The Eigenvalue Method for Extremal Problems on Infinite Vertex-Transitive Graphs

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op

woensdag 27 mei 2015 om 12:30 uur

door

Philip Evan Bigras DECORTE

Master of Science (University of Toronto)
geboren te Ottawa, Ontario, Canada.

Dit proefschrift is goedgekeurd door de promotors:
Prof. dr. F. Vallentin en Prof. dr. ir. K.I. Aardal

Samenstelling promotiecommissie:

Rector Magnificus	
Prof. dr. ir. K.I. Aardal,	Technische Universiteit Delft, promotor
Prof. dr. F. Vallentin	Universität zu Köln, promotor

Onafhankelijke leden:

Prof. dr. C. Bachoc	Université de Bordeaux
Prof. dr. E. de Klerk	Universiteit van Tilburg
Prof. dr. M. Laurent	Universiteit van Tilburg
Prof. dr. J.M.A.M. van Neerven	Technische Universiteit Delft
Prof. dr. ir. G.J.T. Leus	Technische Universiteit Delft, reservelid

Andere leden:

Prof. dr. O. Pikhurko	University of Warwick

Het onderzoek beschreven in dit proefschrift is mede gefinancierd door de Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO), onder projectnummer 639.032.917.

Cover image created with Mathematica.

Printed and bound by Ipskamp Drukkers B.V., The Netherlands.

Contents

Summary

This thesis is about maximum independent set and chromatic number problems on certain kinds of infinite graphs. A typical example comes from the Witsenhausen problem: For $n \geq 2$, let $S^{n-1} := \{x \in \mathbb{R}^n : \|x\|_2 = 1\}$ be the unit sphere in $\mathbb{R}^n$, and let $G = (V, E)$ be the graph with $V = S^{n-1}$, in which two points in S^{n-1} are adjacent if and only if their inner product is equal to 0. What is the largest possible Lebesgue measure of an independent set in G?

The problem is reminiscent of a coding theory problem, in which one asks for the size of a largest set of distinct points in some metric space so that the distance between each pair of points is at least some specified constant d. Such a problem can be framed as a maximum independent set problem: Define a graph whose vertex set is the metric space, and join two points with an edge whenever their distance is less than d. The codes of minimum distance d are then precisely the independent sets in this graph.

In the Witsenhausen problem, rather than asking for a set of points in the sphere in which all the distances less than d are forbidden, we ask

for a set of points in which only one distance is forbidden. And it turns out that the *Delsarte* (also called *linear programming*) upper bounds for the size of codes [Del73] can be adjusted to give upper bounds for the measure of an independent set in the Witsenhausen graph. This was first done in [BNdOFV09] and [dOF09].

The Witsenhausen problem was stated in [Wit74], and in the same note it was shown that the fraction of the n-dimensional sphere which can be occupied by any measurable independent set is upper bounded by the function $1/n$. Frankl and Wilson [FW81] made a breakthrough in 1981 when they proved an upper bound which decreases exponentially in n. Despite this progress on asymptotics, the $1/3$ upper bound in the $n = 3$ case has not moved since the original statement of the problem until now. In Chapter 5 we give one of the main results of the thesis, which is an improvement of this upper bound to 0.313. The proof works by strengthening the Delsarte-type bounds using some combinatorial arguments deduced in Chapters 3 and 4.

The next main result of the thesis answers a natural question about the graphs $G(S^{n-1}, X)$, whose vertex set is S^{n-1} and where two points are joined with an edge if and only if their inner product belongs to the set $X \subset [-1, 1]$ of forbidden inner products. These graphs generalize the Witsenhausen graph, and are called *forbidden inner product graphs*. One may ask, Does there exist a measurable independent set of maximum measure? There is a graph $G = G(S^2, X)$ (many, in fact) having no such independent set. In Chapter 4 we construct for every $\varepsilon > 0$ an independent set in G having measure at least $1/2 - \varepsilon$, but we show that there is no independent set of measure equal to 1/2. In Chapters 6 and 7 we build on the theory of adjacency operators for infinite graphs developed in [BDdOFV14] to prove that maximum measurable independent sets exist in $G(S^{n-1}, X)$ for all $n \geq 3$, and for all sets X. As a relatively easy application of the machinery developed here, we also obtain a third result, which is that the supremum of the measures of independent sets in $G(S^{n-1}, X)$ depends only on the topological closure of X in $[-1, 1]$. In particular, every independent set has measure zero if 1 belongs to the closure of X.

Almost everything in this thesis relates to the Lovász ϑ-function of a graph, introduced in [Lov79]. The Delsarte bounds for binary codes can be regarded as coming from the ϑ-function, and Delsarte's bounds for spherical codes [DGS77] can be thought of as coming from an extension of the ϑ-function to forbidden inner product graphs on the unit sphere. Approaches inspired by the ϑ-function have been successful in improving lower bounds for the measurable chromatic number of Euclidean space (see for instance [BNdOFV09], [dOF09], [dOFV10], [BPT14]).

In Chapters 8 to 13 we develop two extensions of the ϑ-function to (possibly infinite) Cayley graphs over compact groups, which apply respectively to what we call *sparse* and *dense* graphs. Dense Cayley graphs have enough edges to guarantee that their independence numbers are finite, and in this case the applicable ϑ-function gives an upper bound for the cardinality of any independent set. Infinite sparse Cayley graphs have infinite independent sets, and the applicable ϑ-function then gives an upper bound for the Haar measure of any measurable independent set. The extensions we develop are based on the formulations of the ϑ-function for finite Cayley graphs given in [DdLV14]. We also show how many of the ϑ-function approaches taken in the literature can be seen as natural examples of our general framework.

The ϑ-function for finite graphs has formulations both as maximization and as minimization semidefinite programs which are mutually dual. In the approaches mentioned above in which the ϑ-function is extended to infinite graphs, it is also common to make use of duality, although in the infinite case it had not been shown that the primal and dual problems have equal values, a property known as *strong duality*. In this thesis we prove strong duality for our ϑ-functions using a different approach from the known strong duality proofs in the finite case. The definitions and proofs related to the ϑ-function build on a theory of positive type functions and measures which is developed in Chapters 8 to 10.

In [Mon11], Montina gives an application in quantum communication complexity of a natural conjecture about the Witsenhausen problem, the so-called *Double Caps Conjecture*. The extremal example for a spherical set in any dimension avoiding orthogonal pairs of points is conjectured [Kal09] to be the union of two opposite open spherical caps of geodesic radius $\pi/4$. In dimension 3, this configuration occupies about a 0.293-fraction of the unit sphere, so our new upper bound of 0.313 gets roughly halfway from the previous 1/3 upper bound to the Double Caps Conjecture. Assuming the Double Caps Conjecture, Montina is able to deduce a new lower bound on the cost of classically simulating a quantum channel.

Samenvatting

Dit proefschrift gaat over het onafhankelijkheidsgetal en het chromatisch getal van bepaalde soorten oneindige grafen. Een typisch voorbeeld wordt gegeven door het Witsenhausen probleem: Voor $n \geq 2$, zij $S^{n-1} := \{x \in \mathbb{R}^n : \|x\|_2 = 1\}$ de eenheidsbol in $\mathbb{R}^n$, en zij $G = (V, E)$ de graaf met knopenverzameling $V = S^{n-1}$, waarin twee knopen verbonden zijn met een kant als hun inwendig product gelijk is aan 0. Hoe groot kan de Lebesgue-maat zijn van een onafhankelijke verzameling in G?

Het probleem doet denken aan een probleem uit de coderingstheorie, waarin er wordt gevraagd naar het grootste aantal punten in een gegeven metrische ruimte waarvoor de onderlinge afstand tussen elk tweetal punten ten minste een gegeven constante d is. Zo'n probleem kan ook gesteld worden als het bepalen van een onafhankelijkheidsgetal: Definieer een graaf wiens knopenverzameling de metrische ruimte is, en verbind twee knopen met een kant als hun onderlinge afstand kleiner is dan d. De codes met minimum afstand d zijn dan precies de onafhankelijke verzamelingen in deze graaf.

Bij het Witsenhausen probleem vragen we niet naar een verzameling punten in de eenheidsbol waarin alle afstanden kleiner dan d verboden zijn, maar naar een verzameling waarin er precies *één* afstand wordt verboden. Het blijkt dat de *Delsarte* (ook wel *lineaire programmering* genoemd) bovengrenzen voor het aantal punten in een code aangepast kunnen worden aan de Witsenhausen graaf om nieuwe bovengrenzen te geven voor de maat van een onafhankelijke verzameling. Dit werd voor het eerst in [BNdOFV09] en [dOF09] gedaan.

Het Witsenhausen probleem werd in [Wit74] gesteld, en in hetzelfde artikel werd bewezen dat de fractie van de n-dimensionale eenheidsbol waarin een meetbare onafhankelijke verzameling kan zitten van boven begrensd is door de functie $1/n$. In een doorbraak [FW81] van Frankl en Wilson uit 1981 is een bovengrens die exponentiëel afneemt in n ontdekt. Ondanks deze vooruitgang wat betreft het asymptotische gedrag, is de bovengrens van $1/3$ voor $n = 3$ tot nu toe niet verbeterd. In Hoofdstuk 5 geven we een van de hoofdresultaten van het proefschrift, namelijk een verbetering van de bovengrens tot 0.313. Het bewijs kan gezien worden als een verscherping van de uit Delsarte volgende grenzen door deze met wat combinatorische redenering te combineren. Deze redenering wordt in Hoofdstukken 3 en 4 uitgelegd.

Het volgende hoofdresultaat van dit proefschrift geeft antwoord op een natuurlijke vraag over grafen $G(S^{n-1}, X)$ wiens knopenverzameling S^{n-1} is en waarin twee knopen verbonden zijn door een kant als hun onderlinge inwendig product in de verzameling $X \subset [-1, 1]$ van verboden inwendige producten ligt. Deze grafen generaliseren de Witsenhausen graaf en heten *verboden inwendig productgrafen.* Er kan gevraagd worden: Bestaat er een meetbare onafhankelijke verzameling met de grootst mogelijke maat? Er bestaat een graaf $G = G(S^2, X)$ (in feite vele) waarvoor geen zo'n onafhankelijke verzameling bestaat. In Hoofdstuk 4 construeren we voor elke $\varepsilon > 0$ een onafhankelijke verzameling in G met maat minstens $1/2 - \varepsilon$, maar we bewijzen ook dat er geen onafhankelijke verzameling bestaat wiens maat gelijk is aan $1/2$. In Hoofdstukken 6 en 7 gebruiken we de theorie van verbindingsoperatoren van oneindige grafen opgebouwd in [BDdOFV14] om te bewijzen dat meetbare onafhankelijke verzamelingen van zo groot mogelijk maat in $G(S^{n-1}, X)$ feitelijk bestaan voor alle $n \geq 3$ en alle X. Als gevolg van het hier ontwikkelde gereedschap verkrijgen we bovendien een derde resultaat, dat het supremum van de maten van de onafhankelijke verzamelingen in $G(S^{n-1}, X)$ alleen afhangt van de topologische afsluiting van X in $[-1, 1]$. In het bijzonder heeft elke onafhankelijke verzameling maat nul als 1 in de afsluiting van X ligt.

Vrijwel alles in dit proefschrift heeft te maken de Lovász ϑ-functie van een graaf, geïntroduceerd in [Lov79]. De Delsarte grenzen voor binaire codes kunnen gezien worden als speciale gevallen van de ϑ-functie, en de Delsarte grenzen voor sferische codes zijn ontleend aan een uitbreiding van de ϑ-functie op verboden inwendig productgrafen op de eenheidsbol. Ideeën geïnspireerd door de ϑ-functie hebben veel succes gehad bij het verbetering van onder anderen ondergrenzen voor het chromatisch getal van de Euclidische ruimte (zie bijvoorbeeld [BNdOFV09], [dOF09], [dOFV10], [BPT14]).

In Hoofdstukken 8 tot en met 13 ontwikkelen we twee uitbreidingen van de ϑ-functie op (mogelijk oneindige) Cayley-grafen over compacte groepen, die toepassen op wat wij *dunne* en *dikke* grafen noemen. Dikke Cayley-grafen hebben genoeg kanten om te garanderen dat hun onafhankelijkheidsgetallen eindig zijn, en in dit geval geeft de bijpassende ϑ-functie een bovengrens op de cardinaliteit van elke onafhankelijke verzameling. Oneindige dunne Cayley-grafen hebben oneindige onafhankelijke verzamelingen, en de bijpassende ϑ-functie geeft dan een bovengrens op de Haar-maat van een meetbare onafhankelijke verzameling. De hier ontwikkelde uitbreidingen zijn gebaseerd op de formuleringen van de ϑ-functies voor eindige Cayley-grafen gegeven in [DdLV14]. Bovendien laten we zien hoe de meeste ϑ-functie aanpakken die voorkomen in de literatuur gezien kunnen worden als voorbeelden van onze algemene theorie.

De gewone ϑ-functie voor eindige grafen heeft zowel een formulering als een maximaliserings- als een minimaliseringsprobleem in de semidefiniete programmering, die duaal zijn aan elkaar. Bij de bovengenoemde aanpakken waarin de ϑ-functie uitgebreid wordt naar oneindige grafen maakt men vaak gebruik van dualiteit, alhoewel *sterke dualiteit*, namelijk de eigenschap dat de waarden van beide formuleringen gelijk zijn aan elkaar, tot nu toe nooit bewezen was. In dit proefschrift bewijzen we sterke dualiteit voor onze ϑ-functies, en dit doen we op een andere manier dan gebruikelijk in het geval van eindige grafen. De definities en bewijzen die te maken hebben met de ϑ-functie worden overigens opgebouwd op basis van een theorie voor functies en maten van positief type, die wordt ontwikkeld in Hoofdstukken 8 tot en met 10.

In [Mon11] geeft Montina een toepassing in de quantum informatie theorie van een natuurlijk vermoeden dat verbonden is aan het Witsenhausen probleem, het zogenaamde *Twee Bolkappen Vermoeden*. Vermoed wordt [Kal09] dat in elke dimensie, een extremaal voorbeeld van een deelverzameling van de eenheidsbol die geen tweetal punten bevat die loodrecht op elkaar staan gegeven wordt door de vereniging van twee tegengestelde open bolkappen van geodetische straal $\pi/4$. In dimensie 3 bevat deze configuratie ongeveer 0.293 van de oppervlakte van de eenheidsbol, dus onze nieuwe bovengrens van 0.313 ligt nagenoeg middenin de eerder bekende bovengrens van 1/3 en de grens van het Twee Bolkappen Vermoeden. Uitgaande van het Twee Bolkappen Vermoeden kon Montina een nieuwe ondergrens bewijzen voor de kosten van het simuleren van een quantum communicatiekanaal met een klassiek kanaal.

Part I

Overview and preliminaries

Chapter 1

Main results and outline of thesis

The thesis is divided into five parts which we now explain.

Part I lists the main results of the thesis, and also briefly reviews some basic concepts used throughout the thesis. We fix notation and terminology from graph theory and linear algebra. We also introduce the theory of positive semidefinite matrices and semidefinite programming. We then review the Lovász ϑ-function of a graph, which is an important application of semidefinite programming in combinatorial optimization. Almost all of the results presented in this thesis are related in one way or another to the ϑ-function.

Part II is about the maximum measure of a spherical set avoiding orthogonal pairs of points. In the first chapter we calculate the maximum possible measure of a subset of the unit circle avoiding a single forbidden angle, and we determine when a maximizer exists. This solves a problem from [dOF09]. The remainder of Part II is devoted to strengthening the methods in [BNdOFV09] and [dOF09] to upper bound the spherical surface measure of a set I of unit vectors in $\mathbb{R}^3$ having the property that no two points in I are orthogonal. We find the new upper bound of 0.313 times the measure of the unit sphere. This improves the upper bound of 1/3 given in [Wit74], which has remained the best known upper bound for around 40 years. The best known lower bound is $1 - \frac{1}{\sqrt{2}} \approx 0.29$, given by two opposite caps of geodesic radius $\pi/4$; this is conjectured by Gil Kalai [Kal09] to be the optimal configuration.

In Part III, we return to the problem of when a maximizer exists, and we prove one of the main results of the thesis. We ask the following general question: Let $n \geq 2$ and $X \subset [-1, 1]$ be given, and let S^{n-1} be the unit sphere in $\mathbb{R}^n$. A subset $I \subset S^{n-1}$ is called X*-avoiding* if $\xi \cdot \eta \notin X$ for all $\xi, \eta \in I$. Let α denote the supremum of the Lebesgue surface measures of all X-avoiding subsets of S^{n-1}. For which n and X does there exist an X-avoiding subset of S^{n-1} having measure α? We prove that a maximizer exists whenever $n \geq 3$. The proof is a functional analytic compactness argument. Surprisingly, the argument fails when $n = 2$. In this case, the answer depends on X: a maximizer may or may not exist. Parts II and III are based on the article *Spherical sets avoiding a prescribed set of angles* [DP15], which is joint work between the author and Oleg Pikhurko of the University of Warwick. It is under review at the time of writing of this thesis.

Part IV is about positive type functions and measures. The main aim of this part is to provide the analytic foundation needed in Part V, but some of the results presented here may be interesting in their own right. After providing some background in harmonic analysis, we give a proof of Bochner's theorem for functions of positive type on compact groups. The new proof is simple and self-contained. We also state and prove a version of Bochner's theorem for measures of positive type on a compact group. The remainder of the part consists of a number of applications of both versions of Bochner's theorem. The two applications which are most relevant from the optimization point of view are the following: Under the dual pairing of continuous functions with regular Borel measures, (1) the cone of positive type measures is dual to the cone of positive type functions in the sense of conic duality; and (2) the cone of positive type functions is weak-* dense in the cone of positive type measures.

