \langle v, \rho(x) u \rangle$$

Then $f_{\rho; v \otimes \phi}$ is a well-defined square integrable function on G. Functions of this form are called *matrix elements* of G and embed into a subspace (a subalgebra, even) $C_{\text{alg}}(G)$ of $L^2(G)$, which is an invariant subspace for each regular representation L, R, and S.

Finally, note that if V is a finite dimensional complex vector space, then we have the following canonical vector space isomorphism, which we view as an identification.

$$\begin{aligned} V \otimes V^* &\to \mathrm{End}(V) \\ v \otimes \phi &\mapsto (w \mapsto \phi(w)v) \end{aligned}$$

If $\rho \in \widehat{G}$, then $V_\rho \otimes V_\rho^*$ carries the representation $\rho \otimes \rho^t$ of $G \times G$. *I.e.* the left factor of the direct product acts on V_ρ by ρ and the right factor acts on $V_\rho^* = V_{\rho^t}$ by ρ^t. (This is sometimes denoted $\boxtimes$ instead of $\otimes$ to distinguish it from the kind of tensor product of representations we encountered in section 3.4. There will be no confusion here so we will stick to the notation $\otimes$.)

Now we can state and prove what we call the *algebraic version* of the Peter-Weyl Theorem. Its proof is relatively simple, unlike that of the full Peter-Weyl theorem.

Theorem 4.4.2 (Algebraic Peter-Weyl Theorem)**.** Let G be a compact group. Then the linear map

$$\begin{aligned} T : \bigoplus_{\rho \in \widehat{G}} V_\rho \otimes V_\rho^* &\to C_{\mathrm{alg}}(G) \\ v \otimes \phi &\mapsto \sqrt{d_\rho} f_{\rho; v \otimes \phi}(x) \end{aligned}$$

is an isometric isomorphism of representations of $G \times G$.

(On the left-hand-side, $G \times G$ acts by $\bigoplus_{\rho \in \widehat{G}} \rho \otimes \rho^t$ and on the right-hand-side, $G \times G$ acts by the two-sided regular representation S.)

It follows that T is an isomorphism of representations of $G \times 1$ and $1 \times G$ when considering the restricted representations to these subgroups.

Proof. First note that T is a surjective linear map by construction. It suffices to prove that T intertwines the $G \times G$ action, is injective, and is an isometry.

We first show that T intertwines the $G \times G$ action. Let $\rho \in \widehat{G}$. Since V_ρ is finite dimensional, every functional $\phi \in V_\rho^*$ has the form $\phi(v) = \langle v, u \rangle$ for some $u \in V_\rho$, that is, $\phi = u^*$. Now let $v \otimes u^* \in V_\rho \otimes V_\rho^*$ and $(g_1, g_2) \in G \times G$. Then

$$\begin{aligned}
(S(g_1, g_2) \circ T)(v \otimes u^*) &= S(g_1, g_2)(\langle v, \rho(x)u \rangle) \\
&= \langle v, \rho(g_1^{-1} x g_2)u \rangle \\
&= \langle \rho(g_1)v, \rho(x)\rho(g_2)u \rangle && \text{since } \rho \text{ is unitary} \\
&= T(\rho(g_1)v \otimes (\rho(g_2)u)^*) \\
&= T(\rho(g_1)v \otimes (u^* \circ \rho(g_2^{-1}))) && \text{since } \rho \text{ is unitary} \\
&= T(\rho(g_1)v \otimes \rho^t(g_2)u^*) \\
&= (T \circ (\rho \otimes \rho^t)(g_1, g_2))(v \otimes u^*)
\end{aligned}$$

For each $\rho \in \widehat{G}$, it can be shown by Schur's Lemma that $\rho \otimes \rho^t$ is an irreducible representation of $G \times G$ (since ρ is irreducible). Denoting by $T_{\rho \otimes \rho^t} : V_\rho \otimes V_\rho^* \to C_{\text{alg}}(G)_{\rho \otimes \rho^t}$ the domain and range restriction of T to the $\rho \otimes \rho^t$-isotypic components, Schur's Lemma implies that $T_{\rho \otimes \rho^t}$ is (in particular) injective, and hence, T is injective.

Now we show that T is an isometry. For each $\rho \in \widehat{G}$ let $E^\rho = \{e_1^\rho, \cdots, e_{d_\rho}^\rho\}$ be an orthonormal basis for V_ρ. Then let $A_{ij}^\rho \in \text{End}(V_\rho)$ be the operator whose matrix relative to E^ρ has 0s everywhere but in the ijth cell, where it has a 1. Recalling our identification $\text{End}(V_\rho) = V_\rho \otimes V_\rho^*$, we have $T(A_{ij}^\rho) = \sqrt{d_\rho}\rho(x)_{ij}$ (where we are also thinking of $\rho(x)$ as the matrix relative to E^ρ). Now let $\rho, \sigma \in \widehat{G}$. It suffices to show that

$$\langle T(A_{ij}^\rho), T(A_{i'j'}^\sigma) \rangle = \begin{cases} \delta_{ii'}\delta_{jj'} & \text{if } \rho = \sigma \\ 0 & \text{if } \rho \neq \sigma \end{cases}$$

The left-hand-side simplifies to

$$\sqrt{d_\rho d_\sigma} \int_X e_i^{\rho *} \rho(x) e_j^\rho \overline{e_{i'}^{\sigma *} \sigma(x) e_{j'}^\sigma} d\mu(x) = \sqrt{d_\rho d_\sigma}\, \langle \Phi, e_{i'}^\sigma e_i^{\rho *} \rangle$$

where

$$\Phi = \int_X \rho(x) e_j^\rho e_{j'}^{\sigma *} \sigma(x)^{-1} d\mu(x) \,:\, V_\sigma \to V_\rho$$

By Schur's lemma, $\Phi = 0$ if $\rho \neq \sigma$ and Φ is otherwise a homothety, and the trace reveals:

$$\Phi = \frac{1}{d_\rho}\delta_{jj'}$$

as required. ■

We state the Peter-Weyl Theorem without proof, just because it is convenient.

Theorem 4.4.3 (Peter-Weyl Theorem)**.** $C_{\text{alg}}(G)$ embeds into a dense subspace of $L^2(G)$.

For each $\rho \in \widehat{G}$, let E^ρ be an orthonormal basis for V_ρ. The Algebraic Peter-Weyl Theorem says that for each $f \in C_{\text{alg}}(G)$, there exist constants c^ρ_{ij}, all but finitely many nonzero, such that

$$f(x) = \sum_{\rho \in \widehat{G}} \sum_{i,j=1}^{d_\rho} c^\rho_{ij} \rho_{ij}(x) \tag{4.4}$$

And moreover, the functions $\{\sqrt{d_\rho}\rho_{ij}(x) \;:\; \rho \in \widehat{G}, 0 \leq i,j \leq d_\rho\}$ are orthonormal. The full Peter-Weyl Theorem asserts that (4.4) holds for all $f \in L^2(G)$ where we may have infinitely many nonzero constants c^ρ_{ij} and convergence is with respect to the L^2 norm.

Take, for example, $G = S^1$. Then the Peter-Weyl theorem states that any 2π-periodic function has a Fourier series. There are countably many spaces, naturally indexed by $\mathbb{Z}$. For each $r \in \mathbb{Z}$, the space indexed by r is one dimensional and spanned by

$$e^{irx}$$

The hypersphere S^3 also has the structure of a compact group, so we get a similar decomposition of functions into orthogonal matrix coefficients, called hyperspherical harmonics. These are the only two spheres with a group structure, but we will soon see how spherical harmonics arise for spheres of any dimension by realizing S^{d-1} as a quotient. The reader is referred to [Gal] for a detailed account of spherical harmonics.

4.5 Gelfand Pairs

In this section we continue to study 2-point homogeneous spaces $X = G/H$ with the observation that the stabilizer H is a Gelfand subgroup of G. Much of the content in this section applies more generally to any Gelfand pair (defined as follows).

Definition 4.5.1. Let G be a compact group and let H be a subgroup. (G, H) is a *Gelfand pair* if for all $\rho \in \widehat{G}$, the multiplicity of the trivial representation in $\operatorname{Res}_H^G(\rho)$ is at most 1.

Gelfand pairs can be defined this way for other classes of groups besides compact groups given an appropriate definition for $\widehat{G}$. For a given class of groups, there are some equivalent formulations of the Gelfand property. We note one such property for compact groups.

Definition 4.5.2. Let G be a compact group and let H be a subgroup. For our purposes, the *Hecke algebra* of (G, H) is the associative algebra $L^1(H\backslash G/H)$ of integrable functions on $H\backslash G/H$ under *convolution*, which is defined as follows.

$$(f * g)(x) = \int_G f(y)g(y^{-1}x)d\mu(y)$$

Proposition 4.5.1. Let G be a compact group and let H be a subgroup. Then (G, H) is a Gelfand pair if and only if its Hecke algebra is commutative.

We will use the restriction and induction functors along with Frobenius reciprocity, which all still make sense in the context of compact groups.

Proof. Let $\eta = \operatorname{Ind}_H^G(1)$ be the induction of the trivial representation. The map

$$\begin{aligned} L^1(H\backslash G/H) &\to \operatorname{Hom}_G(\eta, \eta) =: A \\ f &\mapsto \int_G f(y)\eta(y)d\mu(y) =: \eta(f) \end{aligned} \tag{4.5}$$

is an isomorphism of associative algebras [Tol44]. By Schur's lemma, A is commutative if and only if η is multiplicity-free. By Frobenius reciprocity, for any $\rho \in \widehat{G}$,

$$\dim \operatorname{Hom}_G(\operatorname{Ind}_H^G(1), \rho) = \dim \operatorname{Hom}_H(1, \operatorname{Res}_H^G(\rho)). \tag{4.6}$$

η is multiplicity-free if and only the left-hand-side of (4.6) is ≤ 1 for all $\rho \in \widehat{G}$, and (G, H) is a Gelfand pair if and only if the right-hand-side of (4.6) is ≤ 1 for all $\rho \in \widehat{G}$. ■

The image $\eta(f)$ of f under the map (4.5) is defined for any $f \in L^1(G)$ and any representation η of G, and is called the *Fourier transform of f at η*. We can use this characterization of the Gelfand property to prove that (G, H) is always a Gelfand pair for 2-point homogeneous space $X = G/H$.

Proposition 4.5.2 (Gelfand's lemma)**.** Let G be a compact group with subgroup H. Let $\sigma : G \to G$ be an antihomomorphism ($\sigma(gg') = \sigma(g')\sigma(g)$ for all $g, g' \in G$) such that $\sigma(g) \in HgH$ for all $g \in G$. Then (G, H) is a Gelfand pair.

Proof. For $f \in L^2(G)$, define $f^\sigma(x) = f(\sigma x)$. Then for any $f, g \in L^1(H\backslash G/H)$ one has $(g * f)^\sigma = f^\sigma * g^\sigma$ and $f^\sigma = f$, and so $f * g = f^\sigma * g^\sigma = (g * f)^\sigma = g * f$. Thus $L^1(H\backslash G/H)$ is commutative. ■

Proposition 4.5.3. Let $X = G/H$ be a 2-point homogeneous space with $H = \mathrm{Stab}(x_0)$, $x_0 \in X$. Then (G, H) is a Gelfand pair.

Proof. Suppose that $y_1, y_2 \in X$ are equidistant from x_0. Since (X, G) is 2-point homogeneous, there exists $h \in H$ such that $y_2 = hy_1$, so the images of y_1 and y_2 in $H\backslash G/H$ are equal, *i.e.*, double cosets HgH depend only on $\tau(x_0, gx_0) = \tau(g^{-1}x_0, x_0) = \tau(x_0, g^{-1}x_0)$. Thus $HgH = Hg^{-1}H$, so we can apply proposition (4.5.2) with σ as the inversion map. ■

For the rest of the section assume that $X = G/H$ is a 2-point homogeneous space, where H is the stabilizer in G of a generic point $x_0 \in X$. It follows that (G, H) is a Gelfand pair.

The complex-valued functions on the quotient set G/H can be naturally identified with the complex-valued functions on G that are constant on left cosets of H. Moreover, the Haar measure μ on G induces a measure, which we also denote by μ, on G/H. When the measure on G/H is suitably normalized, for any measurable subset $S \subseteq G/H$ and any f, the two interpretations of the integral

$$\int_S f(x) d\mu(x)$$

coincide. If V can be thought of as a subspace of $L^2(G)$, we denote by V^H the subspace of functions invariant under right translation by H, and by ${}^H V$ the subspace of functions invariant under left translation by H. If V is a vector space carrying a single G-representation, then V^H or equally ${}^H V$ denotes the H-invariant subspace, and it will be apparent why we choose to write the H on the left or the right.

The set of square integrable functions on G/H inherits an inner product from $L^2(G)$, which agrees with the inner product on $L^2(G)^H$. Thus, we may equate inner product spaces of functions on G/H with inner product spaces of right-H-invariant functions on G.

$$L^2(G/H) = L^2(G)^H \qquad\qquad C_{\text{alg}}(G/H) = C_{\text{alg}}(G)^H$$

Invariance under H goes through the isomorphism of representations given by theorem (4.4.2), and we can simplify using the Gelfand property. That is,

$$C_{\text{alg}}(G/H) = C_{\text{alg}}(G)^H \overset{G\times 1}{\cong} \bigoplus_{\rho\in\widehat{G}} V_\rho \otimes V_\rho^{*H} = \bigoplus_{\rho\in\widehat{G}_H} V_\rho$$

where $\widehat{G}_H = \{\rho \in \widehat{G} \,|\, \dim V_\rho^H = 1\}$. Now, for each $\rho \in \widehat{G}_H$, choose an ordered orthonormal basis $E^\rho = (e_1^\rho, \cdots, e_{d_\rho}^\rho)$ satisfying $\rho(H)(e_\rho^1) = e_\rho^1$. The first column of $\rho(x)$ is invariant under right multiplication by $\rho(H)$ and the first row of $\rho(x)$ is invariant under left multiplication by $\rho(H)$. Thus, $C_{\text{alg}}(G/H)$ is spanned by the functions $\rho_{i1}(x)$ for $\rho \in \widehat{G}_H$ and $i = 1, \cdots, d_\rho$.

By the Peter-Weyl theorem, essentially the same reasoning carries out where we replace the notation C_{alg} with L^2 and introduce closures in the necessary places. This gives us another kind of generalized Fourier series, this time for Gelfand pairs. For each $f \in L^2(G/H)$ there exist constants c_i^ρ such that

$$f(x) = \sum_{\rho\in\widehat{G}_H} \sum_{i=0}^{d_\rho} c_i^\rho \rho_{i1}(x) \tag{4.7}$$

Again the set of functions $\{\sqrt{d_\rho}\rho_{i1}(x) \,|\, \rho \in \widehat{G}_H, 0 \le i \le d_\rho\}$ is orthonormal, and convergence is with respect to the L^2 norm, this time for $L^2(G/H)$. As an example, S^{d-1} can be constructed as such a quotient space, and equation (4.7) gives us a Fourier series in terms of spherical harmonics.

Now consider the function $J_\rho : G \times G \to \mathbb{C}$ given by

$$J_\rho(x, y) = d_\rho \langle \rho_1(x), \rho_1(y) \rangle$$

Since orthonormal bases differ by unitary matrices, J_ρ does not depend on our choices of $\widehat{G}$ or orthonormal bases E^ρ. Since $\rho_1(x)$ is right invariant by H, $J_\rho(xH, yH) = J_\rho(x, y)$, so we can view J_ρ as a function $X \times X \to \mathbb{C}$, where $X = G/H$. Also since ρ is a unitary representation, $J_\rho(gx, gy) = J_\rho(x, y)$ for any $g \in G$.

Note that $J_\rho(x, y)$ depends only on $\tau(x, y)$. Then, since $\tau(x, y) = \tau(y, x)$, it follows that $J_\rho(x, y) = J_\rho(y, x)$, so J_ρ is real-valued. That is, there is a function $z_\rho : \mathbb{R} \to \mathbb{R}$, (which also does not depend on our choices of $\widehat{G}$ or E^ρ) such that

$$z_\rho(\tau(x, y)) = J_\rho(x, y)$$

Alternately, $J_\rho(x, y)$ depends only on $x^{-1}y$, so for instance, we may write

$$J_\rho(x, y) = z_{\rho,x_0}(x_0 x^{-1} y)$$

for some function $z_{\rho,x_0} : X \to \mathbb{R}$. It will momentarily become apparent why we make the function z_{ρ,x_0} dependent on x_0 in this seemingly artificial way. First, we get the familiar property

$$z_{\rho,x_0}(y) = z_\rho(\tau(x_0, y))$$

By the definition of J_ρ it is clear that $z_{\rho,x_0} \in C_{\text{alg}}(X)$. But the main reason for defining z_{ρ,x_0} like this is because $x_0 = \pi(1)$ where $\pi : G \to G/H = X$ is the quotient map. From this, it follows that

$$z_{\rho,x_0}(Hy) = z_{\rho,x_0}(y)$$

Therefore $z_{\rho,x_0} \in C_{\text{alg}}(H \backslash G/H)$. Again, the H invariance factors through the isomorphism given by Theorem (4.4.2), that is,

$$C_{\text{alg}}(H \backslash G/H) = {}^H C_{\text{alg}}(G)^H \overset{\mathbb{C}}{\cong} \bigoplus_{\rho \in \widehat{G}} {}^H V_\rho \otimes V_\rho^{*H} = \bigoplus_{\rho \in \widehat{G}_H} \mathbb{C} z_{\rho,x_0}$$

The functions z_{ρ,x_0} for $\rho \in \widehat{G}_H$ are therefore orthogonal with respect to the inner product on $C_{\text{alg}}(X)$.

Let us now observe that the addition rule is satisfied by z_ρ for all $\rho \in \widehat{G}_H$. Unwinding definitions, we see that

$$z_{\rho,x_0}(x) = d_\rho \rho(x)_{11}\overline{\rho(x_0)_{11}} + \cdots + d_\rho \rho(x)_{1d_\rho}\overline{\rho(x_0)_{1d_\rho}}$$

and $\{\sqrt{d_\rho}\,\rho(x)_{11}, \cdots, \sqrt{d_\rho}\,\rho(x)_{1d_\rho}\}$ is an orthonormal basis for the subspace of $C_{\text{alg}}(X)$ identified with V_ρ. It follows that $z_{\rho,x_0}(x)$ is the reproducing kernel at x_0 for the ρ-isotypic subspace $C_{\text{alg}}(X)_\rho$. Since additionally, the V_ρ are orthogonal, the addition rule holds. In order to conclude that X is a Q-polynomial code space, we now have to show that the zonal orthogonal polynomials z_r, defined in Section 4.3, are sums of these z_ρ for various $\rho \in \widehat{G}_H$. (Caution: *a priori*, we don't know that the z_r satisfy the addition rule.) I could only derive a *criterion* for this, which is albeit satisfied in all our spaces of interest. In general, it is conceivable that X is merely distance regular without being Q-polynomial. However, it could be that I didn't find a full proof only because I didn't try using the Laplacian, which can be defined for these spaces (generalizing the case of the sphere).