In Part V we present generalizations of the Lovász ϑ-function and Schrijver's ϑ'-function which apply to infinite Cayley graphs over compact groups. We distinguish between two sorts of graphs: sparse and dense. Roughly speaking, dense Cayley graphs on compact groups have enough edges to guarantee that the independence number is finite, while sparse graphs have so few edges that they have independent sets of positive Haar measure. Definitions of the ϑ- and ϑ'-function are given for both the dense and sparse case. The main contribution is a duality theory, which includes proofs of strong duality. Additionally, we investigate which properties of the usual ϑ-function hold when the graph is infinite, and we then work through some examples, recovering several seemingly disjoint results in the literature from a common framework; cf. [BNdOFV09], [BDdOFV14], [DdLV14], [DGS77], [dOF09], [dOFV10].

The idea of focussing on Cayley graphs began with the writing of the article *Fourier analysis on finite groups and the Lovasz theta-number of Cayley graphs* [DdLV14], which was joint work between the author, David de Laat, and the author's thesis advisor Frank Vallentin. Only finite graphs are dealt with in [DdLV14]. The main initial interest in Cayley graphs came from the fact that they provide a good setting in which to write down the frequency domain formulation of the ϑ-number, but restricting to Cayley graphs also allows for the application of harmonic analysis when proving theorems, which becomes particularly interesting when the graph is infinite.

At numerous places in the thesis, inspiration has been taken from the article *Spectral bounds for the independence ratio and the chromatic number of an operator* ([BDdOFV14]), which was joint work between the author, Christine Bachoc, Fernando Mario de Oliveira Filho, and Frank Vallentin. In particular, the idea of "adjacency operator" used in Parts II and III for forbidden inner product graphs on the unit sphere came from [BDdOFV14], and especially important in this thesis was a sufficient condition for the compactness of this operator, first given in [BDdOFV14], and applied in Part III of this thesis to obtain one of the main results.

While the primary contribution of Part V is intended to be the duality theory for our ϑ-functions, a secondary contribution is the further development of the line of thought started in [BDdOFV14]. In Part V the main ideas of [BDdOFV14] are streamlined for the most interesting applications by eliminating the language of operators and measurable graphs, and most of the results from [BDdOFV14] are recovered with easier proofs.

Chapter 2

General preliminaries

2.1 Graph theory

A *graph* is an ordered pair (V, E), where V is any set, called the *vertex set*, and E is a collection of subsets of V of cardinality 2. The set E is called the *edge set*.

If $G = (V, E)$ is a graph, an *independent set* I in G is a subset of V such that $\{x, y\} \notin E$ for any $x, y \in I$. A *clique* Q in G is a subset of V such that $\{x, y\} \in E$ for every $x, y \in Q$. The *independence number* $\alpha(G)$ of G is defined as the cardinality of a largest independent in G if this number is finite, and ∞ otherwise. A *colouring* of G is a partition of V into independent sets. The smallest number $\chi(G)$ of independent sets required is called the *chromatic number* of G. We write $\chi(G) = \infty$ when this number is infinite. The *compementary graph* G^c of G is the graph $G^c = (V, E')$, where $E' = \{\{u, v\} \subset V : \{u, v\} \notin E, u \neq v\}$.

An *automorphism* of a graph $G = (V, E)$ is bijection $a : V \to V$ satisfying

$$\{a(u), a(v)\} \in E \iff \{u, v\} \in E, \quad (u, v \in V).$$

The set of automorphisms of G forms a group under composition, which we denote by $\mathrm{Aut}(G)$. We say that G is *vertex-transitive* if $\mathrm{Aut}(G)$ acts transitively on V; that is, if for every $u, v \in V$, there exists $a \in \mathrm{Aut}(G)$ such that $a(u) = v$. We say that G is *edge-transitive* if $\mathrm{Aut}(G)$ acts transitively on E; that is, if for every $\{u_1, u_2\}, \{v_1, v_2\} \in E$, there exists $a \in \mathrm{Aut}(G)$ such that $\{v_1, v_2\} = \{a(u_1), a(u_2)\}$. The complement of the 7-cycle is an example of a vertex-transitive graph which is not edge-transitive, and a star is an example of an edge-transitive graph that is not vertex-transitive.

2.2 Linear algebra

Let $\mathcal{F}$ be either $\mathbb{R}$ or $\mathbb{C}$. The space of matrices with m rows and n columns with entries from $\mathcal{F}$ will be denoted $\mathcal{F}^{m\times n}$ or simply by $\mathcal{F}^m$ when $n = 1$.

The transpose of a matrix $A \in \mathcal{F}^{m\times n}$ is denoted by A^t, and its conjugate transpose is denoted by A^*. We will use $\overline{A}$ to denote the entrywise complex conjugation of A. We say that A is *symmetric* when $A^t = A$, and we say A is *Hermitian* when $A^* = A$. The trace of A is denoted $\mathrm{Tr}(A)$. The $n \times n$ identity matrix will be denoted $I_{n\times n}$.

For $\mathcal{F}$ equal to either $\mathbb{R}$ or $\mathbb{C}$, we think of $\mathcal{F}^n$ as a Hilbert space with the inner product $\langle u, v\rangle = v^*u$, and norm $\|v\| = \sqrt{\langle v, v\rangle}$.

For two $n \times n$ matrices $A = (a_{ij})_{i,j=1}^n$, $B = (b_{ij})_{i,j=1}^n$, we define their *trace inner product* by

$$\langle A, B \rangle := \mathrm{Tr}(B^* A) = \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} \overline{b_{ij}}.$$

The *Frobenius norm* of A is defined as

$$\|A\|_2 := \sqrt{\langle A, A \rangle}.$$

A matrix $A \in \mathbb{C}^{n \times n}$ is called *Hermitian positive semidefinite*, or simply *positive semidefinite*, if

$$v^* A v \geq 0 \quad \text{for all } v \in \mathbb{C}^n. \tag{2.1}$$

Using the polarization identity [Fol95, A1.1], one can show that (2.1) implies that A is Hermitian. Therefore a matrix is positive semidefinite if and only if it is Hermitian and all its eigenvalues are nonnegative. If A has entries from $\mathbb{R}$, then by the spectral theorem A is positive semidefinite if and only if it is symmetric and

$$v^t A v \geq 0 \quad \text{for all } v \in \mathbb{R}^n. \tag{2.2}$$

Note that (2.2) alone does not imply that A is symmetric; consider for instance $A = \left(\begin{smallmatrix} 1 & 1 \\ -1 & 1 \end{smallmatrix}\right)$.

There are many equivalent ways of defining positive semidefinite matrices, a few of which we summarize below. (Cf. [Lov03])

Proposition 2.1. *Let $A \in \mathbb{C}^{n \times n}$. Then the following are equivalent:*

1. *A is positive semidefinite; that is* $v^*Av \geq 0$ *for all* $v \in \mathbb{C}^n$*;*

2. *A is diagonalizable and all its eigenvalues are real and nonnegative;*

3. $A = B^*B$ *for some* $B \in \mathbb{C}^{n\times n}$*;*

4. *A is a nonnegative linear combination of matrices of the form* vv^**, with* $v \in \mathbb{C}^n$*;*

5. *The determinant of every principal submatrix of A is real and nonnegative.*

For a symmetric matrix $A \in \mathbb{R}^{n\times n}$*, the equivalence of 1-5 holds if* $\mathbb{C}$ *is replaced by* $\mathbb{R}$*.*

The set of positive semidefinite matrices in either $\mathbb{R}^{n\times n}$ or $\mathbb{C}^{n\times n}$ forms a cone (see Section 10.1), meaning that it is closed under addition, and under multiplication by nonnegative (real) scalars.

The following are some important facts about positive semidefinite matrices which follow easily from Proposition 2.1.

Proposition 2.2. 1. *If* $A, B \in \mathbb{C}^{n\times n}$ *are positive semidefinite matrices, then* $\langle A, B\rangle \geq 0$*;*

2. *If* $A \in \mathbb{C}^{n\times n}$*, then A is positive semidefinite if and only if* $\langle A, B\rangle \geq 0$ *for every positive semidefinite matrix* $B \in \mathbb{C}^{n\times n}$*;*

3. *If* $A \in \mathbb{R}^{n\times n}$ *is symmetric, then* A *is positive semidefinite if and only if* $\langle A, B\rangle \geq 0$ *for every positive semidefinite matrix* $B \in \mathbb{R}^{n\times n}$*.*

The usual matrix product of two positive semidefinite matrices need not be Hermitian, let alone positive semidefinite. However, it is a fact that if A and B are positive semidefinite, then AB is positive semidefinite if and only if it is Hermitian; this is because if $A = C^*C$ as in Proposition 2.1, then AB and CBC^* have the same nonnegative eigenvalues.

If $A = (a_{ij})_{i,j}, B = (b_{ij})_{i,j}$ are two $m \times n$ matrices, the *Hadamard product* or *entrywise product* of A and B is defined as the $m \times n$ matrix whose ij-entry is $a_{ij}b_{ij}$. We have the following nice fact about the Hadamard product of semidefinite matrices, which is known as the Schur product theorem; it is proven in [Sch11].

Proposition 2.3 (Schur product theorem)**.** *The Hadamard product of two positive semidefinite matrices is positive semidefinite.*

2.3 Semidefinite programming

For $A \in \mathbb{C}^{n\times n}$, we write $A \succeq 0$ to mean that A is Hermitian positive semidefinite. A semidefinite program is an optimization problem of the following form:

$$\text{minimize} \quad c^t x \tag{2.3}$$

$$\text{subject to} \quad x_1 A_1 + \cdots + x_n A_n - B \succeq 0 \tag{2.4}$$

$$x = (x_1, \ldots, x_n) \in \mathbb{R}^n \tag{2.5}$$

where $c \in \mathbb{R}^n$, $B, A_1, \ldots, A_n \in \mathbb{R}^{n \times n}$ are given symmetric matrices. Since a diagonal matrix is positive semidefinite if and only if all the entries on the main diagonal are nonnegative, it is easy to see that semidefinite programming generalizes linear programming.

Modulo some technicalities (which almost never present a problem in practice), a semidefinite program with rational coefficients can be solved to any fixed degree of precision in time growing no faster than a polynomial in the input size. This can be accomplished using the ellipsoid method, though in practice interior point methods are used because of the practical inefficiency of the ellipsoid method. For this reason, semidefinite programming has proven useful in developing approximation algorithms for hard combinatorial optimization problems. An excellent survey on this topic can be found in [Lov03].

Throughout this thesis, we assume the reader has some basic familiarity with linear programming. A good reference is the book by Matoušek and Gärtner [GM07].

2.4 The ϑ- and ϑ'-functions

The Lovász ϑ-function was introduced in [Lov79] as an upper bound for the so-called Shannon capacity of a graph. Using the ϑ-function he determined the Shannon capacity of the 5-cycle, settling a problem of Shannon that had remained open more than 20 years.

For a graph $G = (V, E)$ with $V = \{1, \dots, n\}$, the Lovász ϑ-function $\vartheta(G)$ of G is defined as the value of the following semidefinite program in the matrix variable $A = (a_{ij}) \in \mathbb{R}^{n \times n}$.

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} \\ \text{subject to} \quad & \operatorname{Tr}(A) = 1 \\ a_{ij} = 0 \text{ when } & \{i, j\} \in E \\ A & \succeq 0 \end{aligned} \tag{2.6}$$

The ϑ'-function $\vartheta'(G)$ of G, introduced by Schrijver in [Sch79], is the value of the program obtained from program (2.6) by adding the constraints $a_{ij} \geq 0$ for all $i, j = 1, \dots, n$. Clearly one has $\vartheta'(G) \leq \vartheta(G)$.

The ϑ-function is explored in detail in the original paper of Lovász [Lov79] and in the survey article [Knu94] by Knuth. The most important property for us is what has come to be known as the "sandwich theorem":

$$\alpha(G) \leq \vartheta'(G) \leq \vartheta(G) \leq \chi(G^c).$$

The formulations of ϑ and ϑ' just given show that they can be computed in polynomial time using semidefinite programming solvers, and they therefore provide polynomial time computable bounds for the two NP-hard graph parameters $\alpha(G)$ and $\chi(G)$. This is the main reason these functions are of interest in combinatorial optimization.

In this thesis, the main interest in ϑ and ϑ' is not their low computational complexity, but rather the fact that they can be regarded as *spectral* or *eigenvalue* bounds for α and χ. This thesis deals with infinite graphs for which the definitions of α and χ are extended in a reasonable way. Calculating α and χ exactly for the types of graphs discussed here does not seem possible with today's mathematical technology; a good example is the Hadwiger-Nelson problem (see [Soi09]), which asks for the chromatic number c of the graph over the vertex set $\mathbb{R}^2$ in which two points are joined by an edge precisely when their Euclidean distance is equal to 1. It has been known since 1950 that $4 \leq c \leq 7$, and since then this inequality has not been improved.

Exact values of α and χ have only been found in a few very special cases, for instance the kissing numbers in dimensions 2, 3, 4, 8, and 24 ([Mus08], [CS93], [Lev79], and [PZ04]) and single forbidden distance graphs on the circle (this is Theorem 4.1).

Therefore, rather than trying to compute α and χ exactly, one might try to extend known eigenvalue methods to these infinite graphs in order to obtain bounds. This idea actually goes back at least as far as 1977 when Delsarte, Goethels, and Seidel [DGS77] extended the eigenvalue upper bounds for binary codes from Delsarte's Ph.D. thesis [Del73] to spherical codes.

A nice introduction to the eigenvalue method in extremal combinatorics is given by Ellis in the lecture notes [Ell11].

2.5 Cones and cone programming

A *dual pair* is a pair of $\mathbb{R}$-vector spaces V, V', together with a bilinear mapping $\langle \cdot, \cdot \rangle : V \times V' \to \mathbb{R}$ satisfying

1. If $\langle v, v' \rangle = 0$ for all $v' \in V'$, then $v = 0$;
2. If $\langle v, v' \rangle = 0$ for all $v \in V$, then $v' = 0$.

The mapping $\langle \cdot, \cdot \rangle$ is called the *bilinearity* of the the pair (V, V').

Let V be an $\mathbb{R}$-vector space. A subset $K \subset V$ is called a *cone* if it satisfies the following properties:

1. $K + K \subset K$; and
2. $tK \subset K$ for all $t \geq 0$.

Here we use the notation $K + K = \{k + k' : k, k' \in K\}$ and $tK = \{tk : k \in K\}$. Cones are always convex. If a cone K also satisfies the property $K \cap (-K) = \{0\}$, then we say K is *pointed*. Each pointed cone K in V defines a partial order relation $\geq_K$ on V by $x \geq_K y \iff x - y \in K$.

Let $V, V', \langle \cdot, \cdot \rangle_V$, and $W, W', \langle \cdot, \cdot \rangle_W$ be two dual pairs of $\mathbb{R}$-vector spaces, and let $K \subset V$ and $L \subset W$ be cones. Let $b \in W, c \in V'$, and let

$A : V \to W$ be a linear operator. A *cone program* or *conic program* is an optimization problem of the following form

$$\begin{aligned} \text{minimize } & \langle x, c \rangle \\ \text{subject to } & b - Ax \geq_L 0 \\ & x \geq_K 0. \end{aligned} \tag{2.7}$$

Linear programs and semidefinite programs are both cone programs. Any $x \in K$ satisfying $b - Ax \geq_L 0$ is called a *feasible solution* to (2.7). If (2.7) has a feasible solution, we say it is *feasible*, and otherwise we say it is *infeasible*. The *objective value* of a feasible solution x is $\langle x, c \rangle$. The *value of the program* (2.7) is

$$\inf\{\langle x, c \rangle : b - Ax \in L, x \in K\}. \tag{2.8}$$

We say that a program is *bounded* when its value is finite. In order to save space, we typically rewrite programs of the form (2.7) like (2.8).

Part II

Upper bounds for measures of spherical sets avoiding orthogonal pairs of points

Chapter 3

Background and combinatorial upper bound

3.1 Background

H. S. Witsenhausen [Wit74] in 1974 presented the following problem: Let S^{n-1} be the unit sphere in $\mathbb{R}^n$ and suppose $I \subset S^{n-1}$ is a Lebesgue measurable set having the property that $\langle \xi, \eta \rangle \neq 0$ for all $\xi, \eta \in I$. What is the largest possible Lebesgue surface measure of I?

Let $\alpha(n)$ denote this maximum divided by the total surface measure of S^{n-1}. Also in [Wit74], Witsenhausen deduced that $\alpha(n) \leq 1/n$. In 1981 [FW81, Theorem 6] Frankl and Wilson proved the famous result named after them, and as an application they gave the first exponentially decreasing upper bound: $\alpha(n) \leq (1 + o(1))(1.13)^{-n}$ using a combinatorial argument. Later in 1999, Raigorodskii [Rai99] improved the bound to $(1 + o(1))(1.225)^{-n}$ using a refinement of the Frankl-Wilson method. In 2009, Gil Kalai conjectured in his weblog [Kal09] that $\alpha(n) = (\sqrt{2} + o(1))^{-n}$, achieved by two opposite caps, each of geodesic radius $\pi/4$.

Besides the existing interest in the double caps conjecture, it is also interesting because if true, it would imply new lower bounds for the measurable chromatic number of Euclidean space, which we now discuss. Let $c(n)$ be the smallest integer k such that $\mathbb{R}^n$ can be partitioned into sets $X_1, \ldots, X_k$, with $\|x - y\| \neq 1$ for each $x, y \in X_i$, $1 \leq i \leq k$. The number $c(n)$ is called the *chromatic number of* $\mathbb{R}^n$, since the sets $X_1, \ldots, X_k$ can be thought of as the colour classes for the graph on the vertex set $\mathbb{R}^n$, in which we join two points when they have distance 1.