We will again use the function spaces $\text{Zon}(X)$ and $\text{Harm}(X, r)$ defined in section 4.1, with inner product inherited from their embedding in $L^2(X)$.

Theorem 4.5.1. $\text{Harm}(X, r)$ is a representation of G under left translation.

Proof. We show that $\text{Harm}(X, r)$ is G-invariant using strong induction on r. $\text{Harm}(X, 0) = \mathbb{C}$, which is invariant. Suppose $\text{Harm}(X, i)$ is invariant for $i = 0, \cdots, r$. Let $f_x \in \text{Harm}(X, r+1)$ be an element of the spanning set consisting of the zonal harmonic polynomial functions at varying points $x \in X$. Let $h_y \in \text{Pol}(X, r) \cap \text{Zon}(X) = \sum_{i=1}^{r} \text{Harm}(X, r)$, similarly, and let $g \in G$. It now follows from the induction hypothesis that $g^{-1}h_y \in \text{Pol}(X, r) \cap \text{Zon}(X)$. Now since the action is unitary, $\langle gf_x, h_y\rangle = \langle f_x, g^{-1}h_y\rangle = 0$. ■

In all of our spaces of interest, it turns out that $\text{Harm}(X, r)$ is irreducible for all $r \in \mathbb{N}$.

To see that $\text{Harm}(X, r)$ is irreducible for a particular 2-point homogeneous space X, consider the orthogonal projection f of z_{r,x_0} onto a nonzero invariant subspace V. Without loss of generality assume $f(\alpha_0) > 0$, and renormalize f such that $\langle f, f\rangle = f(\alpha_0)$. It suffices to show f is zonal at x_0, as this implies $f = z_{r,x_0}$ and hence $V = \text{Harm}(X, r)$. Using the fact that f is left H invariant, this could be easily verified in the cases of $S^{d-1}, \mathbb{RP}^{d-1}$, and $\mathbb{CP}^{d-1}$, for example.

The rest of the results of this section require $\mathrm{Harm}(X, r)$ to be finite dimensional. The following proposition requires the condition of irreducibility, c.f. Theorem (4.4.1).

Proposition 4.5.4 (Criterion for the addition rule)**.** If $\mathrm{Harm}(X, r)$ is irreducible, then z_{r,x_0} is its reproducing kernel.

Proof. From our classification of all finite dimensional irreducible representations of G, which are in fact *all* irreducible representations of G, by Theorem (4.4.1), it follows that ${}^H\mathrm{Harm}(X, r)$ is one dimensional. Since $\mathrm{Harm}(X, r)$ contains both z_{r,x_0} and $k_{(x_0)}$, there exists nonzero $c \in \mathbb{C}$ such that $z_{r,x_0} = ck_{(x_0)}$. But then

$$z_r(\alpha_0) = \langle z_{r,x_0}, z_{r,x_0}\rangle = \langle z_{r,x_0}, ck_{(x_0)}\rangle = \bar{c}z_r(\alpha_0)$$

hence $c = 1$. ■

Note that $\mathrm{Harm}(X, r)$ is not necessarily irreducible for general Gelfand pairs (G, H) [Mac98]. In any case, let us establish a few relations among our various sets of functions.

Theorem 4.5.2. If $\mathrm{Harm}(X, r)$ is finite dimensional for all $r \in \mathbb{N}$, then $\mathrm{Zon}(X) \subseteq C_{\mathrm{alg}}(X)$.

Proof. Assume $\mathrm{Harm}(X, r)$ is finite dimensional for all $r \in \mathbb{N}$. Now fix $r \in \mathbb{N}$. There exist finitely many $\rho_1, \cdots, \rho_t \in \widehat{G}$, presumably with repetitions, and an isometric G-isomorphism

$$\phi : \mathrm{Harm}(X, r) \to \bigoplus_{i=1}^{t} W(\rho_i)$$

where each $W(\rho_i)$ is an irreducible left translation invariant subspace of $C_{\mathrm{alg}}(G)$ of type ρ_i. It suffices to show that ϕ is the identity map and that each ρ_i is left H invariant. Recall that $\eta = \mathrm{Ind}_H^G(1)$ is multiplicity-free. Then by Frobenius reciprocity,

$$\begin{aligned} |\{i \in [t] \,:\, \rho_i \in \widehat{G}_H\}| &= \dim {}^H\mathrm{Harm}(X, r) \\ &= \dim \mathrm{Hom}_H(\mathrm{Res}_H^G \mathrm{Harm}(X, r), 1) \\ &= \dim \mathrm{Hom}_G(\mathrm{Harm}(X, r), \eta) \\ &= |\{i \in [t] \,:\, \rho_i \in \widehat{G}\}| \end{aligned}$$

It follows by counting that for each $i \in [t]$, $\rho_i \in \widehat{G}_H$. Let $f_1, \cdots, f_t \in \mathrm{Harm}(X, r)$ be the functions mapping to $z_{\rho_1,x_0}, \cdots, z_{\rho_t,x_0} \in \bigoplus W(\rho_i)$, respectively. Left H-invariance factors through ϕ. Let $k_{(x_0)}$ be the reproducing kernel for $\mathrm{Harm}(X, r)$ at x_0. One can easily show that for all $g \in G$, $gk_{(x_0)} = k_{(gx_0)}$. In particular, by uniqueness of the reproducing kernel, $k_{(x_0)}$ is left H invariant. Since $\mathrm{span}\{f_1, \cdots, f_t\} = {}^H\mathrm{Harm}(X, r)$, we can write $k_{(x_0)}$ as a linear combination of $f_1, \cdots, f_t$, and solving for the constants we readily find

$$k_{(x_0)} = f_1 + \cdots + f_t$$

Define $z_{\rho,x_0} := z_{\rho_1,x_0} + \cdots + z_{\rho_t,x_0}$. This is the reproducing kernel of $\bigoplus W(\rho_i)$ at x_0 and the image of $k_{(x)}$. It follows that for all $f \in \mathrm{Harm}(X, r)$ and $x \in X$ we have

$$f(x) = \langle f, k_{(x)} \rangle = \langle \phi(f), z_{\rho,x} \rangle = \phi(f)(x)$$

so ϕ is the identity map. ■

Note, by the way, that it follows that $\rho_1, \cdots, \rho_t$ are distinct (and hence non-isomorphic as they come from $\widehat{G}$). I am not sure whether the opposite inclusion holds, but if so, then $\mathrm{Zon}(X), \mathrm{Pol}(X)$, and $C_{\mathrm{alg}}(X)$ should all coincide, at least if X is finite:

Theorem 4.5.3. If X is finite and $\mathrm{Zon}(X) = C_{\mathrm{alg}}(X)$, then $\mathrm{Zon}(X) = \mathrm{Pol}(X)$.

As mentioned in Section (4.1), having $\mathrm{Zon}(X) = \mathrm{Pol}(X)$ is helpful for calculating the absolute bound. This property holds in all our spaces of interest (even the infinite ones).

Proof. It is proved in [Mac98] that for all $x, y \in X$ and $\rho \in \widehat{G}_H$,

$$z_{\rho,x} \cdot z_{\rho,y} = \frac{1}{|H|} \sum_{h \in H} z_{\rho,xhy}$$

While the sum can undoubtedly be replaced with a Haar integral for general compact groups, it is only clear that the right-hand-side is an element of $C_{\mathrm{alg}}(X)$ (and hence $\mathrm{Zon}(X)$, by hypothesis) when the sum is finite. Since $z_{r,x}$ is a linear combination of functions $z_{\rho,x}$ for various ρ, the result follows. ■

4.6 Examples

In order to apply the LP bound in a Q-polynomial code space X, one first uses the inner product on $\text{Zon}(X)$ to derive the zonal orthogonal polynomials, and then applies Theorem (4.3.1). This involves computing inner products of some zonal polynomial functions centered at the same point x_0. In section 3, we saw that for a code space X, relatively few restrictions must be placed on the inner product on $\text{Zon}(X)$ in order for the LP bounds to hold in X. It's just that for a bad inner product the LP bounds might be weak or even completely trivial.

In the three examples shown here, we apply the LP bounds in the spaces $X = S^{d-1}, \mathbb{RP}^{d-1}$, and $\mathbb{CP}^{d-1}$. As 2-point homogeneous spaces, we have previously defined the inner product on $\text{Zon}(X)$ to be the one induced by the Haar measure on the automorphism groups $G = \text{SO}(d), \text{SO}(d)$, and $\text{SU}(d)$. We do not (yet) have an explicit formula to compute this inner product, as we have only alluded to the *existence* of Haar measure. As noted above, it would instead be sufficient to somehow come up with a *good* inner product (one that makes X Q-polynomial and that will lead to a meaningful LP bound). We can do one better: with the following uniqueness theorem, we can come up with *the same* inner product, but presented in a way so as to allow the explicit calculation of the inner product of zonal polynomials. This has the advantage of demonstrating that the inner product induced by Haar measure is, in fact, good.

Proposition 4.6.1. Let X be a 2-point homogeneous space with automorphism group G. Let μ be the measure on X inherited from the Haar measure on G, and let σ be another normalized measure on X left-invariant under G. Then $\sigma = \mu$.

We omit the proof to save space as it is an easy but space-consuming application of Fubini's Theorem. (Interestingly, uniqueness of Haar measure is not used!)