A conjecture of Erdös states that $c(n)$ increases exponentially. Frankl and Wilson also prove this conjecture [FW81, Theorem 3] with a combinatorial argument, showing that $c(n) \geq (1 + o(1))(1.2)^n$. Raigorodskii [Rai00] improved the lower bound to $(1 + o(1))(1.239)^n$.

Requiring the classes $X_1, \ldots, X_k$ to be measurable yields the *measurable chromatic number* $c_m(n)$. Clearly $c_m(n) \geq c(n)$. It was proven recently in [BPT14] that $c_m(n) \geq (1.268+o(1))^n$. Assuming Kalai's double caps conjecture, it is not hard to prove that $c_m(n) \geq (\sqrt{2}+o(1))^n$.

3.2 Preliminaries

If $u, v \in \mathbb{R}^n$ are two vectors, their standard inner product will be denote $\langle u, v \rangle$. All vectors will be assume to be column vectors. The transpose of a matrix A will be denoted A^t. We denote by $SO(n)$ the group of $n \times n$ matrices A over $\mathbb{R}$ having determinant 1, for which $A^t A$ is equal to the identity matrix. We will think of $SO(n)$ as a compact topological group, and we will always assume its Haar measure is normalized so that $SO(n)$ has measure 1. We denote by S^{n-1} the set of unit vectors in $\mathbb{R}^n$:

$$S^{n-1} = \{x \in \mathbb{R}^n : \langle x, x \rangle = 1\}.$$

We equip S^{n-1} with its usual topology. The Lebesgue measure λ on S^{n-1} is always taken to be normalized so that $\lambda(S^{n-1}) = 1$. Where we need to refer to the standard surface measure of S^{n-1}, we use ω_n. The Lebesgue σ-algebra on S^{n-1} will be denoted $\mathcal{L}$. When $(X, \mathcal{M}, \mu)$ is a measure space and $1 \leq p < \infty$, we use

$$L^p(X) = \{f : f \text{ is an } \mathbb{R}\text{-valued } \mathcal{M}\text{-measurable function and } \int_X |f|^p \, d\mu < \infty\}.$$

For $f \in L^p(X)$, we define $\|f\|_p := \left(\int_X |f|^p \, d\mu\right)^{1/p}$. Identifying two functions when they agree μ-almost everywhere, $L^p(X)$ becomes a Banach space with the norm $\| \cdot \|_p$.

We will use bold letters (for example $\boldsymbol{X}$) for random variables. The expectation of a function f of a random variable $\boldsymbol{X}$ will be denoted $\mathbb{E}_{\boldsymbol{X}}[f(\boldsymbol{X})]$, or just $\mathbb{E}[f(\boldsymbol{X})]$. The probability of an event E will be denote $\mathbb{P}[E]$.

When X is a set, we use $\mathbb{1}_X$ to denote its characteristc function; that is $\mathbb{1}_X(x) = 1$ if $x \in X$ and $\mathbb{1}_X(x) = 0$ otherwise. When X is a subset of some topological space, $\overline{X}$ will denote its closure.

For $X \subset [-1,1]$, we define $G(S^{n-1}, X)$ to be the graph with vertex set S^{n-1}, where $\xi, \eta \in S^{n-1}$ are joined with an edge if and only if $\langle \xi, \eta \rangle \in X$. The graphs $G(S^{n-1}, X)$ are called *forbidden inner product graphs* on S^{n-1}. We allow the possibility that $1 \in X$, which would correspond to each point in S^{n-1} having a self-loop.

We define the *independence ratio* of $G = G(S^{n-1}, X)$ by

$$\begin{aligned}\tilde{\alpha}(G) &:= \sup\{\lambda(I) : I \in \mathcal{L} \text{ is an independent set in } G\} \\ &= \sup\{\lambda(I) : I \in \mathcal{L} \text{ and } \langle \xi, \eta \rangle \notin X \text{ for any } \xi, \eta \in I\}.\end{aligned}$$

In case $1 \in X$, we have $\tilde{\alpha}(G) = 0$.

3.3 Combinatorial upper bound

Let us begin by deriving a simple "combinatorial" upper bound for the independence ratio of a forbidden inner product graph.

Proposition 3.1. *Let $n \geq 2$. If $G = G(S^{n-1}, X)$ contains a finite subgraph H, then $\tilde{\alpha}(G) \leq \alpha(H)/|V(H)|$.*

Proof. Assume that $V(H) \subset S^{n-1}$. Let I be an independent set, and take a uniform $\boldsymbol{O} \in SO(n)$. Let the random variable $\boldsymbol{Y}$ be the number of $\xi \in V(H)$ with $\boldsymbol{O}\xi \in I$. Since $\boldsymbol{O}\xi \in S^{n-1}$ is uniformly distributed for every $\xi \in V(H)$, we have by the linearity of expectation that $\mathbb{E}(\boldsymbol{Y}) = |V(H)|\,\lambda(I)$. On the other hand, $\boldsymbol{Y} \leq \alpha(H)$ for every outcome $\boldsymbol{O}$, since the points $\boldsymbol{O}\xi$ landing inside I form an independent set of the subgraph of G induced by all the points $\boldsymbol{O}\xi$, and this induced subgraph is isomorphic to H. Thus $\lambda(I) \leq \alpha(H)/|V(H)|$. □

Chapter 4

Circular sets avoiding a given inner product

We next use Proposition 3.1 to find the largest possible Lebesgue measure of a subset of the unit circle in $\mathbb{R}^2$ in which no two points lie at some fixed forbidden angle. This could also be phrased as the problem of finding $\tilde{\alpha}(G)$ for some appropriate forbidden inner product graph G.

Theorem 4.1. *Fix $t \in (0,1)$ and let $G = (V, E)$, where $V = [0,1)$, and where E is defined by declaring $(x, y) \in E$ if and only if $x - y \equiv \pm t \pmod 1$. Let*

$$\tilde{\alpha}(G) = \sup\{\lambda(I) : I \subset [0,1) \text{ is a Lebesgue measurable independent set in } G\},$$

where λ denotes Lebesgue measure. If t is rational and $t = p/q$ with p

and q coprime integers, then

$$\tilde{\alpha}(G) = \begin{cases} 1/2 & \textit{if } q \textit{ is even} \\ (q-1)/(2q) & \textit{if } q \textit{ is odd} \end{cases}.$$

In this case $\tilde{\alpha}(G)$ is attained as a maximum. If t is irrational then $\tilde{\alpha}(G) = 1/2$, but G has no independent set I with $\lambda(I) = 1/2$.

Proof. Consider the interval $[0, 1)$ as a group with the operation of addition modulo 1. Notice that $I \subset V$ is an independent set in G if and only if $I \cap (t + I) = \emptyset$. This implies immediately that $\tilde{\alpha}(G) \leq 1/2$ for all values of t.

Now suppose $t = p/q$ with p and q coprime integers, and suppose that q is even. Let S be any open subinterval of $[0, 1)$ of length $1/q$, and define $T_t : [0, 1) \to [0, 1)$ by $T_t x = x + t \mod 1$. From the fact that p and q are coprime, it follows that that intervals $S, T_t^2 S, \ldots, T_t^{q-4}, T_t^{q-2} S$ are disjoint, and therefore that their union, which we denote by I, has measure $1/2$. Also I is independent since $T_t I = T_t S \cup T_t^3 S \cup \cdots \cup T_t^{q-3} S \cup T_t^{q-1} S$ is disjoint from I. Therefore $\tilde{\alpha}(G) = 1/2$.

Next suppose q is odd. With notation as before, a similar argument shows that $S \cup T_t^2 S \cup \cdots \cup T_t^{q-3} S$ is an independent set in G of measure $(q-1)/(2q)$. Now applying Proposition 3.1 to an induced cylce of length q shows that this is largest possible.

Finally suppose that t is irrational. By Dirichlet's approximation theorem there exist infinitely many pairs of coprime integers p and q such that $|t - p/q| < 1/q^2$. For each such pair, let $\varepsilon = \varepsilon(q) = |t - p/q|$. Using an open interval I of length $\frac{1}{q} - \varepsilon$ and applying the same construction as above with $T_{p/q}$, one obtains an independent set of measure at least $((q-1)/2)(1/q - \varepsilon) = 1/2 - o(q)$. Therefore $\tilde{\alpha}(G) = 1/2$.

However this supremum can never be attained. Indeed, if $I \subset V$ is an independent set with $\lambda(I) = 1/2$, then $I \cap T_t I = \emptyset$ and $T_t I \cap T_t^2 I = \emptyset$. Since $\lambda(I) = 1/2$, this implies that I and $T_t^2 I$ differ by a nullset, contradicting the ergodicity of the irrational rotation T_t^2. □

Chapter 5

Spherical sets avoiding orthogonal pairs of points

The aim of this chapter is to prove that any set occupying more than a 0.313 fraction of the unit sphere in $\mathbb{R}^3$ must contain a pair of orthogonal vectors; in other words $\tilde{\alpha}(G(S^{n-1}, \{0\})) \leq 0.313$.

5.1 Gegenbauer polynomials and Schoenberg's theorem

Before beginning the core of the chapter, we briefly review the Gegenbauer polynomials and Schoenberg's theorem from the theory of spherical harmonics. For $\nu > -1/2$, define the *Gegenbauer weight function*

$$w_\nu(t) := (1 - t^2)^{(\nu-1/2)}, \quad (-1 < t < 1).$$

Applying the Gram-Schmidt process to the polynomials $1, t, t^2, \ldots$ using the inner product $\langle f, g\rangle = \int_{-1}^{1} f(t)g(t)w_\nu(t)\ dt$, one obtains the *Gegenbauer polynomials* $C_i^\nu(t)$ for the degrees $i = 0, 1, 2, \ldots$. We always use the normalization $C_i^\nu(1) = 1$. (Cf. [DX13, Section B.2])

For a fixed $n \geq 2$, a continuous function $f : [-1, 1] \to \mathbb{R}$ is called *positive definite* if for every set of distinct points $\xi_1, \ldots, \xi_s \in S^{n-1}$, the matrix $(f(\langle \xi_i, \xi_j\rangle))_{i,j=1}^s$ is positive semidefinite. The following theorem is known as Schoenberg's theorem ([DX13, Theorem 14.3.3]).

Theorem 5.1 (Schoenberg's theorem). *For $n \geq 2$, a continuous function $f : [-1, 1] \to \mathbb{R}$ is positive definite if and only there exist coefficients $a_i \geq 0$, for $i \geq 0$, such that*

$$f(t) = \sum_{i=0}^{\infty} a_i C_i^{(n-2)/2}(t),$$

where convergence on the right-hand side is absolute and uniform.

For a given positive definite function f, the coefficients a_i in Theorem 5.1 are unique and can be computed explicitly; a formula is given in [DX13, Equation 14.3.3].

We are especially interested in the case $n = 3$. Then $\nu = 1/2$, and the first few Gegenbauer polynomials $C_i^{1/2}(x)$ are

$$C_0^{1/2}(x) = 1, \quad C_1^{1/2}(x) = x, \quad C_2^{1/2}(x) = \frac{3}{2}x^2 - \frac{1}{2},$$
$$C_3^{1/2}(x) = \frac{5}{2}x^3 - \frac{3}{2}x, \quad C_4^{1/2}(x) = \frac{35}{8}x^4 - \frac{30}{8}x^2 + \frac{3}{8}.$$

5.2 An adjacency operator for an infinite graph

Let $n \geq 3$. For each $\xi \in S^{n-1}$ and $-1 < t < 1$, let $\sigma_{\xi,t}$ be the unique probability measure on the Borel subsets of S^{n-1} whose support is equal to the set $\xi^t := \{\eta \in S^{n-1} : \langle \eta, \xi \rangle = t\}$, and which is invariant under all rotations fixing ξ. For $f \in L^2(S^{n-1})$, define

$$(A_t f)(\xi) := \int_{\xi^t} f(\eta)\, d\sigma_{\xi,t}(\eta). \tag{5.1}$$

It is shown in Theorem 6.1 that A_t is well-defined and maps L^2 into L^2.

Lemma 5.2. *Let f and g be functions in $L^2(S^{n-1})$, let $\xi, \eta \in S^{n-1}$ be arbitrary points, and write $t = \langle \xi, \eta \rangle$. If $\boldsymbol{O} \in SO(n)$ is chosen uniformly at random with respect to the Haar measure on $SO(n)$, then*

$$\int_{S^{n-1}} f(\zeta)(A_t g)(\zeta)\, d\zeta = \mathbb{E}[f(\boldsymbol{O}\xi) g(\boldsymbol{O}\eta)].$$

Proof. Note that picking a point uniformly at random from S^{n-1} is equivalent to fixing an arbitrary point in S^{n-1}, and then applying to it a rotation $\boldsymbol{O} \in SO(n)$ chosen uniformly at random. We therefore have

$$\begin{aligned}\int_{S^{n-1}} f(\zeta)(A_t g)(\zeta)\, d\zeta &= \int_{SO(n)} f(O\xi)(A_t g)(O\xi)\, dO \\ &= \int_{SO(n)} f(O\xi) \int_{(O\xi)^t} g(\eta)\, d\sigma_{O\xi,t}(\eta)\, dO\end{aligned}$$

If H is the subgroup of all elements in $SO(n)$ which fix ξ, then the above integral can be rewritten

$$\int_{SO(n)} f(O\xi) \int_H g(Oh\eta)\, dh\, dO.$$

By Fubini's theorem, this integral is equal to

$$\begin{aligned}
&\int_H \int_{SO(n)} f(O\xi) g(Oh\eta)\, dO\, dh \\
&= \int_H \int_{SO(n)} f(Oh^{-1}\xi) g(O\eta)\, dO\, dh \\
&= \int_{SO(n)} f(O\xi) g(O\eta)\, dO,
\end{aligned}$$

where we use the right-translation invariance of the Haar integral on $SO(n)$ at the first equality, and the second equality follows by noting that the integrand is constant with respect to h. Recall that all Haar measures are normalized to have measure 1. $\square$

Lemma 5.3. *Suppose $f \in L^2(S^{n-1})$ and define $k_f : [-1,1] \to \mathbb{R}$ by*

$$k_f(t) = \mathbb{E}[f(\boldsymbol{O}\xi) f(\boldsymbol{O}\eta)], \tag{5.2}$$

where the expectation is taken over randomly chosen $\boldsymbol{O} \in SO(n)$, and $\xi, \eta \in S^{n-1}$, are any two points satisfying $\langle \xi, \eta \rangle = t$. Then $k_f(t)$ is defined for each $t \in [-1,1]$, and k_f is continuous and positive definite.

Proof. Fix any point $\xi_0 \in S^{n-1}$ and let $P : [-1,1] \to SO(n)$ be any continuous function satisfying $\langle \xi_0, P(t)\xi_0 \rangle = t$ for each $-1 \le t \le 1$.

We have

$$k_f(t) = \int_{SO(n)} f(O\xi_0) f(OP(t)\xi_0)\; dO \tag{5.3}$$

for each t. Being an inner product in $L^2(SO(n))$, the right-hand side of (5.3) exists for each $t \in [-1, 1]$. For each $O \in SO(n)$, let $R_O : L^2(SO(n)) \to L^2(SO(n))$ be the operator defined by $(R_O f)(O') = f(O'O)$ for each $O' \in SO(n)$, and define $F : SO(n) \to \mathbb{R}$ by $F(O) = f(O\xi_0)$. Since right-translation is continuous on $L^2(SO(n))$ [DE09, Lemma 1.4.2], the function $t \mapsto R_{P(t)}F$ is continuous from $[-1, 1]$ to $L^2(SO(n))$. Therefore

$$k_f(t) = \int_{SO(n)} F(O)(R_{P(t)}F)(O)\; dO.$$

It now follows that $k_f(t)$ is continuous in t.

To see that k_f is a positive definite function, fix arbitrary distinct points
$\xi_1, \ldots, \xi_s \in S^{n-1}$; we need to show that the $s \times s$ matrix $K = (k_f(\langle \xi_i, \xi_j \rangle))_{i,j=1}^s$ is positive semidefinite. But if $v = (v_1, \ldots, v_s)^t \in \mathbb{R}^s$ is any column vector, then

$$v^T K v = \int_{SO(n)} \left(\sum_{i=1}^{s} f(O\xi_i) v_i \right)^2 \; dO \geq 0.$$

□

5.3 A linear programming relaxation for independence ratio

Combining Lemma 5.3 with Schoenberg's theorem allows us to set up a linear program whose value upper bounds the measure of any independent set in $G = G(S^{n-1}, \{0\})$ for any $n \geq 3$. The same result appears in [BNdOFV09] and [dOF09]; our proof is slightly simpler than the ones presented there.

Theorem 5.4. $\tilde{\alpha}(G)$ *is no more than the value of the following infinite-dimensional linear program.*

$$\begin{gathered} \max x_0 \\ \sum_{i=0}^{\infty} x_i = 1 \\ \sum_{i=0}^{\infty} x_i C_i^{(n-2)/2}(0) = 0 \\ x_i \geq 0, \text{ for all } i = 0, 1, 2, \ldots. \end{gathered} \tag{5.4}$$

Proof. Let I be a Lebesgue measurable subset of S^{n-1} with $\lambda(I) > 0$, having the property that $\langle \xi, \eta \rangle \neq 0$ for any $\xi, \eta \in I$. We shall construct a feasible solution to the linear program (5.4) having value $\lambda(I)$. Let $k = k_{\mathbb{1}_I}$ be as in Lemma 5.3. Then k is a positive definite function satisfying $k(1) = \lambda(I)$ and $k(0) = 0$. By Theorem 5.1, k has an expansion in terms of the Gegenbauer polynomials:

$$k(t) = \sum_{i=0}^{\infty} a_i C_i^{(n-2)/2}(t), \tag{5.5}$$

where the convergence of the righthand side is uniform for on $[-1,1]$. Moreover, we have

$$\begin{aligned} a_0 &= \frac{\omega_{n-1}}{\omega_n} \int_{-1}^{1} k(t)(1-t^2)^{(n-3)/2}\, dt \\ &= \int_{S^{n-1}} k(\langle \xi, \xi_0 \rangle)\, d\xi, \end{aligned}$$

where $\xi_0 \in S^{n-1}$ can be any point. Since the above expression is constant with respect to ξ_0, it follows that if $\boldsymbol{O} \in SO(n)$ is picked uniformly at random, then by Lemma 5.2 we have

$$\begin{aligned} \int_{S^{n-1}} k(\langle \xi, \xi_0 \rangle)\, d\xi &= \int_{S^{n-1}} \mathbb{P}[\boldsymbol{O}\xi \in I, \boldsymbol{O}\xi_0 \in I]\, d\xi \\ &= \int_{S^{n-1}} \int_{S^{n-1}} \mathbb{P}[\boldsymbol{O}\xi \in I, \boldsymbol{O}\xi_0 \in I]\, d\xi\, de \\ &= \lambda(I)^2. \end{aligned}$$

We conclude that $a_0 = \lambda(I)^2$. Recall that $C_i^{(n-2)/2}(1) = 1$ for $i \geq 0$. Therefore setting $x_i = a_i/\lambda(I)$ for $i = 0, 1, 2, \ldots$ gives a feasible solution of value $\lambda(I)$ to program (5.4). $\square$

Unfortunately in the case $n = 3$, the value of (5.4) is at least $1/3$, which is the same bound obtained in Witsenhausen's original statement of the problem in [Wit74]. This can be seen from the feasible solution $x_0 = 1/3, x_2 = 2/3$ and $x_i = 0$ for all $i \neq 0, 2$.

5.4 Adding combinatorial constraints

Our aim now is to strengthen (5.4) for the case $n = 3$ by adding combinatorial inequalities coming from Proposition 3.1. We proceed as follows: Let p and q be coprime integers with $1/4 \le p/q \le 1/2$, and let $t_{p,q} = \sqrt{\frac{-\cos(2\pi p/q)}{1-\cos(2\pi p/q)}}$. If $\xi \in S^{n-1}$ is any point, then two orthogonal unit vectors with endpoints in $\xi^{t_{p,q}}$ make angle $2\pi p/q$ in the circle $\xi^{t_{p,q}}$. The circle therefore contains a cycle of length q, and applying Proposition 3.1 to this circle we obtain

$$(A_{t_{p,q}} \mathbb{1}_I)(\xi) \le (q-1)/2q$$

when q is odd. Since the inequality holds for every $\xi \in S^2$, we get

$$\begin{aligned} k(t_{p,q}) &= \int_{S^2} \mathbb{1}_I(\xi)(A_{t_{p,q}} \mathbb{1}_I)(\xi)\, d\xi \\ &\le \lambda(I)\frac{q-1}{2q}, \end{aligned}$$

and it follows that the inequalities

$$\sum_{i=0}^{\infty} x_i C_i^{1/2}(t_{p,q}) \le (q-1)/2q, \tag{5.6}$$

are valid for the relaxation and can be added to (5.4). The same holds for the inequalities $\sum_{i=0}^{\infty} x_i C_i^{1/2}(-t_{p,q}) \le (q-1)/2q$.

We have just proved the following result.

Theorem 5.5. $\tilde{\alpha}(G(S^2,\{0\}))$ *is no more than the value of the follow-*

ing infinite-dimensional linear program.

$$\max x_0$$
$$\sum_{i=0}^{\infty} x_i = 1$$
$$\sum_{i=0}^{\infty} x_i C_i^{1/2}(0) = 0 \tag{5.7}$$
$$\sum_{i=0}^{\infty} x_i C_i^{1/2}(\pm t_{p,q}) \leq (q-1)/2q, \quad (q \textit{ odd}, \ p,q \textit{ coprime})$$
$$x_i \geq 0, \textit{ for all } i = 0,1,2,\ldots.$$

5.5 Main theorem

The next theorem is the main result of Part II. Let $G = G(S^2, \{0\})$. Rather than attempting to find the exact value of the linear program (5.7), the idea will be to discard all but finitely many of the combinatorial constraints, and then to apply the weak duality theorem of linear programming. The dual linear program has only finitely many variables, and any feasible solution gives an upper bound for the value of program (5.7), and therefore also for $\tilde{\alpha}(G)$. At the heart of the proof is the verification of the feasibility of a particular dual solution which we give explicitly. While part of the verification has been carried out by computer in order to deal with the large numbers that appear, it requires only rational arithmetic and can therefore be considered rigorous.

Theorem 5.6. $\tilde{\alpha}(G) < 0.313$.

Proof. Consider the following linear program

$$\max\Big\{x_0 : \sum_{i=0}^{\infty} x_i = 1, \sum_{i=0}^{\infty} x_i C_i^{1/2}(0) = 0, \sum_{i=0}^{\infty} x_i C_i^{1/2}(t_{1,3}) \le 1/3, \tag{5.8}$$
$$\sum_{i=0}^{\infty} x_i C_i^{1/2}(t_{2,5}) \le 2/5, \sum_{i=0}^{\infty} x_i C_i^{1/2}(-t_{2,5}) \le 2/5,$$
$$x_i \ge 0, \text{ for all } i = 0, 1, 2, \dots \Big\}$$

The linear programming dual of (5.8) is the following.

$$\begin{gathered}\min\ b_1 + \frac{1}{3}b_{1,3} + \frac{2}{5}b_{2,5} + \frac{2}{5}b_{2,5-} \\ b_1 + b_0 + b_{1,3} + b_{2,5} + b_{2,5-} \ge 1 \\ b_1 + C_i^{1/2}(0)b_0 + C_i^{1/2}(t_{1,3})b_{1,3} + C_i^{1/2}(t_{2,5})b_{2,5} + C_i^{1/2}(-t_{2,5})b_{2,5-} \ge 0 \\ \text{for } i = 1, 2, \dots \\ b_1, b_0 \in \mathbb{R},\ b_{1,3}, b_{2,5}, b_{2,5-} \ge 0 \end{gathered} \tag{5.9}$$

By linear programming duality, any feasible solution for program (5.9) gives an upper bound for (5.8), and therefore also for $\tilde{\alpha}(G)$. So in order to prove the claim $\tilde{\alpha}(G) < 0.313$, it suffices to give a feasible solution to (5.9) having objective value no more than 0.313. Let

$$b = (b_1, b_0, b_{1,3}, b_{2,5}, b_{2,5-}) = \frac{1}{10^6}(128614, 404413, 36149, 103647, 327177).$$

It is easily verified that b satisfies the first constraint of (5.9) and that its objective value less than 0.313. To verify the infinite family of constraints

$$b_1 + C_i^{1/2}(0)b_0 + C_i^{1/2}(t_{1,3})b_{1,3} + C_i^{1/2}(t_{2,5})b_{2,5} + C_i^{1/2}(-t_{2,5})b_{2,5-} \geq 0 \tag{5.10}$$

for $i = 1, 2, \ldots,$ we apply Theorem 8.21.11 from [Sze92], which implies

$$|C_i^{1/2}(\cos\theta)| \leq \frac{\sqrt{2}}{\sqrt{\pi}\sqrt{\sin\theta}} \frac{\Gamma(i+1)}{\Gamma(i+3/2)} + \frac{1}{\sqrt{\pi}2^{3/2}(\sin\theta)^{3/2}} \frac{\Gamma(i+1)}{\Gamma(i+5/2)} \tag{5.11}$$

for each $0 < \theta < \pi$, where Γ denotes the Euler Γ-function. Note that $t_{1,3} = 1/\sqrt{3}$ and $t_{2,5} = 5^{-1/4}$. When

$$\theta \in A := \{\pi/2, \arccos t_{1,3}, \arccos t_{2,5}, \arccos -t_{2,5}\},$$

we have $\sin\theta \in \{1, \sqrt{\frac{2}{3}}, \gamma\}$, where $\gamma = \frac{2}{\sqrt{5+\sqrt{5}}}$. The righthand side of equation (5.11) is maximized at $\sin\theta = \gamma$ for each fixed i, and since the righthand side is decreasing in i, one can verify using rational arithmetic only that it is no greater than $128614/871386 = b_1/(b_0 + b_{1,3} + b_{2,5} + b_{2,5-})$ when $i \geq 40$, by evaluating at $i = 40$. Therefore,

$$\begin{aligned} &b_1 + C_i^{1/2}(0)b_0 + C_i^{1/2}(t_{1,3})b_{1,3} + C_i^{1/2}(t_{2,5})b_{2,5} + C_i^{1/2}(-t_{2,5})b_{2,5-} \\ &\geq b_1 - (b_0 + b_{1,3} + b_{2,5} + b_{2,5-}) \max_{\theta\in A}\{|C_i^{1/2}(\cos\theta)|\} \\ &\geq 0 \end{aligned}$$

when $i \geq 40$. It now suffices to check that b satsifies the constraints (5.10) for $i = 0, 1, \ldots, 39$. This can also be accomplished using rational arithmetic only. □

The rational arithmetic calculations required in the above proof were carried out with Mathematica. When verifying the upper bound for the righthand side of (5.11), it is helpful to recall the identity $\Gamma(i + 1/2) = (i - 1/2)(i - 3/2)\cdots(1/2)\sqrt{\pi}$. When verifying the constraints (5.10) for $i = 0, 1, \ldots, 39$, it can be helpful to observe that $t_{1,3}$ and $t_{2,5}$ are roots of the polynomials $x^2 - 1/3$ and $x^4 - 1/5$ respectively; this can be used to cut down the degree of the polynomials $C_i^{1/2}(x)$ to at most 3 before evaluating them. The ancillary folder of the `arxiv.org` version of [DP15] contains a *Mathematica* notebook that verfies all calculations.

The combinatorial inequalities of the form (5.6) we chose to include in the strengthened linear program (5.8) were found as follows: Let L_0 denote the linear program (5.4). We first find an optimal solution σ_0 to L_0. We then proceed recursively; having defined the linear program L_{i-1} and found an optimal solution σ_{i-1}, we search through the inequalities (5.6) until one is found for which σ_{i-1} is infeasible for L_{i-1}, and we strengthen L_{i-1} with that inequality to produce L_i. At each stage, an optimal solution to L_i is found by first solving the dual minimization problem, and then applying the complementary slackness theorem from linear programming to reduce L_i to a linear programming maximization problem with just a finite number of variables.

Adding more inequalities of the form (5.6) appears to give no improvement on the upper bound. Also adding the constraints $\sum_{i=0}^{\infty} x_i C_i^{1/2}(t) \geq 0$ for $-1 \leq t \leq 1$ appears to give no improvement. A small (basically insignificant) improvement can be achieved by allowing the odd cycles to embed into G in more general ways, for instance with the points lying on two different latitudes rather than just one.

Part III

Existence of measurable maximum independent sets in forbidden inner product graphs

Chapter 6

Revisiting adjacency operators of forbidden inner product graphs on S^{n-1}

In this chapter we investigate the adjacency operator defined in Equation (5.1), proving some properties of it which we will require in Chapter 7.

6.1 Boundedness and self-adjointness

Recall the definition (5.1) of the adjacency operator from Section 5.2:

$$(A_t f)(\xi) := \int_{\xi^t} f(\eta)\ d\sigma_{\xi,t}(\eta).$$

Theorem 6.1. *For every $t \in (-1,1)$, A_t is a bounded, self-adjoint operator mapping $L^2(S^{n-1})$ to $L^2(S^{n-1})$, having operator norm equal to* 1.

Proof. The right-hand side of (5.1) involves integration over nullsets of a function $f \in L^2(S^{n-1})$ which is only defined almost everywhere, and so strictly speaking one should argue that (5.1) really makes sense. In other words, given a particular representative f from its L^2-equivalence class, we need to check that the integral on the right-hand side of (5.1) is defined for almost all $\xi \in S^{n-1}$, and that the L^2-equivalence class of $A_t f$ does not depend on the particular choice of representative f.

Our main tool will be Minkowski's integral inequality (see e.g. [Fol99, Theorem 6.19]).

Let $e_n = (0, \ldots, 0, 1)$ be the n-th basis vector in $\mathbb{R}^n$ and let

$$S = \{(x_1, x_2, \ldots, x_n) : x_n = 0, x_1^2 + \cdots + x_{n-1}^2 = 1\}$$

be a copy of S^{n-2} inside $\mathbb{R}^n$. Considering f as a particular measurable function (not an L^2-equivalence class), we define $F : SO(n) \times S \to \mathbb{R}$ by

$$F(\rho, \eta) = f\left(\rho\left(te_n + \sqrt{1-t^2}\,\eta\right)\right), \qquad \rho \in SO(n),\ \eta \in S.$$

Let us formally check all the hypotheses of Minkowski's integral inequality applied to F, where $SO(n)$ is equipped with the Haar measure, and where S is equipped with the normalised Lebesgue measure; this will show that the function $\tilde{F} : SO(n) \to \mathbb{R}$ defined by $\tilde{F}(\rho) = \int_S F(\rho, \eta)\,\mathrm{d}\eta$ belongs to $L^2(SO(n))$.

Clearly the function F is measurable. To see that the function $\rho \mapsto F(\rho, \eta)$ belongs to $L^2(SO(n))$ for each fixed $\eta \in S$, simply note that

$$\int_{SO(n)} |F(\rho, \eta)|^2 \,\mathrm{d}\rho = \int_{SO(n)} \left| f(\rho(te_n + \sqrt{1-t^2}\,\eta)) \right|^2 \,\mathrm{d}\rho = \|f\|_2^2.$$

That the function $\eta \mapsto \|F(\cdot, \eta)\|_2$ belongs to $L^1(S)$ then also follows easily (in fact, this function is constant):

$$\int_S \left(\int_{SO(n)} |F(\rho, \eta)|^2 \,\mathrm{d}\rho \right)^{1/2} \mathrm{d}\eta = \int_S \|f\|_2 \,\mathrm{d}\eta = \|f\|_2.$$

Minkowski's integral inequality now gives that the function $\eta \mapsto F(\rho, \eta)$ belongs to $L^1(S)$ for a.e. ρ, that the function $\tilde{F}$ belongs to $L^2(SO(n))$, and that its norm can be bounded as follows:

$$\begin{aligned} \|\tilde{F}\|_2 &= \left(\int_{SO(n)} \left| \int_S F(\rho, \eta)\,\mathrm{d}\eta \right|^2 \mathrm{d}\rho \right)^{1/2} \\ &\leq \int_S \left(\int_{SO(n)} |F(\rho, \eta)|^2 \,\mathrm{d}\rho \right)^{1/2} \mathrm{d}\eta = \|f\|_2. \end{aligned} \tag{6.1}$$

Applying (6.1) to $f - g$ where g is a.e. equal to f, we conclude that the L^2-equivalence class of $\tilde{F}$ does not depend on the particular choice of representative f from its equivalence class.

Now $(A_t f)(\xi)$ is simply $\tilde{F}(\rho)$, where $\rho \in SO(n)$ can be any rotation such that $\rho e_n = \xi$. This shows that the integral in (5.1) makes sense for almost all $\xi \in S^{n-1}$.

We have $\|A_t\| \leq 1$ since for any $f \in L^2(S^{n-1})$,

$$\|A_t f\|_2 = \left(\int_{S^{n-1}} |(A_t f)(\xi)|^2 \,\mathrm{d}\xi\right)^{1/2} = \left(\int_{SO(n)} |(A_t f)(\rho e_n)|^2 \,\mathrm{d}\rho\right)^{1/2}$$
$$= \left(\int_{SO(n)} \left|\tilde{F}(\rho)\right|^2 \,\mathrm{d}\rho\right)^{1/2} \leq \|f\|_2,$$

by (6.1).

Applying A_t to the constant function 1 shows that $\|A_t\| = 1$. To see that A_t is self-adjoint, fix $\xi, \eta \in S^{n-1}$ that satisfy $\langle \xi, \eta \rangle = t$. Then Lemma 5.2 implies that for any $f, g \in L^2(S^{n-1})$,

$$\langle A_t f, g \rangle = \mathbb{E}_{\boldsymbol{O} \in SO(n)}[f(\boldsymbol{O}\xi) g(\boldsymbol{O}\eta)] = \langle f, A_t g \rangle.$$

□

6.2 Eigenvalues and eigenvectors

For $n \geq 2$ and $d \geq 0$, let $\mathscr{H}_d^n$ be the vector space of homogeneous polynomials $p(x_1, \ldots, x_n)$ of degree d in n variables belonging to the kernel of Laplace operator; that is

$$\frac{\partial^2 p}{\partial x_1^2} + \cdots + \frac{\partial^2 p}{\partial x_1^2} = 0.$$

Note that each $\mathscr{H}_d^n$ is finite-dimensional. The restrictions of the elements of $\mathscr{H}_d^n$ to the surface of the unit sphere are called the *spherical* *harmonics*. For fixed n, we have $L^2(S^{n-1}) = \oplus_{d=0}^{\infty}\mathscr{H}_d^n$ ([DX13, Theorem 2.2.2]); that is, each function in $L^2(S^{n-1})$ can be written as the infinite sum of elements from $\mathscr{H}_d^n$, $d = 0, 1, 2, \ldots$, with convergence in the L^2 norm.

The next lemma says that the eigenfunctions of the operators A_t are exactly the spherical harmonics. It extends the Funk-Hecke formula ([DX13, Theorem 1.2.9]) to the Dirac measures, obtaining the eigenvalues of A_t explicitly.