Example 4.6.1. As our first example, we apply the LP bounds to the sphere S^{d-1}. Cartesian coordinates $x_1, \cdots, x_d$ and spherical coordinates $r, \theta_{d-1}, \cdots, \theta_2, \phi_1$ for $\mathbb{R}^d$ are related by the system of equations

$$\begin{aligned} x_1 &= r\cos\theta_{d-1} \\ x_2 &= r\sin\theta_{d-1}\cos\theta_{d-2} \\ x_3 &= r\sin\theta_{d-1}\sin\theta_{d-2}\cos\theta_{d-3} \\ &\vdots \\ x_{d-1} &= r\sin\theta_{d-1}\cdots\sin\theta_2\cos\phi_1 \\ x_d &= r\sin\theta_{d-1}\cdots\sin\theta_2\sin\phi_1 \end{aligned}$$

In presence of rotational symmetry it is sometimes useful to change from Cartesian to spherical coordinates. When doing so, one must multiply the integrand by the Jacobian determinant

$$J_d = r^{d-1}\sin^{d-2}\theta_{d-1}\cdots\sin^1\theta_2$$

Consider now the inner product (c.f. Section 4.3) of two zonal polynomial functions f_{x_0}, g_{x_0} on the sphere centered at x_0. (Recall that by regularity, this quantity does not depend on x_0, only on f and g, and defines a possibly degenerate inner product on $\mathbb{R}[t]$, which we denote $(\cdot,\cdot)$.)

$$(f,g) = \langle f_{x_0}, g_{x_0}\rangle = \int_{S^{d-1}} f(\langle x_0, x\rangle)g(\langle x_0, x\rangle)d\mu(x)$$

By Proposition (4.6.1), we can replace $d\mu(x)$ by $\frac{1}{s_{d-1}}dx_1\cdots dx_d$ or better yet by $\frac{1}{s_{d-1}}J_d dr d\theta_{d-1}\cdots d\theta_2\phi_1$, where s_{d-1} is the surface area of the unit $d-1$ sphere. Without loss of generality, assume $\theta_{d-1} = \langle x_0, x\rangle =: t$. Taking note of the recurrence $J_d = r\sin^{d-2}\theta_{d-1}J_{d-1}$, we get the following formula. The integral is over the sphere, so we drop factors of r and dr.

$$\begin{aligned} (f,g) &= \frac{1}{s_{d-1}}\int_{S^{d-1}} f(\theta_{d-1})g(\theta_{d-1})\sin^{d-2}\theta_{d-1}J_{d-1}\,d\theta_{d-1}\cdots d\theta_2 d\phi_1 \\ &= \frac{1}{s_{d-1}}\int_{S^{d-2}} J_{d-1}rd\theta_{d-2}\cdots d_{\theta_2}d\phi_1\int_{-1}^{1} f(t)g(t)(1-t^2)^{\frac{d-3}{2}}dt \\ &= \frac{s_{d-2}}{s_{d-1}}\int_{-1}^{1} f(t)g(t)(1-t^2)^{\frac{d-3}{2}}dt \end{aligned}$$

It follows that the zonal orthogonal polynomials are the *Gegenbauer polynomials*, which obey the following recurrence relation [DGS77].

$$\begin{aligned}\lambda_{r+1}z_{r+1}(t) = tz_r(t) - (1-\lambda_{r-1})z_{r-1}(t) \text{ where } \lambda_r = \frac{r}{d+2r-2}\\ z_0(t) = 1\\ z_1(t) = dt\end{aligned}$$

The next couple Gegenbauer polynomials are

$$\begin{aligned}z_2(t) &= \frac{1}{2}(d+2)(dt^2-1)\\ z_3(t) &= \frac{1}{6}d(d+4)((d+2)x^3-3x)\end{aligned}$$

The Gegenbauer polynomials are the Jacobi polynomials $J_r^{a,b}(t)$ with $a = b = (d-3)/2$. The following two examples, first shown in [DGS77], demonstrate the power of the LP bounds.

Consider a $(-1,-\alpha]$-code $\mathcal{C}$ of cardinality n in S^{d-1}, where $\alpha \in (0,1)$. This means $\mathcal{C}$ is a unit simplex of n points in $\mathbb{R}^d$, and in particular this includes the case of a $-\alpha$-code, i.e. regular simplex. Set $f(t) = t+\alpha$. Then $f(\gamma) \leq 0$ for all $\gamma \in (-1,-\alpha]$, and $f(t) = \frac{1}{d}z_1(t) + \alpha z_0(t)$. Since $\alpha > 0$ and $\frac{1}{d} \geq 0$, Theorem (4.3.1) yields

$$n \leq 1 + \frac{1}{\alpha}$$

The bound is attained by a regular unit simplex with $d+1$ vertices centered at the origin of $\mathbb{R}^d$. (This is a $-\alpha$-code with $\alpha = \frac{1}{d}$.) We will see other such *simplex bounds* below, for the space $\mathbb{RP}^{d-1}$ and $\mathbb{CP}^{d-1}$.

Continuing with the sphere, we derive a slight generalization of the Welch bound for real equiangular lines. Let $-1 < \alpha_1 \leq \alpha_2 < 1$ with $\alpha_1 + \alpha_2 \leq 0$ and $\alpha_1\alpha_2 \geq -\frac{1}{d}$. Consider an $[\alpha_1,\alpha_2]$-code $\mathcal{C}$ of cardinality n in S^{d-1}. Set $f(t) = (t-\alpha_1)(t-\alpha_2)$, so that $f(\gamma) \leq 0$ for all $\gamma \in [\alpha_1,\alpha_2]$. This time we have

$$f(t) = \frac{2}{d(d+2)}z_2(t) - \frac{\alpha_1+\alpha_2}{d}z_1(t) + \left(\alpha_1\alpha_2 + \frac{1}{d}\right)z_0(t)$$

The Gegenbauer coefficients are non-negative and the coefficient f_0 of the constant term is positive, so by Theorem (4.3.1), we get

$$n \leq \frac{d(1-\alpha_1)(1-\alpha_2)}{d\alpha_1\alpha_2 + 1} \tag{4.8}$$

If $\alpha_1 = -\alpha_2$ we obtain the first Welch bound for real equiangular lines. (Of course, the exact same bound with $\alpha_1 = -\alpha_2$ holds for complex equiangular lines, but note that we have only proved it here for the real case. While the real case can be dealt with spherically, the complex case must be dealt with projectively.) Note that equality holds in (4.8) if and only if $\mathcal{C}$ is an ETF. It is not hard to see that the polynomial $f(t)$ derived above is optimal, and therefore, an equiangular code $\mathcal{C}$ is an ETF if and only if it is a tight spherical code.

Two other kinds of spherical $\{\alpha_1, \alpha_2\}$-codes arise in this thesis. First, in Section 3.3 we encounter one with $\alpha_1 = -1/t$ and $\alpha_2 = 0$. Plugging this in gives $n \leq d + d/t$. Second are MUBs, which also happen to be related to this first kind. Suppose $\{\mathcal{B}_1, \cdots, \mathcal{B}_t\}$ is a set of t MUBs for $\mathbb{C}^d$. This means the projectors $uu^\dagger$ for $u \in \bigcup \mathcal{B}_i$ form a $\{0, 1/d\}$-code in the unit sphere S^{d^2-1} in $H_{\mathbb{C},d}$. However, since $\alpha_1 + \alpha_2 \not\leq 0$, we cannot directly apply (4.8). It turns out that the projectors $uu^\dagger$ form a spherical code in the trace 1 hyperplane centered at the maximally mixed state $\frac{1}{d}I$. This suggests that we might instead consider the following associated normalized traceless matrices, and indeed, these satisfy $\alpha_1 + \alpha_2 \leq 0$.

$$\Gamma_u = \frac{d}{d-1}\left(uu^\dagger - \frac{1}{d}I\right)$$

$\{\Gamma_u \,|\, u \in \bigcup \mathcal{B}_i\}$ is a $\{-1/(d-1), 0\}$-code in the unit sphere S^{d^2-2} of the traceless subspace of $H_{\mathbb{C},d}$, (putting us back in the first regime,) and taking $n = dt$, it follows by (4.8) that $t \leq d+1$. This is the MUB bound.

Example 4.6.2. Real and complex projective space can be handled as we did the sphere. Recall that we use the separation function $\tau([u],[v]) = |\langle u, v\rangle|^2$ (the *squared* overlap). In the end, it turns out that the zonal orthogonal polynomials for $X = K\mathbb{P}^{d-1}$ are given by the Jacobi polynomials $P_r^{a,b}(t)$:

$$z_r(t) = P_r^{a,b}(2t-1)$$

with $(a,b) = ((d-3)/2, -1/2)$ resp. $(a,b) = (d-2, 0)$ for $K = \mathbb{R}$ resp. $\mathbb{C}$ [TZ12]. The first couple Jacobi polynomials are

$$\begin{aligned} z_0(t) &= 1 \\ z_1(t) &= (a+1) + (a+b+2)(t-1) \end{aligned}$$

As in the case of the sphere, we can derive a bound on the maximum number of points in a simplex in $K\mathbb{P}^{d-1}$ with prescribed distances/angles. Consider a $[0,\beta)$-code in $K\mathbb{P}^{d-1}$. Set $f(t) = t - \beta$. Then

$$\begin{aligned} f(t) &= \frac{2}{d} z_1(t) + \left(\frac{1}{d} - \beta\right) z_0(t) && \text{if } K = \mathbb{R}, \text{ and} \\ f(t) &= \frac{1}{d} z_1(t) + \left(\frac{1}{d} - \beta\right) z_0(t) && \text{if } K = \mathbb{C}. \end{aligned}$$

We thereby obtain the first Welch bound for real resp. complex equiangular lines by applying Theorem (4.3.1). Interestingly, the bound does not depend on whether $K = \mathbb{R}$ or $\mathbb{C}$.