Proposition 6.2. *Let $t \in [-1, 1]$. Then for every spherical harmonic Y_d of degree d,*

$$(A_t Y_d)(\xi) := \int_{\xi^t} Y_d(\eta)\ d\sigma_{\xi,t}(\eta) = \mu_d(t) Y_d(\xi), \quad \xi \in S^{n-1},$$

where $\mu_d(t)$ is the constant

$$\mu_d(t) = C_d^{(n-2)/2}(t)(1-t^2)^{(n-3)/2} \Big/ C_d^{(n-2)/2}(1).$$

Proof. Let ds be Lebesgue measure on $[-1, 1]$ and let $\{f_\alpha\}_\alpha$ be a net of functions in $L^1([-1, 1])$ such that $\{f_\alpha\ ds\}$ converges to the Dirac point mass δ_t at t in the weak-* topology on the set of Borel measures on $[-1, 1]$. By [DX13, Theorem 1.2.9], we have

$$\int_{S^{n-1}} Y_d(\eta) f_\alpha(\langle \xi, \eta \rangle)\ d\eta = \mu_{d,\alpha} Y_d(\xi),$$

where

$$\mu_{d,\alpha} = \int_{-1}^{1} C_d^{(n-2)/2}(s)(1-s^2)^{(n-3)/2} f_\alpha(s)\ ds \Big/ C_d^{(n-2)/2}(1)$$

and taking limits finishes the proof. □

The next lemma is a general fact about weakly convergent sequences in a Hilbert space.

Lemma 6.3. *Let $\mathcal{H}$ be a Hilbert space and let $K : \mathcal{H} \to \mathcal{H}$ be a compact operator. Suppose $\{x_i\}_{i=1}^{\infty}$ is a sequence in $\mathcal{H}$ converging weakly to $x \in \mathcal{H}$. Then*

$$\lim_{i\to\infty} \langle Kx_i, x_i\rangle = \langle Kx, x\rangle.$$

Proof. Let C be the maximum of $\|x\|$ and $\sup_{i\geq 1} \|x_i\|$, which is finite by the principle of uniform boundedness. Let $\{K_m\}_1^{\infty}$ be a sequence of finite rank operators such that $K_m \to K$ in the operator norm as $m \to \infty$. Clearly

$$\lim_{i\to\infty} \langle K_m x_i, x_i\rangle = \langle K_m x, x\rangle$$

for each $m = 1, 2, \ldots$. Let $\varepsilon > 0$ be given and choose m_0 so that $\|K - K_{m_0}\| < \varepsilon$. Choosing i_0 so that $|\langle K_{m_0} x_i, x_i\rangle - \langle K_{m_0} x, x\rangle| < \varepsilon$ whenever $i \geq i_0$, we have

$$\begin{aligned}
&|\langle Kx_i, x_i\rangle - \langle Kx, x\rangle| \\
\leq &|\langle Kx_i, x_i\rangle - \langle K_{m_0} x_i, x_i\rangle| + |\langle K_{m_0} x_i, x_i\rangle - \langle K_{m_0} x, x\rangle| \\
&\qquad + |\langle K_{m_0} x, x\rangle - \langle Kx, x\rangle| \\
\leq &\|K - K_{m_0}\| C^2 + \varepsilon + \|K - K_{m_0}\| C^2 \\
< &(2C^2 + 1)\varepsilon,
\end{aligned}$$

and the lemma follows. □

6.3 Compactness

The next corollary is a result obtained in [BDdOFV14].

Corollary 6.4. *If $n \geq 3$ and $t \in (-1, 1)$, then A_t is compact.*

Proof. The operator A_t is diagonalizable by Proposition 6.2, since the spherical harmonics form an orthonormal basis for $L^2(S^{n-1})$. It therefore suffices to show that its eigenvalues have only zero as a cluster point.

By [Sze92, Theorem 8.21.8] and Proposition 6.2, the eigenvalues $\mu_d(t)$ tend to zero as $d \to \infty$. The eigenspace corresponding to the eigenvalue $\mu_d(t)$ is precisely the vector space of spherical harmonics of degree d, which is finite dimensional. Therefore A_t is compact. □

Chapter 7

Attainment of the independence ratio of forbidden inner product graphs on S^{n-1}

Let $n \geq 2$ and $X \subset [-1, 1]$, and let $G = G(S^{n-1}, X)$. From Theorem 4.1 we know that the supremum in the definition of $\tilde{\alpha}(G)$ is sometimes attained as a maximum, and sometimes not. It is therefore interesting to ask when a maximizer exists. The main positive result in this direction is Theorem 7.3, which says that a largest measurable independent set always exists when $n \geq 3$. Remarkably, this result holds under no additional restrictions (not even Lebesgue measurability) on the set X of forbidden inner products.

In this chapter we first prove a technical result, and we then put it together with the facts established about adjacency operators in Chapter 6 in order to prove the main result, Theorem 7.3, which says that $G(S^{n-1}, X)$ has maximum measurable independent sets for any X, provided $n \geq 3$. We then conclude by proving that the independence ratio of $G(S^{n-1}, X)$ does not change if X is replaced with its closure.

For the remainder of this section we suppose $n \geq 3$.

7.1 A lemma concerning pairs of Lebesgue density points

This aim of this section is to prove Lemma 7.1, which is crucial in the proof of the existence of measurable maximum independent sets in forbidden inner product graphs. Essentially it says that if some pair of Lebesgue density points of a subset I of S^{n-1} make inner product t, then there are "many" pairs of points in I making inner product t.

Lemma 7.1. *Suppose $n \geq 3$ and let $I \subset S^{n-1}$ be a Lebesgue measurable set with $\lambda(I) > 0$. Define $k : [-1, 1] \to \mathbb{R}$ by*

$$k(t) = \mathbb{E}[\mathbb{1}_I(\boldsymbol{O}\xi)\mathbb{1}_I(\boldsymbol{O}\eta)]$$

as in Lemma 5.3 with $f = \mathbb{1}_I$. If $\xi_1, \xi_2 \in S^{n-1}$ are Lebesgue density points of I, then $k(\langle \xi_1, \xi_2 \rangle) > 0$.

For each $\xi \in S^{n-1}$, let $C_h(\xi)$ be the open spherical cap of height h in S^{n-1} centred at ξ. Recall that $C_h(\xi)$ has volume proportional to

$\int_{1-h}^{1}(1-t^2)^{(n-3)/2}\,\mathrm{d}t$.

Lemma 7.2. *For each* $\xi \in S^{n-1}$*, we have* $\lambda(C_h(\xi)) = \Theta(h^{(n-1)/2})$*, and*
$\lambda(C_{h/2}(\xi)) \geq \lambda(C_h(\xi))/2^{(n-1)/2} - o(h^{(n-1)/2})$ *as* $h \to 0^+$.

Proof. If $f(h) = \int_{1-h}^{1}(1-t^2)^{(n-3)/2}\,\mathrm{d}t$, then we have $\frac{df}{dh}(h) = (2h - h^2)^{(n-3)/2}$. Since $f(0) = 0$, the smallest power of h occurring in $f(h)$ is of order $(n-1)/2$. This gives the first result. For the second, note that the coefficient of the lowest order term in $f(h)$ is $2^{(n-1)/2}$ times that of $f(h/2)$. ☐

Proof of Lemma 7.1. Let $t = \langle \xi_1, \xi_2 \rangle$. If $t = 1$, then the conclusion holds since $k(1) = \lambda(I) > 0$. If $t = -1$, then $\xi_2 = -\xi_1$, and by the Lebesgue density theorem we can choose $h > 0$ small enough that $\lambda(C_h(\xi_i) \cap I) > \frac{2}{3}\lambda(C_h(\xi_i))$ for $i = 1, 2$. Therefore,

$$\begin{aligned} k(-1) &= \mathbb{E}[\mathbb{1}_I(\boldsymbol{O}\xi_1)\mathbb{1}_I(\boldsymbol{O}(-\xi_1))] \\ &\geq \mathbb{E}[\mathbb{1}_{I\cap C_h(\xi_2)}(\boldsymbol{O}\xi_1)\mathbb{1}_{I\cap C_h(\xi_2)}(\boldsymbol{O}(-\xi_1))] \geq \frac{1}{3}\lambda(C_h(\xi_1)). \end{aligned}$$

From now on we may therefore assume $-1 < t < 1$. By Lemma 5.2 we have $k(t) = \int_{S^{n-1}} f(\zeta)(A_t g)(\zeta)\ d\zeta$. Let $h > 0$ be a small number which will be determined later. Suppose $x \in C_h(\xi_1)$. The intersection $x^t \cap C_h(\xi_2)$ is a spherical cap in the $(n-2)$-dimensional sphere x^t having height proportional to h; this is because $C_h(\xi_2)$ is the intersection of S^{n-1} with a certain halfspace H, and $x^t \cap C_h(\xi_2) = x^t \cap H$. We have $\sigma_{x,t}(x^t \cap C_h(\xi_2)) = \Theta(h^{(n-2)/2})$ by Lemma 7.2, and it follows that there exists $D > 0$ such that $\sigma_{x,t}(x^t \cap C_h(\xi_2)) \le Dh^{(n-2)/2}$ for sufficiently small $h > 0$.

If $x \in C_{h/2}(\xi_1)$, then $x^t \cap C_{h/2}(\xi_2) \neq \emptyset$ since x^t is just a rotation of the hyperplane ξ_1^t through an angle equal to the angle between x and ξ_1. Therefore $x^t \cap C_h(\xi_2)$ is a spherical cap in x^t having height at least $h/2$.

Thus there exists $D' > 0$ such that $\sigma_{x,t}(x^t \cap C_h(\xi_2)) \ge D'h^{(n-2)/2}$ for all $x \in C_{h/2}(\xi_1)$, by Lemma 7.2.

Now choose $h > 0$ small enough that $\lambda(C_h(\xi_i) \cap I) \ge (1 - \frac{D'}{2^n D})\lambda(C_h(\xi_i))$ for $i = 1, 2$; this is possible by the Lebesgue density theorem since ξ_1 and ξ_2 are density points. We have by Lemma 5.2 that

$$k(t) = \mathbb{P}[\boldsymbol{\eta}_1 \in I, \boldsymbol{\eta}_2 \in I],$$

if $\boldsymbol{\eta}_1$ is chosen uniformly at random from S^{n-1}, and if $\boldsymbol{\eta}_2$ is chosen uniformly at random from $\boldsymbol{\eta}_1^t$. Then

$$\begin{aligned} k(t) &\ge \mathbb{P}[\boldsymbol{\eta}_1 \in I \cap C_h(\xi_1), \boldsymbol{\eta}_2 \in I \cap C_h(\xi_2)] \\ &\ge \mathbb{P}[\boldsymbol{\eta}_1 \in C_h(\xi_1), \boldsymbol{\eta}_2 \in C_h(\xi_2)] - \mathbb{P}[\boldsymbol{\eta}_1 \in C_h(\xi_1) \setminus I, \boldsymbol{\eta}_2 \in C_h(\xi_2)] \\ &\qquad - \mathbb{P}[\boldsymbol{\eta}_1 \in C_h(\xi_1), \boldsymbol{\eta}_2 \in C_h(\xi_2) \setminus I]. \end{aligned}$$

The first probability is at least

$$D'h^{(n-2)/2}\lambda(C_{h/2}(\xi_1)) \ge \frac{D'}{2^{(n-1)/2}} h^{(n-2)/2}\lambda(C_h(\xi_1)) - o(h^{(2n-3)/2})$$

by Lemma 7.2. The second and third probabilities are each no more than

$$\frac{D'}{2^n D}\lambda(C_h(\xi_1))Dh^{(n-2)/2} = \frac{D'}{2^n}\lambda(C_h(\xi_1))h^{(n-2)/2}$$

for sufficiently small $h > 0$, and therefore by the first part of Lemma 7.2,

$$k(t) \geq \frac{D'}{2^{(n-1)/2}} \lambda(C_h(\xi_1)) h^{(n-2)/2} - o(h^{(2n-3)/2}) - \frac{D'}{2^{n-1}} \lambda(C_h(\xi_1)) h^{(n-2)/2},$$

and this is strictly positive for sufficiently small $h > 0$. □

7.2 Attainment of the independence ratio for $n \geq 3$

Theorem 7.3. *Suppose $n \geq 3$ and let X be any subset of $[-1, 1]$. Then $G(S^{n-1}, X)$ has an independent set $I \subset S^{n-1}$ such that $\lambda(I) = \tilde{\alpha}(G)$.*

Proof. Let $G = G(S^{n-1}, X)$. If $1 \in X$ then every independent set for G is empty. Otherwise, let $\{I_i\}_{i=1}^{\infty}$ be a sequence of measurable independent sets of G such that $\lim_{i\to\infty} \lambda(I_i) = \tilde{\alpha}(G)$. Passing to a subsequence if necessary, we may suppose that the sequence $\{\mathbb{1}_{I_i}\}$ of characteristic functions converges weakly in $L^2(S^{n-1})$; let h be its limit. Then $0 \leq h \leq 1$ almost everywhere since $0 \leq \mathbb{1}_{I_i} \leq 1$ for every i.

Denote by I' the set $h^{-1}((0, 1])$, and let I be the set of Lebesgue density points of I'. We claim that I is an independent set for G.

For all $t \in X \setminus \{-1\}$, the operator $A_t : L^2(S^{n-1}) \to L^2(S^{n-1})$ is compact and self-adjoint by Theorem 6.1 and Corollary 6.4. Since $\langle A_t \mathbb{1}_{I_i}, \mathbb{1}_{I_i} \rangle = 0$ for each i, Lemma 6.3 implies $\langle A_t h, h \rangle = 0$. Since $0 \leq h \leq 1$, it follows from the definition of A_t that $\langle A_t \mathbb{1}_{I'}, \mathbb{1}_{I'} \rangle = 0$, and therefore also that $\langle A_t \mathbb{1}_I, \mathbb{1}_I \rangle = 0$. But if there exist points $\xi, \eta \in I$ with $t_0 = \langle \xi, \eta \rangle \in X \setminus \{-1\}$, then $\langle A_{t_0} \mathbb{1}_I, \mathbb{1}_I \rangle > 0$ by Lemma 7.1.

Therefore I is an independent set in $G(S^{n-1}, X \setminus \{-1\})$, so we will be done if we can show that there is no pair of points $\xi, -\xi \in I$ when $-1 \in X$. Since ξ and $-\xi$ are Lebesgue density points of I, there is a spherical cap C centred at ξ such that $\lambda(I \cap C) > \frac{2}{3}\lambda(C)$ and $\lambda(I \cap (-C)) > \frac{2}{3}\lambda(C)$. The same applies to I_i for all large i. But this contradicts the fact that I_i and its reflection $-I_i$ are disjoint for every i.

We conclude that I is an independent set, and finally, we have

$$\lambda(I) = \lambda(I') \geq \int h = \langle \mathbb{1}_{S^{n-1}}, h \rangle = \lim_{i \to \infty} \langle \mathbb{1}_{S^{n-1}}, \mathbb{1}_{I_i} \rangle = \lim_{i \to \infty} \lambda(I_i) = \tilde{\alpha}(G),$$

whence $\lambda(I) = \tilde{\alpha}(G)$ by the definition of $\tilde{\alpha}$. $\square$

7.3 Invariance of the independence ratio under taking the closure of the forbidden inner product set

Again let $n \geq 2$, $X \subset [-1, 1]$, and let $G = G(S^{n-1}, X)$. When X is replaced with its closure $\overline{X}$, the resulting graph $\overline{G} = G(S^{n-1}, \overline{X})$ obviously has fewer independent sets. It is therefore surprising that the following theorem should be true.

Theorem 7.4. *Let X be an arbitrary subset of* $[-1,1]$. *Let* $G = G(S^{n-1}, X)$ *and* $\overline{G} = (S^{n-1}, \overline{X})$. *Then* $\tilde{\alpha}(G) = \tilde{\alpha}(\overline{G})$. *In particular* $\tilde{\alpha}(G) = 0$ *if* $1 \in \overline{X}$.

Proof. Clearly $\tilde{\alpha}(G) \geq \tilde{\alpha}(\overline{G})$. For the reverse inequality, let $I' \subset S^{n-1}$ be any measurable independent set for G, let $I \subset I'$ be the set of Lebesgue density points of I', and define $k : [-1,1] \to \mathbb{R}$ by $k(t) = k_{\mathbb{1}_I}(t) = \int_{S^{n-1}} f(\xi)(A_t f)(\xi)\, d\xi$ as in Lemma 5.3. Then k is continuous, and since $k(t) = 0$ for every $t \in X$, it follows that $k(t) = 0$ for every $t \in \overline{X}$. Lemma 7.1 now implies that I is an independent set for $\overline{G}$. Since I' was arbitrary and since $\lambda(I) = \lambda(I')$, we have proven $\tilde{\alpha}(G) \leq \tilde{\alpha}(\overline{G})$. □

Part IV

Positivity for the analyst

Chapter 8

Functions and measures of positive type on compact groups

The aim of this chapter is to review the basic facts from representation theory and harmonic analysis which we will use later on. The layout of the chapter is as follows: In Section 8.1 we recall a number of basic facts from the representation theory and harmonic analysis of abstract compact groups, and we fix notation for the rest of the chapter. Section 8.2 defines functions of positive type, and recalls several key facts about them. Section 8.3 treats the basic Fourier analysis of measures and defines measures of positive type.

8.1 Representation theory and harmonic analysis preliminaries

A *topological group* Γ is a topological space, endowed with the structure of a group such that the group operations are continuous; that is, $x \mapsto x^{-1}$ is a continuous map from Γ to Γ and $(x, y) \mapsto xy$ is a continuous map from $\Gamma \times \Gamma$ to Γ.

By a *compact group*, we shall mean a topological group whose underlying topological space is compact and Hausdorff. All compact groups discussed in this thesis will be assumed to be metrizable.[1] Every compact group Γ has a regular finite Borel measure λ which is both left- and right-translation invariant, meaning that $\lambda(E) = \lambda(xE) = \lambda(Ex)$ for all $x \in \Gamma$ and all Borel subsets E of Γ. The measure λ is called the *Haar measure* of Γ, and it is unique up to scaling by positive numbers. We always take λ to be scaled so that $\lambda(\Gamma) = 1$. We will sometimes write dx in place of $d\lambda(x)$, and we may use $\int f \; d\lambda$ or just $\int f$ to mean $\int f(x) \; dx$. For a detailed treatment of the Haar measure, see for instance [Fol95, Chapter 2].