$$n \leq \frac{d(1-\beta)}{1-d\beta}$$

Again it is not hard to see that f is optimal, and so a set of equiangular lines forms an ETF if and only if it is a tight projective code. Let us now prove the entire sequence of Welch bounds. In order to avoid explicitly calculating the Jacobi polynomials and finding all the zonal coefficients of f, we use the orthogonality relations, which allow us to calculate only f_0. The Jacobi polynomials are orthogonal with respect to the inner product

$$(f,g) = c_{a,b} \int_{-1}^{1} f(t)g(t)(1-t)^a(1+t)^b dt$$

where the constant $c_{a,b}$ is given in terms of the gamma function:

$$c_{a,b} = \frac{\Gamma(a+b+2)}{2^{a+b+1}\Gamma(a+1)\Gamma(b+1)}$$

$c_{a,b}$ is chosen such that $z_0 = 1$ has length 1. In fact, as with the sphere, one can argue by Proposition (4.6.1) that the above inner product is the one defined in Section 4.3. Let us derive the generalized Welch bound for the complex case, as the formulas are slightly simpler. The same process can be carried out in real projective space or the sphere, but it is unnecessary since the bound holding in the complex case implies that it holds in the real case through the inclusion $\mathbb{R}^d \subseteq \mathbb{C}^d$. Thus we set $a = d-2$ and $b = 0$. Then $c_{a,b} = (d-1)/2^{d-1}$. Let $u = 2t-1$. Set $f(t) = t^k - \beta^k$. Then

$$f_0 = (f, 1) = \frac{d-1}{2^{d-1}} \int_{-1}^{1} \left(\left(\frac{u+1}{2} \right)^k - \beta^k \right) (1-u)^{d-2} du$$

In the $k = 1$ case it is easy to see (substituting $v = 1-u$) that the expression evaluates to $1/d - \beta$, just as we found before. Using integration by parts, it can be shown by induction on k that in general,

$$f_0 = \frac{1}{\binom{d+k-1}{k}} - \beta > 0$$

Assuming the condition that for each $r = 1, \cdots, k$ that $(f, z_r) \geq 0$, Theorem (4.3.1) implies that the general Welch bounds hold. Checking this condition involves a complicated integral and could be checked on a computer.

Remark 4.6.1. Futher examples (involving the Hamming and Johnson schemes) have been omitted due to lack of space and time.

Chapter 5

Constructions

In this section we provide examples of codes, which are referred to throughout the entire thesis. Some of these examples provide a tight lower bound on $N(X)$ or $N(X, A)$, the maximum cardinality of a simplex in a Q-polynomial code space X, or the maximum cardinality of a code with distances in a set A. We do not include *all* the constructions needed to fully justify all of our claims, but the missing constructions can readily be found in the literature. The examples shown here are intended to demonstrate certain interesting features.

5.1 Real equiangular configurations

Example 5.1.1. The simplest example of rigid equiangular lines is three lines in $\mathbb{R}^2$.

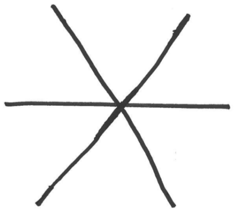

There are no four equiangular lines in $\mathbb{R}^2$. The three lines in $\mathbb{R}^2$ are part of a family of configurations: the lines through the $d+1$ points of a regular simplex centered at the origin of $\mathbb{R}^d$. This configuration of lines corresponds to the graphs switching equivalent to the empty graph, namely the complete bipartite graphs.

Example 5.1.2. A configuration of six vectors, one spanning each of the lines through the six pairs of antipodal vertices of a regular icosahedron in $\mathbb{R}^3$, has the following Seidel matrix and graph.

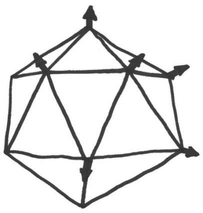

$$\begin{pmatrix} & + & + & + & + & + \\ + & & + & - & - & + \\ + & + & & + & - & - \\ + & - & + & & + & - \\ + & - & - & + & & + \\ + & + & - & - & + & \end{pmatrix}$$

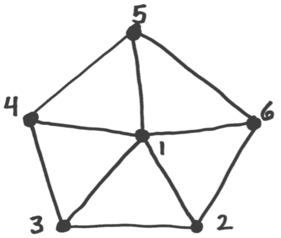

This is an ETF since 6 is the absolute bound. This ETF is of the same configuration as its complementary ETF. This is the only configuration of 6 equiangular lines in $\mathbb{R}^3$. This can be seen by the following proposition and the following two observations. First note that a graph Y containing an induced subgraph switching equivalent to K_4 has dimension ≥ 4. Second, it is easy to check (there are just four cases) that if Y has more than 4 vertices and contains an induced subgraph switching equivalent to $\overline{K}_4$, then it has dimension ≥ 4.

Proposition 5.1.1. Let $d \geq 3$ and $n \geq 6$. Then a graph on n vertices contains an induced subgraph switching equivalent to either $\overline{K}_4$, K_4, or $C_5 \sqcup K_1$.

Proof. Consider a subgraph induced by six of the vertices. Assume one of the vertices is isolated and that the graph contains no induced subgraph switching equivalent to $\overline{K}_4$ or K_4. Then the following sequence of moves shows that the graph is $C_5 \sqcup K_1$.

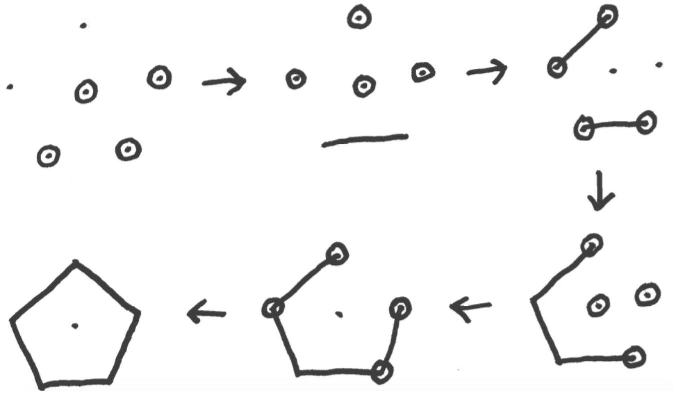

■

Corollary 5.1.0.1. Let Y be a graph on $n \geq 6$ vertices with $d(Y) \leq 5$ and $\alpha(Y) \neq 1/\sqrt{5}$. Then Y has an induced subgraph switching equivalent to $\overline{K}_4$ or K_4.

Example 5.1.3. In $\mathbb{R}^4$, a different configuration of four equiangular lines is available, which of course has rank 4. Tremain gave the following representation, in which the non-uniform matroid structure is visually apparent [Tre08]. This is the smallest example of a non-uniform matroid of real equiangular vectors.

$$L = \begin{pmatrix} & & & & \sqrt{2/3} & -\sqrt{2/3} \\ & & \sqrt{2/3} & -\sqrt{2/3} & & \\ \sqrt{2/3} & -\sqrt{2/3} & & & & \\ \sqrt{1/3} & \sqrt{1/3} & \sqrt{1/3} & \sqrt{1/3} & \sqrt{1/3} & \sqrt{1/3} \end{pmatrix}$$

The angle is $1/3$, and any four of the lines that span a three dimensional subspace are of course in the tetrahedral configuration.

Example 5.1.4. The defining representation of S_8 on $\mathbb{R}^8$ has two invariant subspaces: the trivial representation, which is spanned by the all 1s vector, plus the standard representation, which is 7 dimensional. Thus the 28 vectors obtained by permuting entries of the vector $\frac{1}{\sqrt{14}}(-3, -3, 1, 1, 1, 1, 1, 1)$ lie in a 7 dimensional subspace. They are equiangular with angle $1/3$. Thus, the absolute bound is attained in $\mathbb{R}^7$. This configuration also comes from the weights associated to the 56 dimensional representation of E_7 [citation needed].

Example 5.1.5. The Peterson graph

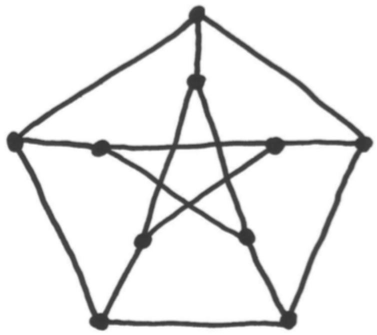

The eigenvalues of its Seidel matrix are ±3 each with multiplicity 5. Thus we have an

ETF of 10 vectors in dimension 5 with angle 1/3. This meets the relative bound, so the Peterson graph is switching equivalent to an SRG $k = 2c$ graph plus an isolated vertex. This SRG is the Paley graph P_9, which we define below.

Example 5.1.6. Consider the $(2d-2)\times(2d-2)$ Seidel matrix X with the following blocks on the diagonal and 1s elsewhere.

$$\begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}$$

Let J be the all 1s matrix and let I be the identity matrix. Then $A = X - J + I$ has characteristic polynomial $(\lambda+2)^{d-1}(\lambda-2)^{d-1}$. Some eigenvectors of A for -2 and 2 are

$$v_{-2}^1 = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix},\ v_{-2}^2 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix},\cdots \qquad v_{2}^1 = \begin{pmatrix} 1 \\ -1 \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix},\ v_{2}^2 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ -1 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix},\cdots$$

The vectors v_2^i are eigenvectors of J each with eigenvalue 0, so they are eigenvectors of X with eigenvalue 1. The all 1s vector $j = \sum_i v_{-2}^i$ is an eigenvector of J with eigenvalue $2(d-1)$, so it is an eigenvector of X with eigenvalue $2d-5$. There are $d-2 = (2d-2)-d$ remaining eigenvectors to account for. The differences $v_{-2}^i - v_{-2}^j$ for $i < j$ are eigenvectors of A with eigenvalue -2 and of J with eigenvalue 0, so they are the $d-2$ remaining eigenvectors of X and have eigenvalue -3. This is the least eigenvalue of X, so we see that X is the Seidel matrix of $2d-2$ equiangular vectors in $\mathbb{R}^d$ with angle 1/3.