For $1 \leq p < \infty$, the set $L^p(\Gamma)$ will be defined as the set of Borel measurable functions $f : \Gamma \to \mathbb{C}$ for which $\int |f(x)|^p \; dx < \infty$. By identifying functions that agree λ-almost everywhere, one makes $L^p(\Gamma)$ into a Banach space with the norm $\|f\|_p := \left(\int |f(x)|^p \; dx\right)^{1/p}$. Strictly speaking, one should consider the elements of $L^p(\Gamma)$ as equivalence classes of functions; here, when we say $f \in L^p(\Gamma)$, it will always be

[1]This is in order to ensure that $L^2(\Gamma)$ is separable.

clear from context whether f is being thought of as a function or an equivalence class. If $f : \Gamma \to \mathbb{C}$ is Borel measurable, we define

$$\|f\|_\infty := \inf \Big\{ a \geq 0 : \lambda(\{x \in \Gamma : |f(x)| > a\}) = 0 \Big\},$$

with the convention that $\inf \emptyset = \infty$. We next define $L^\infty(\Gamma)$ to be the set of Borel measurable functions $f : \Gamma \to \mathbb{C}$ for which $\|f\|_\infty < \infty$. Again, $L^\infty(\Gamma)$ becomes a Banach space when one identifies functions agreeing λ-almost everywhere.

Given two functions $f, g \in L^1(\Gamma)$, we define their *convolution product* as

$$f * g(x) = \int f(y) g(y^{-1}x) \; dy.$$

This expression is valid for λ-almost all $x \in \Gamma$, and one can show that $f * g \in L^1(\Gamma)$; in fact, one has the inequality $\|f * g\|_1 \leq \|f\|_1 \|g\|_1$. We also define the *involution* f^* of $f \in L^1(\Gamma)$ by the formula

$$f^*(x) = \overline{f(x^{-1})}, \;\; (x \in \Gamma).$$

These convolution and involution operations turn $L^1(\Gamma)$ into a Banach $*$-algebra, called the *group algebra of* Γ.

A *unitary representation* of a compact group Γ is a group homomorphism π from Γ into the group $U(\mathcal{H}_\pi)$ of unitary operators on some nontrivial complex Hilbert space $\mathcal{H}_\pi$, which is continuous in the strong operator topology; that is, the map $x \mapsto \pi(x)u$ is continuous from Γ to $\mathcal{H}_\pi$ for each fixed $u \in \mathcal{H}_\pi$. We will refer to unitary representations simply as *representations*. Representations always satisfy the identities $\pi(xy) = \pi(x)\pi(y)$ and $\pi(x^{-1}) = \pi(x)^*$, where $\pi(x)^*$ denotes the adjoint of $\pi(x)$. The *dimension* or the *degree* of π is the dimension of $\mathcal{H}_\pi$, which will be denoted d_π.

A vector subspace $\mathcal{M} \subset \mathcal{H}_\pi$ is called π-invariant if $\pi(x)u \in \mathcal{M}$ for all $u \in \mathcal{M}$ and all $x \in \Gamma$. The representation π is called *irreducible* if the only π-invariant subspaces are $\{0\}$ and $\mathcal{H}_\pi$ itself. Two representations π and π' are called *(unitarily) equivalent* if there exists a unitary operator $U : \mathcal{H}_\pi \to \mathcal{H}_{\pi'}$ such that $\pi'(x) = U\pi(x)U^{-1}$ for all $x \in \Gamma$.

The next theorem is Theorem 5.2 in [Fol95].

Theorem 8.1. *If* Γ *is compact, then every irreducible representation of* Γ *is finite-dimensional.*

For each unitary equivalence class of irreducible representations, choose a representative π, and let $\widehat{\Gamma}$ be the collection of all such π.

For $f \in L^1(\Gamma)$, we define the *Fourier transform* $\hat{f}$ of f by

$$\hat{f}(\pi) = \int f(x)\pi(x)^*\,dx, \quad \text{for all } \pi \in \widehat{\Gamma}.$$

The value of the above integral is an operator on $\mathcal{H}_\pi$; see [Fol95, Appendix 3] for a general treatment of vector-valued integrals. For each $x \in \Gamma$, one can express $\pi(x)$ in terms of some fixed basis for $\mathcal{H}_\pi$, and for this reason there is usually no harm in thinking of $\hat{f}(\pi)$ as a $d_\pi \times d_\pi$ matrix when d_π is finite.

The next theorem is an immediate consequence of the Peter-Weyl theorem (see for instance [Fol95, Theorem 5.12]):

Theorem 8.2 (Fourier inversion)**.** *If* Γ *is compact, then for* $f \in L^2(\Gamma)$, *we have*

$$f(x) = \sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}(\hat{f}(\pi)\pi(x)),$$

where Tr *denotes the trace, and where the sum converges in the* L^2 *sense.*

The next result pulls together some basic facts about the Fourier transform which we will use later on without further mention:

Proposition 8.3. *For every* $f, g \in L^1(\Gamma)$, $\pi \in \widehat{\Gamma}$, *and* $a, b \in \mathbb{C}$, *we have*

(a) $(af + bg)\widehat{\ }(\pi) = a\widehat{f}(\pi) + b\widehat{g}(\pi)$

(b) $(f^*)\widehat{\ }(\pi) = \widehat{f}(\pi)^*$

(c) $(f * g)\widehat{\ }(\pi) = \widehat{g}(\pi)\widehat{f}(\pi)$

(d) *(Parseval's identity) If* $f \in L^2$, *then*

$$\|f\|_2^2 = \sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}(\hat{f}(\pi)^* \hat{f}(\pi)).$$

Note the following useful consequence of Proposition 8.3 (d):

$$\langle f, g \rangle = \frac{1}{2} \sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}\left[\hat{f}(\pi)^* \hat{g}(\pi) + \hat{g}(\pi)^* \hat{f}(\pi)\right], \qquad (8.1)$$

for $f, g \in L^2$.

For any function $f : \Gamma \to \mathbb{C}$, we can define the left and right translates of f as

$$(L_y f)(x) = f(y^{-1}x), \text{ and } (R_y f)(x) = f(xy).$$

The following fact, which is Proposition 2.41 in [Fol95], will be useful for us later.

Proposition 8.4. *If $1 \le p < \infty$ and $f \in L^p(\Gamma)$ then $\|L_y f - f\|_p$ and $\|R_y f - f\|_p$ tend to zero as y tends to e.*

8.2 Functions of positive type

The aim of this section is to define functions of positive type on a compact group, and to collect several useful facts about them which we will need later on. With some minor modifications, we follow the exposition given in [Fol95, Section 3.3]; all the proofs can be found there, except for the proof of Corollary 8.9, which we prove here. Let Γ be a compact group and let λ be its Haar measure. We say that $\phi \in L^\infty(\Gamma)$ is of *positive type* if

$$\int (f^* * f)\phi \, d\lambda \ge 0, \tag{8.2}$$

for every $f \in L^1(\Gamma)$. Note that by a straightforward density argument, when checking that a function is of positive type, it suffices to check (8.2) only for $f \in C(\Gamma)$.

Proposition 8.5. *For any $f \in L^2(\Gamma)$, the function $f^* * f$ is continuous and of positive type.*

The next theorem gives a complete characterization of positive type functions.

Theorem 8.6. *Let π be a unitary representation of Γ, take $u \in \mathcal{H}_\pi$, and let $\phi(x) = \langle \pi(x)u, u \rangle$. Then ϕ is a function of positive type. Conversely, if ϕ is a function of positive type on Γ, then there exists a unitary representation π of Γ and a vector $u \in \mathcal{H}_\pi$ such that $\phi(x)$ is equal to $\langle \pi(x)u, u \rangle$ for λ-almost every $x \in \Gamma$.*

The main importance of Theorem 8.6 for us will actually come from its corollaries.

Corollary 8.7. *The following statements hold:*

(a) *Every function of positive type on Γ agrees almost everywhere with a continuous function.*

(b) *If ϕ is a continuous function of positive type, then $\|\phi\|_\infty = \phi(e)$ and $\phi(x^{-1}) = \overline{\phi(x)}$. In particular, real-valued continuous functions of positive type are even.*

A function $\phi : \Gamma \to \mathbb{C}$ is called *positive definite* if for every positive integer n and distinct points $x_1, \ldots, x_n \in \Gamma$, the $n \times n$ matrix $(\phi(x_j^{-1}x_i))_{i,j=1}^n$ is positive semidefinite. The relation between positive definite functions and functions of positive type is given in the next proposition.

Proposition 8.8. *A continuous function $\phi : \Gamma \to \mathbb{C}$ is positive definite if and only if it is of positive type.*

An important consequence of Proposition 8.8 for us is the next result, which will become crucial in determining the dual cone of the cone of positive type functions in Chapter 10.

Corollary 8.9. *If ϕ and η are two continuous functions of positive type, then their pointwise product $\phi\eta$ is also of positive type.*

Proof. This follows from Propositions 8.8: The matrix $(\phi(x_j^{-1}x_i)\eta(x_j^{-1}x_i))_{i,j}$ is positive semidefinite for every $x_1, \ldots, x_n \in \Gamma$, being the Hadamard (entrywise) product of the positive semidefinite matrices $(\phi(x_j^{-1}x_i))_{i,j}$ and $(\eta(x_j^{-1}x_i))_{i,j}$. □

8.3 Fourier analysis of measures and measures of positive type

Recall that if X is any locally compact Hausdorff space and μ is a (positive) measure on the Borel σ-algebra $\mathcal{M}$ of X, then μ is called *regular* [Rud87, Definition 2.15] if

1. $\mu(E) = \inf\{\mu(V) : E \subset V, V \text{ open}\}$ for every $E \in \mathcal{M}$; and
2. $\mu(E) = \sup\{\mu(K) : K \subset E, K \text{ compact}\}$ for every open set E, and for every $E \in \mathcal{M}$ with $\mu(E) < \infty$.

If μ is a complex Borel measure, its *total variation* $|\mu|$ is defined by the equation

$$|\mu|(E) = \sup \sum_i |\mu(E_i)|,$$

where the supremum is taken over all countable partitions $\{E_i\} \subset \mathcal{M}$ of E. It is a fact [Rud87, Theorem 6.2] that $|\mu|$ is actually a measure and [Rud87, Theorem 6.4] that $|\mu|(X) < \infty$. We say that μ is *regular* if $|\mu|$ is regular in the sense defined above.

Now suppose Γ is a compact group. Let $C(\Gamma)$ denote the vector space of continuous functions $f : \Gamma \to \mathbb{C}$, and let $\mathcal{M}(\Gamma)$ denote the vector space of complex regular Borel measures on Γ equipped with the total variation norm; that is $\|\nu\| = |\nu|(\Gamma)$ for all $\nu \in \mathcal{M}(\Gamma)$.

A measure $\nu \in \mathcal{M}(\Gamma)$ is of *positive type* if $\int f^* * f \; d\nu \geq 0$ for all $f \in C(\Gamma)$. A function $\phi \in L^\infty(\Gamma)$ is therefore of positive type if $\phi \; d\lambda$ is a measure of positive type. The *involution* μ^* of a measure $\mu \in \mathcal{M}(\Gamma)$ is defined by $\mu^*(E) = \overline{\mu(E^{-1})}$ for each Borel subset $E \subset \Gamma$, where $E^{-1} = \{x^{-1} : x \in E\}$. The Riesz representation theorem (see e.g. [Rud87, Theorem 6.19]) says that every bounded linear functional on $C(\Gamma)$ is represented by an element of $\mathcal{M}(\Gamma)$; more precisely, for each bounded linear functional ψ on $C(\Gamma)$, there exists a $\omega \in \mathcal{M}(\Gamma)$ such that $\psi(f) = \int f \; d\omega$ for all $f \in C(\Gamma)$. If $\nu \in \mathcal{M}(\Gamma)$, then the *convolution product* $\mu * \nu$ is defined via the Riesz representation theorem as the unique regular Borel measure ω satisfying

$$\int f(x) \; d\omega(x) = \int \int f(xy) \; d\mu(x) \; d\nu(x),$$

for all $f \in C(\Gamma)$.

The *Fourier transform* of $\mu \in \mathcal{M}(\Gamma)$ is defined by

$$\hat{\mu}(\pi) := \int \pi(x)^* \; d\mu(x) \quad \text{for all } \pi \in \widehat{\Gamma}. \tag{8.3}$$

As with Fourier transforms of functions, $\hat{\mu}(\pi)$ can be thought of either as an operator on $\mathcal{H}_\pi$, or as a $d_\pi \times d_\pi$ matrix with respect to some fixed basis for $\mathcal{H}_\pi$. One proves that the integral (8.3) converges as follows: Let $u, v \in \mathcal{H}_\pi$ be arbitrary; then $x \mapsto \langle \pi(x)^* u, v\rangle$ is a continuous function, and $|\langle \pi(x)^* u, v\rangle| \leq \|u\|\|v\|$. It follows that $\int \langle \pi(x)^* u, v\rangle \; d\mu(x)$ exists, and $\hat{\mu}(\pi)$ is then defined as the unique operator on $\mathcal{H}_\pi$ for which $\langle \hat{\mu}(\pi)u, v\rangle = \int \langle \pi(x)^* u, v\rangle \; d\mu(x)$ for all $u, v \in \mathcal{H}_\pi$.

The following facts about Fourier transforms of measures are easily verified:

Proposition 8.10. *For any $\mu, \nu \in \mathcal{M}(\Gamma)$, $\pi \in \widehat{\Gamma}$, and $a, b \in \mathbb{C}$, we have*

(a) $(a\mu + b\nu)\widehat{\ }(\pi) = a\widehat{\mu}(\pi) + b\widehat{\nu}(\pi)$

(b) $\widehat{\mu * \nu}(\pi) = \widehat{\nu}(\pi)\widehat{\mu}(\pi)$

(c) $(\mu^*)\widehat{\ } = \widehat{\mu}^*$

(d) μ *is of positive type if and only if* $\overline{\mu}$ *is of positive type.*

Chapter 9

Bochner's theorem for compact groups

In this chapter we state and prove two versions of Bochner's theorem we will need later on. Section 9.1 is about the first version, which applies to continuous functions on compact groups. This theorem is well-known, but we present a streamlined proof which uses only the tools we have developed here. The proof appears simpler than the one in [HR94]. Section 9.2 then discusses a version for measures of positive type. While the proof is not difficult, the result does not appear to be contained in the literature. We also present several consequences of the Bochner theorems which are interesting from the optimization perspective and for our purposes.

9.1 Bochner's theorem for continuous functions

There are two parts to Bochner's theorem; the first simply asserts that the Fourier coefficients of a positive type function are positive semidefinite matrices. The second part says that the Fourier series converges uniformly. We will prove the first part, and then after establishing some technical results, we will use the first part to give a proof of the second.

Theorem 9.1 (Bochner's theorem for compact groups)**.** *Suppose $\phi \in C(\Gamma)$. Then*

(a) *ϕ is of positive type if and only if $\widehat{\phi}(\pi)$ is a positive semidefinite matrix for every $\pi \in \widehat{\Gamma}$; and*

(b) *if ϕ is of positive type, then*

$$\phi(x) = \sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}(\widehat{\phi}(\pi)\pi(x)),$$

where the right-hand side converges absolutely as a series in $C(\Gamma)$ with the supremum-norm, and therefore also uniformly in x.

Proof of part (a). First suppose ϕ is a continuous function of positive type. Given $\pi \in \widehat{\Gamma}$, and a vector $v \in \mathcal{H}_\pi$, we have

$$\begin{aligned} \langle \widehat{\phi}(\pi)v, v\rangle &= \left\langle \left(\int \phi(x)\pi(x)^* \, d\lambda(x) \right) v, v \right\rangle \\ &= \int \phi(x)\langle \pi(x)^* v, v\rangle \, d\lambda(x), \end{aligned} \tag{9.1}$$

which is the integral of the pointwise product of two positive type functions by Proposition 8.6. The integral of a positive type function ψ is always nonnegative, since $\int \psi \, d\lambda = \int \psi(\mathbb{1}^* * \mathbb{1}) \, d\lambda \geq 0$, where $\mathbb{1}$ is the function identically equal to 1. The last integral in (9.1) is therefore nonnegative by Corollary 8.9, and this shows that $\widehat{\phi}(\pi)$ is positive semidefinite.

Next suppose that ϕ is a continuous function and that $\widehat{\phi}(\pi)$ is positive semidefinite for every $\pi \in \widehat{\Gamma}$. By equation (8.1) and Proposition 2.2, for every $g \in C(\Gamma)$ we have

$$\int (g^* * g)\phi \, d\lambda = \sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}(B_\pi \widehat{\phi}(\pi)) \geq 0,$$

where each $B_\pi = \widehat{g^* * g}(\pi) = \widehat{g}(\pi)\widehat{g}^*(\pi)$ is positive semidefinite. □

We now establish the tools we require for part (b).