Example 5.1.7. A *conference graph*, c.f. [GR01], is an SRG(n, k, a, c) with equal multiplicities $m_\theta = m_\tau$ of the two eigenvalues θ and τ other than the eigenvalue k. It follows that $k = (n-1)/2$, $a = (n-5)/4$, and $c = (n-1)/4$ [GR01]. Since the eigenvalues of a strongly regular graph are a function of its parameters, these two conditions are equivalent. It turns out that $n \equiv 1 \bmod 4$ and n is a sum of two squares. A *conference matrix* is the Seidel matrix of a conference graph.

Example 5.1.8. A *C-matrix* is an $n \times n$ Seidel matrix X such that $X^2 = cI$ for some $c \in \mathbb{R}$. It follows that $c = n - 1$. A *C-graph* is the graph corresponding to a C-matrix.

Consider an ETF of $n = 2d$ vectors in $\mathbb{R}^d$ with angle α. Let X be its Seidel matrix and let G be its Gram matrix. Then

$$X^2 = \frac{1}{\alpha^2}(G^2 - 2G + I) = \frac{1}{\alpha^2}\left(\frac{n}{d}G - 2G + I\right) = \frac{1}{\alpha^2}I$$

so X is a C-matrix and $\alpha = 1/\sqrt{n-1}$. Conversely, a C-matrix is a Seidel matrix with exactly two eigenvalues, so is an ETF. The relative bound is an equality, and simplifies to

$$0 = (n - 2d)(1 - dn)$$

Since n and d are integers greater than 1, $n = 2d$. $C_5 \sqcup K_1$ and the Peterson graph are examples of C-graphs.

Example 5.1.9 (Example (Paley graphs))**.** Let $\mathbb{F}_q$ be the field of order $q = p^r$, with p prime. Assume $q \equiv 1 \bmod 4$ so that -1 is a square, and so $x - y$ is a square if and only if $y - x$ is a square. The *Paley graph* of order q, P_q, is the graph with vertex set $\mathbb{F}_q$ with an edge from x to y if and only if $x - y$ is a square. Exactly half of the elements of $\mathbb{F}_q$ are squares, so P_q is k-regular with $k = (q-1)/2$. Moreover, it is shown in [Els12, Pei01] that Paley graphs are symmetric and self-complementary, and that such graphs are automatically strongly regular. See [Pei01] for a classification of the symmetric self-complementary graphs. Alternately, strong regularity of Paley graphs can be proved directly using Legendre symbols. This is shown in [zotb] when q is a prime, but their exact proof generalizes to prime powers. P_q is strongly regular with parameters

$$\left(q, \frac{q-1}{2}, \frac{q-5}{4}, \frac{q-1}{4}\right)$$

Thus all Paley graphs are conference graphs. $P_9 \sqcup K_1$ is switching equivalent to the Peterson graph. We now summarize the proof that $a = \frac{q-5}{4}$ from [zotb] adapted for any prime power $q \equiv 1 \bmod 4$. The proof that $c = \frac{q-1}{4}$ is similar.

Proof. For $x \in \mathbb{F}_q$, define the *Legendre symbol*

$$\left(\frac{x}{\mathbb{F}_q}\right) = \begin{cases} 0 & \text{if } x = 0, \text{ and otherwise,} \\ 1 & \text{if } x \text{ is a quadratic residue} \\ -1 & \text{if } x \text{ is not a quadratic residue} \end{cases}$$

This is *not* the Jacobi symbol. The property needed for our purpose is that the Legendre symbol is multiplicative (in the numerator). Let $x, y \in \mathbb{F}_q$ be distinct elements. Set

$$\begin{aligned} X &= \{x' \,|\, x - x' \text{ is a quadratic residue}\} \\ Y &= \{y' \,|\, y - y' \text{ is a quadratic residue}\} \\ Z &= (X \cap Y) \cup (X^c \cap Y^c) \end{aligned}$$

Then for any $z \in \mathbb{F}_q \setminus \{x, y\}$, $z \in Z$ if and only if the following conditions are equivalent:

$$\left(\frac{x - z}{\mathbb{F}_q}\right) = 1 \qquad\qquad \left(\frac{y - z}{\mathbb{F}_q}\right) = 1$$

or equivalently,

$$\left(\frac{(x - z)(y - z)}{\mathbb{F}_q}\right) = 1$$

It follows that

$$|Z| = \frac{1}{2} \sum_{z \in \mathbb{F}_q \setminus \{x,y\}} \left(1 + \left(\frac{(x - z)(y - z)}{\mathbb{F}_q}\right)\right)$$

This can be simplified (by doing a variable substitution) to show that $|Z| = \frac{q-3}{2}$. Then by set theory, $|Z| = 2|X \cap Y| + 1$, and so $|X \cap Y| = \frac{q-5}{4}$. ∎

It is required to carry out the similar proof for c; the SRG relation cannot be invoked to determine c since *a priori* we don't know that the graph is an SRG.

Recall that the eigenvalues of a Seidel matrix with $n > 2d$ are odd integers. For this reason, C-matrices, (which have the eigenvalue $1/\alpha = \sqrt{n-1}$) behave differently than graphs obtained from SRGs by adding an isolated vertex with $n > 2d$. It should also be noted that there exist conference graphs that are not Paley graphs. For instance, Waldron claims that such a graph exists that has 46 vertices [Wal09].

5.2 Complex equiangular configurations

We have not discussed the theory surrounding complex equiangular lines or SIC-POVMs in this thesis, but SICs in dimensions 2 and 3 can be constructed with very little machinery.

Example 5.2.1. There exists a unique unordered configuration of four equiangular lines in $\mathbb{C}^2$. A sequence of equiangular lines in $\mathbb{C}^2$ is represented by a regular tetrahedron inscribed in S^2 with vertices labeled 1,2,3,4. The group of unitaries preserving the set of lines maps to the alternating group on the four corresponding points on the Bloch sphere, A_4, (and not the full symmetric group S_4, since SU(2) maps onto SO(3) and not O(3)). Thus there are two ordered configurations of equiangular lines in $\mathbb{C}^2$. It turns out that there is an infinite (3 parameter) family of ordered configurations of equiangular vectors in $\mathbb{C}^2$. This phenomenon cannot happen in the real case, since there are finitely many Seidel matrices of a given size. The vectors form an ETF, and the eigenvalues of the Seidel matrix X are $\pm\sqrt{3}$, so the minimal polynomial is $x^2 - 3$. X has the form

$$X = \begin{pmatrix} 0 & a & b & c \\ \overline{a} & 0 & d & e \\ \overline{b} & \overline{d} & 0 & f \\ \overline{c} & \overline{e} & \overline{f} & 0 \end{pmatrix}$$

The relations imposed by the off-diagonal equations of $X^2 = 3I$ are equivalent to

$$\begin{aligned} d &= \pm i\overline{a}b \\ e &= \mp i\overline{a}c \\ f &= \pm i\overline{b}c \end{aligned}$$

The only other restrictions on the complex variables a, b, c are that $|a| = |b| = |c| = 1$. Thus ordered configurations of four equiangular vectors in $\mathbb{C}^2$ correspond to quadruples (a, b, c, ϵ) where $a, b, c \in S^1$ and $\epsilon = \pm 1$. ϵ determines the *orientation* of the configuration. By switching off one of the vectors, we obtain one of two possible 3×3 signature matrices for the configuration of lines.

Example 5.2.2. There is a 1-parameter family of ordered configurations of 9 equiangular lines in $\mathbb{C}^3$. This is thought to be a phenomenon unique to dimension 3. It is easy to verify that for any $\theta \in \mathbb{R}$, the vector

$$\begin{pmatrix} 1 \\ 0 \\ e^{i\theta} \end{pmatrix}$$

is *fiducial*, that is, the lines spanned by the vectors in its Heisenberg orbit form a SIC. Recall that the angle for 9 equiangular lines in $\mathbb{C}^3$ must be $1/(3+1)$. Thus, the distance between any two of these elements in $\mathbb{CP}^2$ is $\sqrt{2-2/4} = \sqrt{3/2}$. A regular triangle with side length $\sqrt{3/2}$ can be inscribed in a great circle of the (unnormalized) sphere $\mathbb{CP}^1$, so it is conceivable that a configuration of 9 equiangular lines exists in which three of the lines occupy a two dimensional subspace. Indeed, this is achieved by $\theta = \pi$, in which case the lines are said to form the *Hesse configuration*. This is the smallest example of complex equiangular lines with a non-uniform matroid structure, pictured below. The lines (each containing 3 points) represent 2 dimensional subspaces.

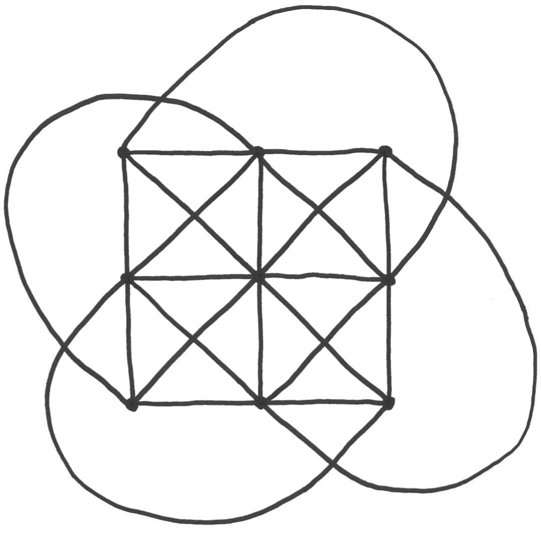

The Hesse configuration also arises in two ways from the *Hesse pencil*, the following family of projective plane curves

$$C_\mu : X^2 + Y^2 + Z^2 + \mu XYZ = 0$$

with $\mu \in \mathbb{C} \cup \{\infty\}$. Let us restrict our attention to the case that C_μ is nonsingular, or equivalently, $-3\mu \notin \{\infty, 1, \zeta_3, \zeta_3^2\}$. As for any projective plane curve, the inflection points of C_μ are its intersection with its Hessian H_μ, defined as the vanishing of the polynomial

$$\det\begin{pmatrix} \frac{\partial^2 f}{\partial X^2} & \frac{\partial^2 f}{\partial X \partial Y} & \frac{\partial^2 f}{\partial X \partial Z} \\ \frac{\partial^2 f}{\partial Y \partial X} & \frac{\partial^2 f}{\partial Y^2} & \frac{\partial^2 f}{\partial Y \partial Z} \\ \frac{\partial^2 f}{\partial Z \partial X} & \frac{\partial^2 f}{\partial Z \partial Y} & \frac{\partial^2 f}{\partial Z^2} \end{pmatrix}$$
$$= \big(XZ\mu^2 - 6\,Y^2\mu\big)Y\mu + \big(XY\mu^2 - 6\,Z^2\mu\big)Z\mu - 6\,\big(X^2\mu^2 - 36\,YZ\big)X$$

where f is the defining polynomial of C_μ. Since the defining polynomials of C_μ and H_μ each have degree 3, Bezout's theorem states that C_μ has $3 \times 3 = 9$ inflection points. Solving the simultaneous equations, (which is easily done in any symbolic algebra package,) yields the nine lines spanned by the Heisenberg orbit of the fiducial $(1, 0, e^{i\pi})$, the Hesse configuration.

The curves C_μ are smooth and have genus 1, and therefore can be made into elliptic curves by specifying an origin [Sil09]. Another way the Hesse configuration arises from the Hesse pencil is as its 3-torsion. Maybe there is a higher dimensional generalization where SICs are obtained from torsion subgroups of some abelian variety generalizing the Hesse cubic.

5.3 Other codes and applications

Example 5.3.1. The Pauli matrices

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \qquad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

are complementary observables, and so their respective eigenbases $\{|+\rangle, |-\rangle\}$,$\{|i\rangle, |-i\rangle\}$, and $\{|0\rangle, |1\rangle\}$ are mutually unbiased. The lines spanned by the vectors in these bases correspond to the points $(\pm1, 0, 0), (0, \pm1, 0), (0, 0, \pm1)$, respectively, on the Bloch sphere. (These are the *qubit stabilizer states*). We briefly describe two applications.

Two of these bases, typically chosen to be $B^1 = \{|0\rangle, |1\rangle\}$ and $B^2 = \{|+\rangle, |-\rangle\}$, are used in the BB84 quantum key distribution protocol. Alice wants to send a private key to Bob. Alice chooses a uniformly random bit string $x = x_1 \cdots x_n$, and for each bit x_i, she chooses a uniformly random basis $B_i \in \{B^1, B^2\}$. Then, for each i, she sends the following qubit $|\psi_i\rangle$ over a public quantum channel (or perhaps teleports it across a public classical channel):

x_i	B_i	$\lvert\psi_i\rangle$
0	B^1	$\lvert 0\rangle$
0	B^2	$\lvert +\rangle$
1	B^1	$\lvert 1\rangle$
1	B^2	$\lvert -\rangle$

Then, for each i, Bob chooses a uniformly random basis $C_i \in \{B^1, B^2\}$ and measures $|\psi_i\rangle$ relative to C_i. Bob then informs Alice, over a public channel, of his sequence of bases, and Alice reports back which ones he got right. Alice and Bob then compute their shared *sifted key*, the subsequence of bits x_i for which $B_i = C_i$. If Eve intercepts some of the $|\psi_i\rangle$ and measures them, then she will cause errors in the sifted keys obtained by Alice and Bob. This leads to a security proof based on the no-cloning theorem. Cryptographic protocols generalizing BB84 have also been devised that exploit a set of $d+1$ MUBs in $\mathbb{C}^d$ [CBKG02].

Next we describe a widely used method of quantum state tomography using the Pauli matrices. Since $B = \{I, \sigma_x, \sigma_y, \sigma_z\}$ is a basis for $\mathbb{C}^{2\times 2}$, $B^{\otimes n}$ is a basis for $\mathbb{C}^{2^n \times 2^n}$. Thus, for an n qubit system, the exact values of the observables $A \in B^{\otimes n} \smallsetminus \{I\}$ allow for complete state reconstruction. That is, an n qubit state must have the form

$$\rho = \frac{1}{2^n}\left(I + \sum_{A \in B^{\otimes n} \smallsetminus \{I\}} a_A A\right)$$

The values of these observables can be approximated by performing many measurements of each observable on identically-prepared states ρ. Note however, that the observables $B^{\otimes n} \smallsetminus \{I\}$ are not complementary for $n > 1$, since this set has cardinality $4^n - 1$, which exceeds the absolute bound for MUBs of $2^n + 1$.

Also note that the number of observables grows exponentially with the size of the system. This is unavoidable in any form of complete state determination since the dimension of the Hilbert space grows exponentially with the number of qubits.

Example 5.3.2. We review the construction for MUBs in odd prime power dimensions introduced by [WF89]. Let $q = p^r$ for some odd prime p and $r \geq 1$ and let ω be a nontrivial

pth root of unity. Each $y \in \mathbb{F}_q$ defines an $\mathbb{F}_p$-linear map $\mathbb{F}_q \to \mathbb{F}_q$ by left multiplication, and its trace is denoted $\operatorname{tr} y$. Set

$$|\psi_{a,b}\rangle = \frac{1}{\sqrt{q}} \sum_{x \in \mathbb{F}_q} \omega^{\operatorname{tr}\left(ax^2+bx\right)} |x\rangle$$

Then the sets $B_a = \{|\psi_{a,b}\rangle \,|\, b \in \mathbb{F}_q\}$ for $a \in \mathbb{F}_q$ are MUBs. To see this, we must show that

$$|\langle\psi_{a,b}|\psi_{a',b'}\rangle| = \begin{cases} 1 & a = a' \text{ and } b = b' \\ 0 & a = a' \text{ and } b \neq b' \\ \text{constant } c & a \neq a' \end{cases}$$

It follows that $c = 1/\sqrt{q}$. These q bases along with the standard basis are MUBs for $\mathbb{C}^q$.

Proof. In the $a = a'$ case we get

$$\langle\psi_{a,b}|\psi_{a,b'}\rangle = \frac{1}{q} \sum_{x \in \mathbb{F}_q} \omega^{\operatorname{tr}((b-b')x)} = \begin{cases} 1 & b = b' \\ \frac{1}{q} \sum_{x \in \mathbb{F}_q} \omega^{\operatorname{tr} x} & b \neq b' \end{cases}$$

Since the trace is a group homomorphism $\mathbb{F}_q \to \mathbb{F}_p$, its preimages (additive cosets of its kernel) have equal cardinality. Thus, in the $a = a'$, $b \neq b'$ case above, the sum is equal to 0. Now let $\alpha = a - a' \neq 0$ and $\beta = b - b'$. Note that for any $c \in \mathbb{F}_q$, the map $x \mapsto x + c$ is a bijection of $\mathbb{F}_q \to \mathbb{F}_q$, and for any (nonzero) quadratic residue α, the map $x \mapsto ax$ is a bijection of $(\mathbb{F}_q^\times)^2 \to (\mathbb{F}_q^\times)^2$. Then

$$\begin{aligned}
|\left\langle \psi_{a,b} \middle| \psi_{a',b'} \right\rangle| &= \frac{1}{q}\left|\sum_{x\in\mathbb{F}_q} \omega^{\operatorname{tr}\left(\alpha x^2+\beta x\right)}\right| \\
&= \frac{1}{q}\left|\sum_{x\in\mathbb{F}_q} \omega^{\operatorname{tr}\left(\alpha\left(\left(x+\frac{\beta}{2\alpha}\right)^2-\frac{\beta^2}{4\alpha^2}\right)\right)}\right| \\
&= \frac{1}{q}\left|\omega^{-\operatorname{tr}\left(\frac{\beta^2}{4\alpha^2}\right)}\sum_{x\in\mathbb{F}_q} \omega^{\operatorname{tr}\left(\alpha x^2\right)}\right| \\
&= \frac{1}{q}\left|\sum_{x\in\mathbb{F}_q} \omega^{\operatorname{tr}\left(\alpha x^2\right)}\right|
\end{aligned}$$

Note that completing the square required $\alpha \neq 0$ and $p \neq 2$. Now,

$$|\left\langle \psi_{a,b} \middle| \psi_{a',b'} \right\rangle| = \begin{cases} \frac{1}{q}\left|\sum_{x\in\mathbb{F}_q}\omega^{\operatorname{tr} x^2}\right| & \alpha \in (\mathbb{F}_q^\times)^2 \\ \frac{1}{q}\left|\sum_{x\in\mathbb{F}_q}\omega^{\operatorname{tr} x} - \sum_{x\in\mathbb{F}_q}\omega^{\operatorname{tr} x^2}\right| = \frac{1}{q}\left|\sum_{x\in\mathbb{F}_q}\omega^{\operatorname{tr} x^2}\right| & \alpha \notin (\mathbb{F}_q^\times)^2 \end{cases}$$