Proposition 9.2. *If $\phi : \Gamma \to \mathbb{C}$ is a continuous function of positive type, then*

$$0 \leq \sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}(\widehat{\phi}(\pi)) < \infty.$$

Proof. Let $\mathcal{F}$ be the collection of finite subsets of $\widehat{\Gamma}$. For each $F \in \mathcal{F}$, let $\phi_F(x) = \sum_{\pi \in F} d_\pi \operatorname{Tr}(\widehat{\phi}(\pi)\pi(x))$. Then $\phi - \phi_F$ and also ϕ_F are of positive type by Theorem 9.1(a), and so Corollary 8.7 implies $\phi(e) \geq \phi_F(e) \geq 0$. But since $\phi_F(e) = \sum_{\pi \in F} d_\pi \operatorname{Tr}(\widehat{\phi_F}(\pi))$, it follows that the set

$$\left\{ \sum_{\pi \in F} d_\pi \operatorname{Tr}(\widehat{\phi_F}(\pi)) : F \in \mathcal{F} \right\}$$

is bounded from above. The conclusion now follows by noting that each term $d_\pi \operatorname{Tr}(\widehat{\phi_F}(\pi))$ is nonnegative, since each $\widehat{\phi_F}(\pi)$ is a positive semidefinite matrix. □

Lemma 9.3. *Let $A \in \mathbb{C}^{n\times n}$ be a positive semidefinite matrix, and let $U \in \mathbb{C}^{n\times n}$ be a unitary matrix. Then $|\operatorname{Tr}(AU)| \le \operatorname{Tr}(A)$.*

Proof. Let B be an $n \times n$ matrix such that $A = B^*B$. Then by the Cauchy-Schwartz inequality

$$|\operatorname{Tr}(AU)| = |\operatorname{Tr}(B^*BU)| \le \sqrt{\langle UB, UB\rangle}\sqrt{\langle B, B\rangle} = \langle B, B\rangle = \operatorname{Tr}(A),$$

where the inner product used here is the trace inner product. □

Theorem 9.1(b) now follows easily.

Proof of part (b). Since $\pi(x)$ is unitary, we have $|\operatorname{Tr}(\widehat{\phi}(\pi)\pi(x))| \le \operatorname{Tr}(\widehat{\phi}(\pi))$ for all $x \in \Gamma$ by Lemma 9.3 and Theorem 9.1(a). From Proposition 9.2 it then follows that the series $\sum_{\pi\in\widehat{\Gamma}} d_\pi \operatorname{Tr}(\widehat{\phi}(\pi)\pi(\cdot))$ converges absolutely in the supremum norm, and hence it converges uniformly. □

The next proposition with its corollary will be helpful to us in proving that the cone of continuous positive type functions and the cone of positive type measures are mutually dual.

Proposition 9.4. *If ϕ is any continuous function of positive type, then $\phi = g^* * g$ for some $g \in L^2(\Gamma)$.*

Proof. For each $\pi \in \widehat{\Gamma}$, let B_π be a $d_\pi \times d_\pi$ matrix satisfying $B_\pi^* B_\pi = \widehat{\phi}(\pi)$; this is possible by Theorem 9.1. By Proposition 8.3 and Theorem 8.2, it is enough to show that the series

$$\sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}(B_\pi \pi(x)) \tag{9.2}$$

converges in L^2. Let $\pi_1, \pi_2, \dots$ be an enumeration of the elements of $\widehat{\Gamma}$.[1]

Then the series (9.2) is equal to

$$\sum_{i=1}^{\infty} d_{\pi_i} \operatorname{Tr}(B_{\pi_i} \pi_i(x)),$$

which converges if and only if it is Cauchy. But for $M \geq m \geq 1$, we have by Parseval's identity (Proposition 8.3) that

$$\left\| \sum_{i=m}^{M} d_{\pi_i} \operatorname{Tr}(B_{\pi_i} \pi_i(x)) \right\|_2^2 = \sum_{i=m}^{M} d_{\pi_i} \operatorname{Tr}(B_{\pi_i}^* B_{\pi_i}) = \sum_{i=m}^{M} d_{\pi_i} \operatorname{Tr}(\widehat{\phi}(\pi_i)),$$

and by Proposition 9.2 the latter tends to 0 as $m, M \to \infty$. □

Corollary 9.5. *A measure $\nu \in \mathcal{M}(\Gamma)$ is of positive type if and only if $\int \phi \, d\nu \geq 0$ for all continuous functions $\phi : \Gamma \to \mathbb{C}$ of positive type.*

---

[1] Such an enumeration exists for the following reason: Given an orthonormal basis $\{e_{\pi,i}\}_i$ for each Hilbert space $\mathcal{H}_\pi$, $(\pi \in \widehat{\Gamma})$, the Peter-Weyl theorem (Theorem 5.12 [Fol95]) says that the functions $x \mapsto \sqrt{d_\pi}\langle \pi(x) e_i, e_j \rangle$ form an orthonormal basis B for $L^2(\Gamma)$. Since Γ is metrizable, it is second countable, and therefore $L^2(\Gamma)$ is separable. This implies that B is countable, and hence $\widehat{\Gamma}$ is countable.

Proof. One implication follows immediately from the definition of positive type measures. For the other, suppose ν is of positive type. If ϕ is a continuous function of positive type, then apply Proposition 9.4 to write $\phi = f^* * f$ with $f \in L^2(\Gamma)$. Let $\{f_n\}$ be a sequence of continuous functions converging to f in L^2. Then $f_n^* * f_n \to f^* * f$ pointwise as $n \to \infty$: Indeed, we have

$$\begin{aligned} &|f^* * f(x) - f_n^* * f_n(x)| \\ &\leq |\langle R_x f, f\rangle - \langle R_x f_n, f_n\rangle| \\ &\leq |\langle R_x f, f\rangle - \langle R_x f, f_n\rangle| + |\langle R_x f, f_n\rangle - \langle R_x f_n, f_n\rangle| \\ &\leq \|f\|_2 \|f - f_n\|_2 + \|f - f_n\|_2 \|f_n\|_2 \end{aligned}$$

for each $x \in \Gamma$. Therefore $\int \phi \, d\nu = \lim_{n\to\infty} \int f_n^* * f_n \, d\nu \geq 0$ by the theorem of Dominated Convergence. □

Lastly we prove a result for measures which is analogous to Corollary 8.7 (b).

Proposition 9.6. *If $\mu \in \mathcal{M}(\Gamma)$ is of positive type, then $\mu(E^{-1}) = \overline{\mu(E)}$, for all Borel subsets E of Γ, where we use the notation $E^{-1} = \{x^{-1} : x \in E\}$.*

Proof. It suffices to show that

$$\int \phi(x^{-1}) \, d\mu(x) = \int \phi(x) \, d\overline{\mu}(x) \tag{9.3}$$

for all $\phi \in C(\Gamma)$. First suppose ϕ is of positive type. Then $\overline{\phi}$ is also of positive type, and so Corollary 8.7 (b) and Corolary 9.5 together imply (9.3). The general case now follows from [Fol95, Proposition 3.33], which says that the linear span of the positive type functions is uniformly dense in the space of continuous functions. □

9.2 Bochner's theorem for measures

The purpose of this section is to prove a version of Bochner's theorem for measures of positive type on a compact group. After proving the main theorem, we give some consequences which display an analogy between measures of positive type and positive semidefinite matrices.

Theorem 9.7 (Bochner's theorem for measures)**.** *Suppose $\nu \in \mathcal{M}(\Gamma)$. Then ν is of positive type if and only if $\widehat{\nu}(\pi) := \int \pi(x)^* \, d\nu(x)$ is positive semidefinite for every $\pi \in \widehat{\Gamma}$.*

Proof. First suppose ν is of positive type. Given $\pi \in \widehat{\Gamma}$ and $v \in \mathcal{H}_\pi$, we have

$$\langle \widehat{\nu}(\pi) v, v \rangle = \Big\langle \Big(\int \pi(x)^* \, d\nu(x) \Big) v, v \Big\rangle = \int \langle \pi(x)^* v, v \rangle \, d\nu(x),$$

which is nonnegative by Proposition 8.6 and Corollary 9.5. This shows that the operator $\widehat{\nu}(\pi)$ is positive semidefinite.

For the other direction, assume $\widehat{\nu}(\pi)$ is positive semidefinite for every $\pi \in \widehat{\Gamma}$, and let $\phi \in C(\Gamma)$ be of positive type. Then by Theorem 9.1,

$$\phi(x) = \sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}(\widehat{\phi}(\pi)\pi(x)),$$

where the sum converges uniformly in x, and where each $\widehat{\phi}(\pi)$ is positive semidefinite. We have

$$\int \phi(x)\ d\nu(x) = \sum_{\pi\in\widehat{\Gamma}} d_\pi \operatorname{Tr}\left(\widehat{\phi}(\pi) \int \pi(x)\ d\nu(x)\right),$$

and the operators $\int \pi(x)\ d\nu(x)$ are positive semidefinite, being the complex conjugates of the positive semidefinite operators $\widehat{\nu}(\overline{\pi})$. Since the trace inner product of two positive semidefinite operators is nonnegative, so is the above sum. □

As an immediate consequence of Theorem 9.7 and Proposition 8.10, we get the following:

Corollary 9.8. *For every $\mu \in \mathcal{M}(\Gamma)$, the measure $\mu^* * \mu$ is of positive type.*

For the next result, we recall the notion of approximate identity. This is treated in [Fol95, Section 2.5]. Let $\mathcal{U}$ be a neighbourhood base at e in Γ, and for each $U \in \mathcal{U}$, let ψ_U be a continuous function such that $\operatorname{supp}(\psi_U)$ is contained in U, $\psi_U \geq 0$, $\psi_U(x) = \psi_U(x^{-1})$ for all $x \in \Gamma$, and $\int \psi_U = 1$. Then the family $\{\psi_U\}$ is called an *approximate identity.* Approximate identities exist for every locally compact group, and they satisfy the properties

1. $\|f * \psi_U - f\|_p \to 0$ as $U \to \{e\}$ if $1 \leq p < \infty$ and $f \in L^p(\Gamma)$;
2. $\|\psi_U * f - f\|_p \to 0$ as $U \to \{e\}$ if $1 \leq p < \infty$ and $f \in L^p(\Gamma)$; and

3. If $f \in C(\Gamma)$ then $\psi_U * f \to f$ and $f * \psi_U \to f$ in the supremum norm as $U \to \{e\}$.

Above we use the notation $U \to \{e\}$ to mean that the neighbourhood U tends to e in the directed set of neighbourhoods of e, ordered by reverse inclusion. For instance, item 1 says that for every $\varepsilon > 0$, there is a neighbourhood V of e such that $\|f * \psi_U - f\|_p < \varepsilon$ whenever $U \subset V$.

Lemma 9.9. *Let $\{\psi_U\}$ be an approximate identity. Then $\widehat{\psi_U}(\pi)$ converges to the $d_\pi \times d_\pi$ identity matrix for each $\pi \in \widehat{\Gamma}$ as $U \to \{e\}$.*

Proof. Fix $\pi_0 \in \widehat{\Gamma}$ and let $g_{\pi_0}(x) = d_{\pi_0} \operatorname{Tr}(\pi_0(x))$. Letting I denote the $d_{\pi_0} \times d_{\pi_0}$ identity matrix, we have

$$\widehat{g_{\pi_0}}(\pi) = \begin{cases} I & \text{if } \pi = \pi_0 \\ 0 & \text{otherwise} \end{cases}.$$

Then $\|\psi_U * g_{\pi_0} - g_{\pi_0}\|_2^2 \to 0$ as $U \to \{e\}$. But by Parseval's identity we have

$$\begin{aligned} \|\psi_U * g_{\pi_0} - g_{\pi_0}\|_2^2 &= \sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}[(\widehat{\psi_U * g_{\pi_0}}(\pi) - \widehat{g_{\pi_0}}(\pi))^*(\widehat{\psi_U * g_{\pi_0}}(\pi) - \widehat{g_{\pi_0}}(\pi))] \\ &= \sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}[(\widehat{g_{\pi_0}}(\pi)\widehat{\psi_U}(\pi) - \widehat{g_{\pi_0}}(\pi))^*(\widehat{g_{\pi_0}}(\pi)\widehat{\psi_U}(\pi) - \widehat{g_{\pi_0}}(\pi))] \\ &= d_{\pi_0} \operatorname{Tr}[(\widehat{\psi_U}(\pi_0)^* - I)(\widehat{\psi_U}(\pi_0) - I)]. \end{aligned}$$

The latter is d_{π_0} times the squared Frobenius norm of the matrix $\widehat{\psi_U}(\pi_0) - I$, and the conclusion follows. □

Recall that the *support* of a Borel measure $\nu \in \mathcal{M}(\Gamma)$ is defined as the set

$$\operatorname{supp}(\nu) = \{x \in \Gamma : |\nu|(U) > 0 \text{ for every open neighbourhood } U \text{ of } x\},$$

where $|\nu|$ denotes the total variation of ν. Note that the support of a Borel measure is always a closed set.

Corollary 9.10. *Suppose that $\nu \in \mathcal{M}(\Gamma)$ is a nonzero measure of positive type. Then $e \in \operatorname{supp}(\nu)$, where e denotes the identity element of Γ.*

Proof. Let $\{\psi_U\}$ be an approximate identity and suppose $e \notin \operatorname{supp}(\nu)$. Let V be an open, symmetric unit neighbourhood such that $V^2 \subset \Gamma \setminus \operatorname{supp}(\nu)$, and set $\phi_V = \psi_V^* * \psi_V$. Then ϕ_V is a continuous function of positive type by Proposition 8.5, and $\operatorname{supp}(\phi_V) \subset V^2$. Therefore $\int \phi_V(x)\, d\nu(x) = 0$. On the other hand, by Theorem 9.1, we have

$$\phi_V(x) = \sum_{\pi \in \widehat{\Gamma}} \operatorname{Tr}(\widehat{\phi_V}(\pi)\ \pi(x))$$

where the sum converges uniformly. Therefore,

$$0 = \int \phi_V(x)\, d\nu(x) = \sum_{\pi \in \widehat{\Gamma}} \operatorname{Tr}\left(\widehat{\phi_V}(\pi) \int \pi(x)\, d\nu(x)\right).$$

Each term in the sum on the right is nonnegative by the same argument used in the proof of Theorem 9.7, and therefore each one is zero. But $\widehat{\phi_V}(\pi)$ converges to the identity matrix as $V \to \{e\}$ by Lemma 9.9, and it follows that $\widehat{\nu}(\pi) = 0$ for every $\pi \in \widehat{\Gamma}$. □

For every $n \times n$ symmetric matrix A over $\mathbb{R}$, there exists an $r \geq 0$ such that $rI_{n\times n} + A$ is positive semidefinite. The next corollary shows that something analogous is true for signed symmetric measures.

Corollary 9.11. *Let $\nu \in \mathcal{M}(\Gamma)$ be a signed (real) measure which is symmetric; that is, $\nu(E) = \nu(E^{-1})$ for each Borel subset $E \subset \Gamma$. Then for each sufficiently large $r > 0$, the measure $r\delta_e + \nu$ is of positive type.*

Proof. Notice first that the conditions on ν imply that $\widehat{\nu}(\pi)$ is Hermitian for every $\pi \in \widehat{\Gamma}$. We have $|\langle \widehat{\nu}(\pi)v, v\rangle| \leq \|\nu\|$ for every unit vector $v \in \mathcal{H}_\pi$, since

$$|\langle \widehat{\nu}(\pi)v, v\rangle| = \left| \int \langle \pi(x)^* v, v\rangle \; d\nu(x) \right|, \tag{9.4}$$

and $|\langle \pi(x)^* v, v\rangle| \leq 1$ by the Cauchy-Schwartz inequality. Therefore if $r \geq \|\nu\|$, the matrix

$$(r\delta_e + \nu)\widehat{\ }(\pi) = rI_{d_\pi \times d_\pi} + \widehat{\nu}(\pi)$$

is positive semidefinite. Theorem 9.7 now implies that $r\delta_e + \nu$ is of positive type. □

Chapter 10

Positive type functions and measures as cones

In this chapter we show how the cone of positive type functions and the cone of positive type measures can be seen as dual to one another in a sense analogous to the self-duality of the cone of positive semidefinite matrices. Section 10.1 reviews the theory of cones and duality in topological vector spaces, and Section 10.2 then proves the main result of the chapter, which says that the cone of real-valued continuous positive type functions and the cone of finite signed positive type Radon measures are each other's duals.

10.1 Cones and duality

In this section, we review cones and duality in general. Most of the material comes from Chapters 2 and 8 of [Ali07].

10.1.1 Dual pairs and weak topologies

Given any dual pair, one defines the *weak topology* on V to be the topology τ generated by the seminorms $p_{v'} : V \to \mathbb{R}$, defined by

$$p_{v'}(v) = |\langle v, v' \rangle|, \quad v \in V,$$

for each $v' \in V'$. The topology τ is denoted $\sigma(V, V')$. A net $\{v_\alpha\}$ in V converges to $v \in V$ in τ if and only if $\langle v_\alpha, v' \rangle \to \langle v, v' \rangle$ for every fixed $v' \in V'$. The topology τ is locally convex and Hausdorff. All linear functionals on V which are continuous in τ have the form $v \mapsto \langle v, v' \rangle$ for some $v' \in V'$; this is a consequence of the Mackey-Arens theorem (Theorem 8.14 in [Ali07]). The weak topology on V' is defined by interchanging the roles of V and V' in the above, and the entire discussion is symmetric.

10.1.2 A separation theorem

Recall that two nonempty subsets X, Y of a vector space V are *strongly separated* by a nonzero linear functional $f : V \to \mathbb{R}$ if there exist $\alpha, \beta \in \mathbb{R}$ such that

$$f(x) \leq \alpha < \beta \leq f(y)$$

for all $x \in X$ and $y \in Y$. There exist many important separation theorems. However we will need just one (Corollary 8.18 in [Ali07]), which is a consequence of the Hahn-Banach theorem:

Theorem 10.1 (Separation of points from closed convex sets)**.** *Every point that lies outside a closed convex subset of a locally convex space can be strongly separated from the set by a nonzero continuous linear functional.*

10.1.3 Cones and their duals

Let $(V, V', \langle , \rangle)$ be a dual pair of $\mathbb{R}$-vector spaces and let $W \subset V$ be a cone. The *dual cone* of W is defined as

$$W^* := \{x' \in V' : \langle x, x' \rangle \geq 0 \text{ for all } x \in W\}.$$

The dual cone W^* is pointed precisely when $W - W$ is $\sigma(V, V')$-dense in V. The dual cone W^* is $\sigma(V', V)$-closed, and the second dual cone W^{**} is equal to the $\sigma(V, V')$-closure of W in V. In particular $W = W^{**}$ if and only if W is $\sigma(V, V')$-closed. The conic dual operation is inclusion-reversing: $W \subset W'$ implies $W^* \supset (W')^*$. See Section 2.2 in [Ali07] for a thorough treatment of conic duality in topological vector spaces.