This expression is constant with respect to a, a', b, and b' (requiring only that $a \neq a'$). ■

As noted, this implies the expression equals $1/\sqrt{q}$. Alternately, we obtain a Gauss sum,

$$\begin{aligned}
\sum_{x\in\mathbb{F}_q}\omega^{\operatorname{tr} x^2} &= 1 + 2\sum_{x\in(\mathbb{F}_q^\times)^2}\omega^{\operatorname{tr} x} \\
&= 1 + \sum_{x\in\mathbb{F}_q^\times}\omega^{\operatorname{tr} x}\left(1+\left(\frac{x}{\mathbb{F}_q}\right)\right) \\
&= \sum_{x\in\mathbb{F}_q}\omega^{\operatorname{tr} x} + \sum_{x\in\mathbb{F}_q}\omega^{\operatorname{tr} x}\left(\frac{x}{\mathbb{F}_q}\right) \\
&= \sum_{x\in\mathbb{F}_q}\omega^{\operatorname{tr} x}\left(\frac{x}{\mathbb{F}_q}\right)
\end{aligned}$$

which has absolute value $\sqrt{q}$. When q is prime, this leads to a proof of quadratic reciprocity

[zota]. It would be interesting to know if this construction of MUBs generalizes such that the absolute value of other Gauss sums is derived, or that other reciprocity laws could be proved. For example, the MUB property and the Gauss sum still work properly when the trace function is replaced with any group homomorphism $\tau : \mathbb{F}_q \to \mathbb{F}_p$, and thus also works where the map $x \mapsto \omega^{\operatorname{tr} x}$ is replaced with any group homomorphism $f : \mathbb{F}_q \to S^1$, $(f = \omega^\tau)$.

The same set of MUBs can also be described as follows. Consider the generalized Paulis

$$\begin{aligned} X_s &= \sum_{x \in \mathbb{F}_q} |x+s\rangle \langle x| && \text{for } s \in \mathbb{F}_q \\ Z_s &= \sum_{x \in \mathbb{F}_q} \omega^{\operatorname{tr}(sx)} |x\rangle \langle x| && \text{for } s \in \mathbb{F}_q \end{aligned}$$

Then for each $a \in \mathbb{F}_q$ we have a distinct set $S_a = \{X_s Z_{2as} \,|\, s \in \mathbb{F}_q\}$ of q commuting $q \times q$ matrices. B_a (defined above) is the unique mutual eigenbasis of unit vectors for S_a.

References

[AAR99] George E. Andrews, Richard Askey, and Ranjan Roy. *Special Functions*. Cambridge University Press, 1999.

[AFF11] D. M. Appleby, Steven T. Flammia, and Christopher A. Fuchs. The Lie Algebraic Significance of Symmetric Informationally Complete Measurements. *Journal of Mathematical Physics*, 52(2):022202, February 2011.

[AFMY17] Marcus Appleby, Steven Flammia, Gary McConnell, and Jon Yard. SICs and Algebraic Number Theory. *Foundations of Physics*, 47(8):1042–1059, August 2017.

[BDKS16] Igor Balla, Felix Dräxler, Peter Keevash, and Benny Sudakov. Equiangular Lines and Spherical Codes in Euclidean Space. *arXiv:1606.06620 [math]*, June 2016.

[Buk15] Boris Bukh. Bounds on equiangular lines and on related spherical codes. *arXiv:1508.00136 [math]*, August 2015.

[CBKG02] Nicolas J. Cerf, Mohamed Bourennane, Anders Karlsson, and Nicolas Gisin. Security of quantum key distribution using d-level systems. *Physical Review Letters*, 88(12):127902, March 2002.

[CKM16] Henry Cohn, Abhinav Kumar, and Gregory Minton. Optimal simplices and codes in projective spaces. *Geometry & Topology*, 20(3):1289–1357, 2016.

[CS98] John Conway and Neil J. A. Sloane. *Sphere Packings, Lattices and Groups*. Springer Science & Business Media, December 1998.

[CW16] Tuan-Yow Chien and Shayne Waldron. A Characterization of Projective Unitary Equivalence of Finite Frames and Applications. *SIAM J. Discrete Math.*, 30:976–994, 2016.

[DGS77] P. Delsarte, J. M. Goethals, and J. J. Seidel. Spherical codes and designs. *Geometriae Dedicata*, 6(3):363–388, September 1977.

[Els12] Ahmed Noubi Elsawy. Paley Graphs and Their Generalizations. *arXiv:1203.1818 [math]*, March 2012.

[Gal] Jean Gallier. Notes on Spherical Harmonics and Linear Representations of Lie Groups. page 56.

[GH92] C. D Godsil and A. D Hensel. Distance regular covers of the complete graph. *Journal of Combinatorial Theory, Series B*, 56(2):205–238, November 1992.

[GKMS16] G. Greaves, J. H. Koolen, A. Munemasa, and F. Szöllősi. Equiangular lines in Euclidean spaces. *Journal of Combinatorial Theory, Series A*, 138:208–235, February 2016.

[God93] Chris Godsil. *Algebraic Combinatorics*. CRC Press, April 1993.

[GR01] Chris Godsil and Gordon F. Royle. *Algebraic Graph Theory*. Graduate Texts in Mathematics. Springer-Verlag, New York, 2001.

[Haa48] Johannes Haantjes. Equilateral point-sets in elliptic two- and three-dimensional spaces. *Niew Arch. Wisk.*, 1948.

[Har59] Jean Harmon. The Psychology of Interpersonal Relations. By Fritz Heider. New York: John Wiley and Sons, Inc., 1958. 322 pp. $6.25. *Social Forces*, 37(3):272–273, March 1959.

[Ivo81] I. D. Ivonovic. Geometrical description of quantal state determination. *Journal of Physics A: Mathematical and General*, 14(12):3241–3245, December 1981.

[JTY$^+$19] Zilin Jiang, Jonathan Tidor, Yuan Yao, Shengtong Zhang, and Yufei Zhao. Equiangular lines with a fixed angle. *arXiv:1907.12466 [math]*, July 2019.

[Kon13] Takis Konstantopoulos. A multilinear algebra proof of the Cauchy–Binet formula and a multilinear version of Parseval's identity. *Linear Algebra and its Applications*, 439(9):2651–2658, November 2013.

[KR05] A. Klappenecker and M. Rotteler. Mutually unbiased bases are complex projective 2-designs. In *Proceedings. International Symposium on Information Theory, 2005. ISIT 2005.*, pages 1740–1744, September 2005.

[LR81] Nathan Linial and Bruce L. Rothschild. *Incidence Matrices of Subsets – a Rank Formula*. 1981.

[LS73] P. W. H Lemmens and J. J Seidel. Equiangular lines. *Journal of Algebra*, 24(3):494–512, March 1973.

[Mac98] I. G. MacDonald. Symmetric functions and Hall polynomials. 1998.

[MS75] C. L. Mallows and N. J. A. Sloane. Two-Graphs, Switching Classes and Euler Graphs are Equal in Number. *SIAM Journal on Applied Mathematics*, 28(4):876–880, 1975.

[Neu89] A. Neumaier. Graph representations, two-distance sets, and equiangular lines. *Linear Algebra and its Applications*, 114-115:141–156, March 1989.

[Pei01] Wojciech Peisert. All Self-Complementary Symmetric Graphs. *Journal of Algebra*, 240(1):209–229, June 2001.

[Rob69] Robert W. Robinson. Enumeration of Euler graphs. 1969.

[Sil09] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*. Graduate Texts in Mathematics. Springer-Verlag, New York, 2 edition, 2009.

[Tol44] Tullio Ceccherini-Silberstein;Fabio Scarabotti;Filippo Tolli. *Representation Theory of the Symmetric Groups: The Okounkov-Vershik Approach, Character Formulas, and Partition Algebras by Tullio Ceccherini-Silberstein*. Cambridge University Press, 1744.

[Tre08] Janet C. Tremain. Concrete Constructions of Real Equiangular Line Sets. *arXiv:0811.2779 [math]*, November 2008.

[TZ12] Juan Alfredo Tirao and Ignacio Nahuel Zurrián. Spherical Functions: The Spheres Vs. The Projective Spaces. *arXiv:1207.0024 [math]*, June 2012.

[Wal09] Shayne Waldron. On the construction of equiangular frames from graphs. *Linear Algebra and its Applications*, 431(11):2228–2242, November 2009.

[Wal18] Shayne F. D. Waldron. *An Introduction to Finite Tight Frames*. Applied and Numerical Harmonic Analysis. Springer New York, New York, NY, 2018.

[WF89] William K Wootters and Brian D Fields. Optimal state-determination by mutually unbiased measurements. *Annals of Physics*, 191(2):363–381, May 1989.

[Wil90] Richard M. Wilson. A Diagonal Form for the Incidence Matrices of t-Subsets vs.k-Subsets. *European Journal of Combinatorics*, 11(6):609–615, November 1990.

[Zau11] Gerhard Zauner. QUANTUM DESIGNS: FOUNDATIONS OF A NONCOMMUTATIVE DESIGN THEORY. *International Journal of Quantum Information*, 09(01):445–507, February 2011.

[zota] A Proof of Quadratic Reciprocity Using Gauss Sums. https://wstein.org/edu/2007/spring/ent/ent-html/node54.html.

[zotb] Spectral Graph Theory, course announcement. http://www.cs.yale.edu/homes/spielman/561/2009/.

[zotc] Strongly regular graphs. https://www.win.tue.nl/˜aeb/graphs/srg/srgtab.html.

[zotd] Strongly Regular Graphs. http://www.maths.gla.ac.uk/˜es/srgraphs.php.