10.2 Positive type functions and measures as mutually dual cones

We now proceed to prove the main new result of this chapter, Theorem 10.2. As in previous chapters, let Γ be a compact group. Let $C_{\mathbb{R}}(\Gamma)$ be the $\mathbb{R}$-vector space of real-valued continuous functions on Γ, and let $\mathcal{M}_{\mathbb{R}}(\Gamma)$ be the $\mathbb{R}$-vector space of signed Borel measures on Γ having finite total variation. We let $S_{\mathbb{R}}(\Gamma)$ be the set of symmetric functions in $C_{\mathbb{R}}(\Gamma)$,

$$S_{\mathbb{R}}(\Gamma) := \{f \in C_{\mathbb{R}}(\Gamma) : f(x^{-1}) = f(x) \text{ for all } x \in \Gamma\},$$

and we let $\Sigma_{\mathbb{R}}(\Gamma)$ denote the set of symmetric measures in $\mathcal{M}(\Gamma)$,

$$\Sigma_{\mathbb{R}}(\Gamma) := \{\nu \in \mathcal{M}_{\mathbb{R}}(\Gamma) : \nu(E) = \nu(E^{-1}) \text{ for all Borel subsets } E \text{ of } \Gamma\}.$$

In the terminology of Section 10.1, we define the duality $\langle , \rangle : S_{\mathbb{R}}(\Gamma) \times \Sigma_{\mathbb{R}}(\Gamma) \to \mathbb{R}$ by $\langle f, \nu \rangle = \int f \; d\nu$. This duality induces the usual weak topology on $S_{\mathbb{R}}(\Gamma)$ regarded as a subspace of $C(\Gamma)$, and the weak-$*$ topology on $\Sigma_{\mathbb{R}}(\Gamma)$. Let $\mathcal{P}$ denote the set of functions of positive type in $C_{\mathbb{R}}(\Gamma)$, and let $\mathcal{Q}$ be the set of measures of positive type in $\mathcal{M}_{\mathbb{R}}(\Gamma)$:

$$\mathcal{P} = \{\phi \in C_{\mathbb{R}}(\Gamma) : \phi \text{ is of positive type}\}$$
$$\mathcal{Q} = \{\nu \in \mathcal{M}_{\mathbb{R}}(\Gamma) : \nu \text{ is of positive type}\}.$$

We have $\mathcal{P} \subset S_{\mathbb{R}}(\Gamma)$ and $\mathcal{Q} \subset \Sigma_{\mathbb{R}}(\Gamma)$ by Corollary 8.7 (b) and Proposition 9.6. Both $\mathcal{P}$ and $\mathcal{Q}$ are closed cones in their respective topologies. Our first result says that $\mathcal{P}$ and $\mathcal{Q}$ are dual to one another, and that, in a sense, $\mathcal{P}$ is dense in $\mathcal{Q}$.

Theorem 10.2. *Let*

$$\widetilde{\mathcal{P}} = \{\phi \; d\lambda : \phi \in \mathcal{P}\}$$

be the image of $\mathcal{P}$ under the map $f \mapsto f\ d\lambda$ from $C_{\mathbb{R}}(\Gamma)$ into $\mathcal{M}_{\mathbb{R}}(\Gamma)$, and let $\mathcal{D}$ denote its weak-$*$ closure in $\mathcal{M}_{\mathbb{R}}(\Gamma)$. Let

$$\mathcal{P}^* = \left\{\mu \in \Sigma_{\mathbb{R}}(\Gamma) : \int \phi\ d\mu \geq 0 \textit{ for all } \phi \in \mathcal{P}\right\}$$

be the dual cone of $\mathcal{P}$. Then $\mathcal{D} = \mathcal{P}^* = \mathcal{Q}$.

Proof. We first prove $\mathcal{D} = \mathcal{P}^*$. For the inclusion $\mathcal{D} \subset \mathcal{P}^*$, it suffices to prove $\tilde{\mathcal{P}} \subset \mathcal{P}^*$ since $\mathcal{P}^*$ is closed. But if $k, \phi \in \mathcal{P}$, then $\int k\phi\ d\lambda \geq 0$ by Corollary 8.9. Now suppose there exists $\nu \in \mathcal{P}^* \setminus \mathcal{D}$. By Theorem 10.1, there exists $\phi \in S_{\mathbb{R}}(\Gamma)$ and $\beta \in \mathbb{R}$ such that

$$\int \phi\ d\mu \geq \beta \text{ for all } \mu \in \mathcal{D}, \text{ but} \tag{10.1}$$

$$\int \phi\ d\nu < \beta. \tag{10.2}$$

We cannot have $\int \phi\ d\mu_0 < 0$ for any $\mu_0 \in \mathcal{D}$, because $\mathcal{D}$ is a cone and we would have by (10.1) that $t \int \phi\ d\mu_0 \geq \beta$ for all $t \geq 0$. Therefore $\int \phi\ d\mu \geq 0$ for all $\mu \in \mathcal{D}$. It now follows that $\phi \in \mathcal{P}$, for if ψ is any (possibly $\mathbb{C}$-valued) function of positive type, then

$$\begin{aligned} 0 \leq \int \phi(x)(\psi(x) + \overline{\psi}(x))\ d\lambda(x) &= \int \phi(x)\psi(x)\ d\lambda(x) + \int \phi(x^{-1})\psi(x)\ d\lambda(x) \\ &= 2\int \phi(x)\psi(x)\ d\lambda(x), \end{aligned}$$

by Corollary 8.7 (b). Now, using $\mu = 0$ in (10.1) we obtain $0 \geq \beta$, so (10.2) becomes $\int \phi\ d\nu < 0$, which contradicts $\nu \in \mathcal{P}^*$. This completes the proof that $\mathcal{D} = \mathcal{P}^*$.

The inclusion $\mathcal{Q} \subset \mathcal{P}^*$ is just Corollary 9.5. For the inclusion $\mathcal{P}^* \subset \mathcal{Q}$, we need to show that $\mu \in \Sigma_{\mathbb{R}}(\Gamma)$ is of positive type whenever it satisfies $\int \psi \, d\mu \geq 0$ for all $\mathbb{R}$-valued continuous functions ψ of positive type. But for arbitrary ψ of positive type, this follows from Corollary 9.5, the symmetry of μ, and the fact that $\psi + \psi^* \in \mathcal{P}$. □

We also record here two useful facts regarding $S_{\mathbb{R}}(\Gamma)$ and the cone $\mathcal{P}$.

Proposition 10.3. *The $\mathbb{R}$-linear span of $\mathcal{P}$ is sup-norm dense in $S_{\mathbb{R}}(\Gamma)$.*

Proof. Let $f, g \in S_{\mathbb{R}}(\Gamma)$ be given. Then the above linear span contains the function

$$\frac{1}{4}[(f+g) * (f+g)^* - (f-g) * (f-g)^*] = \frac{1}{2}[f * g^* + g * f^*].$$

Taking g to run over an approximate identity proves that $\frac{1}{2}(f+f^*) = f$ belongs to the sup-norm closure of the span. □

Proposition 10.4. *If $f \in S_{\mathbb{R}}(\Gamma)$, then $\widehat{f}(\pi)$ is Hermitian for each $\pi \in \widehat{\Gamma}$.*

Proof. Fix $\pi \in \widehat{\Gamma}$. For all $u, v \in \mathcal{H}_\pi$ we have

$$\begin{aligned} \langle \widehat{f}(\pi)u, v\rangle &= \int \langle \pi(x)^* u, v\rangle f(x)\, dx = \int \langle \pi(x)^* u, v\rangle f(x^{-1})\, dx \\ &= \int \langle \pi(x) u, v\rangle f(x)\, dx = \int \langle u, \pi(x)^* v\rangle f(x)\, dx \\ &= \langle u, \widehat{f}(\pi) v\rangle. \end{aligned}$$

□

10.3 Bases and interior points

Understanding whether a cone has interior points and a compact base is often helpful in convex optimization when one wants to know whether duality gaps exist, and whether optimal values are attained. In this section, we ask and answer these questions for $\mathcal{P}$ and $\mathcal{Q}$.

10.3.1 Bases

We begin by recalling a definition. Let V will be an $\mathbb{R}$-vector space. If $K \subset V$ is a cone, then a nonempty convex subset $B \subset K \setminus \{0\}$ is called a *base* for K if for each $x \in K \setminus \{0\}$, there exists a unique $t > 0$ and a unique $b \in B$ for which $x = tb$.

Unfortunately, the cone $\mathcal{P}$ need not have a $\sigma(S_{\mathbb{R}}(\Gamma), \Sigma_{\mathbb{R}}(\Gamma))$-compact base, and $\mathcal{Q}$ need not have a $\sigma(\Sigma_{\mathbb{R}}(\Gamma), S_{\mathbb{R}}(\Gamma))$-compact base. We now demonstrate this with an example.

Let $\mathbb{T}$ be the circle group and denote its Haar measure by λ. Suppose B is a compact base for $\mathcal{P}$. Given an finite subset F of the character group $\widehat{\mathbb{T}}$, let

$$N_F = \{\phi \in B : \int \phi(x)\chi(x)\ d\lambda(x) = 0 \text{ for all } \chi \in F\}.$$

Then each N_F is a weakly closed subset of B, and is nonempty since for any F, we have $\chi \in \mathcal{P} \setminus N_F$ whenever $\chi \notin F$, and $\chi = tb$ for some $t > 0$, $b \in B$. The family $\{N_F\}$ therefore has the finite intersection property; that is, all the finite intersections are nonempty. However the intersection $\cap N_F$ over all finite subsets F is empty since $0 \notin B$.

Therefore B cannot be compact. The proof that $\mathcal{Q}$ does not have a compact base is essentially the same.

It may however be of interest that $\mathcal{P}$ and $\mathcal{Q}$ usually both have closed bases. This can be shown with the help of Theorem 1.47 of [Ali07], which characterizes bases of cones as the positive level sets of *strictly positive linear functionals*, which are linear functionals $f : V \to \mathbb{R}$ satisfying $f(x) > 0$ for all nonzero x in the cone of V.

First $\mathcal{P}$: It follows from Corollary 8.7 (b) that a strictly positive linear functional on $S_{\mathbb{R}}(\Gamma)$ with cone $\mathcal{P}$ is given simply by δ_e, namely $f \mapsto f(e)$. So a natural base for $\mathcal{P}$ is just

$$\{\phi \in \mathcal{P} : \phi(e) = 1\},$$

which is clearly closed in the weak topology.

Constructing a base for $\mathcal{Q}$ is trickier, but possible with the help of Bochner's theorem using the hypothesis that Γ is second countable. In this case $L^2(\Gamma)$ is separable, from which it follows that $\widehat{\Gamma}$ is countable. We may therefore let $\{\pi_n\}$ be an enumeration of $\widehat{\Gamma}$, and set

$$\phi(x) = \sum_{n=1}^{\infty} \frac{1}{d_{\pi_n} 2^n} \operatorname{Tr}(\pi_n(x)).$$

The series converges absolutely to a function in $C(\Gamma)$. Now if $\mu \in \mathcal{Q}$ is nonzero, we have

$$\int \phi \; d\mu = \sum_{n=1}^{\infty} \frac{1}{d_{\pi_n} 2^n} \operatorname{Tr}\left(\int \pi_n(x) \; d\mu(x)\right) = \sum_{n=1}^{\infty} \frac{1}{d_{\pi_n} 2^n} \operatorname{Tr}(\widehat{\overline{\mu}}(\pi_n)) > 0,$$

by Theorem 9.7. We also have $\int \overline{\phi}\ d\mu > 0$, and therefore $\phi + \overline{\phi} \in S_{\mathbb{R}}(\Gamma)$ is a function which defines a strictly positive, continuous linear functional on $\mathcal{Q}$.

10.3.2 Interior points

Also unfortunate from the convex optimization perspective is that $\mathcal{P}$ and $\mathcal{Q}$ need not have weakly interior points. This can be proven using the characterization for interior points given in Lemma 2.5 of [Ali07]. We write $f \geq_{\mathcal{P}} g$ to mean that $f - g \in \mathcal{P}$.

Suppose $\phi \in \mathcal{P}$ is an interior point. Then

$$U = \{f \in S_{\mathbb{R}}(\Gamma) : -\phi \leq_{\mathcal{P}} f \leq_{\mathcal{P}} \phi\}$$

is a neighbourhood of 0 by Lemma 2.5 of [Ali07]. Therefore U contains a basic open neighbourhood of 0,

$$U_1 = \{f \in S_{\mathbb{R}}(\Gamma) : \left|\int f\ d\mu_1\right| < \varepsilon, \ldots \left|\int f\ d\mu_n\right| < \varepsilon\},$$

for some $\mu_1, \ldots, \mu_n \in \mathcal{M}_{\mathbb{R}}(\Gamma)$ and $\varepsilon > 0$. But if $S_{\mathbb{R}}(\Gamma)$ has infinite dimension, then U_1 contains a subspace of infinite dimension. Therefore there exists $f \in U_1$ with $f \neq 0$ and $tf \in U_1 \subset \mathcal{P}$ for all $t \in \mathbb{R}$, which is impossible since $\mathcal{P}$ is a cone. The argument for $\mathcal{Q}$ is essentially the same.

We note however that $\mathcal{Q}$ *does* have a norm interior point, namely δ_e. To see this, let $\mu \in \mathcal{M}_{\mathbb{R}}(\Gamma)$ be any measure with $\|\delta_e - \mu\| < 1$. We want to show $\mu \in \mathcal{Q}$. Let $\phi \in \mathcal{P}$ be arbitrary, and assume $\phi(e) = 1$.

Then $1 > |\int \phi \, d\mu - \int \phi \, d\delta_e| = |\int \phi \, d\mu - 1|$, which implies $\int \phi \, d\mu > 0$. That $\mu \in \mathcal{Q}$ now follows from Theorem 10.2.

Part V

Lovász ϑ duality for Cayley graphs over compact groups

Chapter 11

Introduction

11.1 Background and motivation

Eigenvalue techniques related to the Lovász ϑ-function have been applied to obtain upper bounds for spherical codes [DGS77] (including the kissing numbers), for indepedence numbers of compact packing graphs [dLV13], and for upper bounding densities of packings of congruent copies of a body in $\mathbb{R}^n$ [dOFV13]. They have also been used to upper bound measures of spherical sets avoiding a prescribed set of angles ([BNdOFV09], [BDdOFV14], [DP15]), and to write down infinite-dimensional semidefinite programs which upper bound measures of distance avoiding sets in compact metric spaces ([BNdOFV09]). Similar approaches were used in [dOFV10], [dOF09], and [BPT14] to upper bound upper densities of distance-avoiding sets in $\mathbb{R}^n$; these upper bounds were then turned into lower bounds for the measurable chromatic number of Euclidean n-space.

In the applications involving measures and upper densities, an infinite-dimensional linear or semidefinite relaxation in the spirit of ϑ is set up as a maximization problem whose value upper bounds the quantity of interest. In the absence of methods for solving such infinite-dimensional optimization problems to optimality, one typically symmetrizes the problem ([DKPS07], [BGSV12]) to obtain a linear program, and then writes down the dual linear program. The dual has the property that any feasible solution provides an upper bound for the quantity of interest, so one then only needs to find good feasible solutions. In the coding and packing applications, the relaxation one obtains is already a minimization problem, and it is not immediately clear whether the dual maximization problem provides any interesting information about the original (unrelaxed) problem.

Given the success of these methods in recent years, it is therefore of interest to study the duality of these infinite-dimensional systems. How should one dualize them, and can one dualize *before* symmetrizing to a linear program? Does strong duality hold? Does the maximization dual in the coding and packing cases give any information about the original graphs? Answering these questions is the main objective of Part V of the thesis.

In Part V, we develop primal and dual formulations of the Lovász ϑ-function and the Schrijver ϑ'-function for two kinds of Cayley graphs over compact groups, which are defined below. Most of the definitions are quite natural; the main contributions in this part of the thesis are the strong duality proofs which use abstract harmonic analysis in an essential way. Additionally we show that our extensions satisfy most of the important properties of the usual ϑ- and ϑ'-functions by giving new analytic proofs.

An attempt to unify some of the results mentioned in the first paragraph was made in [BDdOFV14], but until now only de Laat and Vallentin [dLV13] have seriously investigated duality for any of these problems; their work was in the context of what they call topological packing graphs, and they give a delicate functional-analytic argument proving strong duality for an entire hierarchy of infinite-dimensional semidefinite programs, of which ϑ' for their graphs is only the first level. Their proof however does not extend to the other cases mentioned in the first paragraph, or to ϑ.

The main theoretical challenge in this part of the thesis was to develop a theory of the Lovász ϑ-function and Schrijver ϑ'-function for infinite graphs which applies to as many existing (and future) examples from the literature as possible, while still enabling one to prove strong duality and to recover the well-known properties of the ϑ-function which hold on finite graphs. This was made difficult by the fact that the ϑ-function examples from the literature were each developed ad hoc for a specific application, and as such they do not come from a common set of definitions.

The choice was made to develop a theory of the ϑ-function in the context of Cayley graphs over compact groups. This choice was taken because it seems to be the simplest framework which both captures the examples from the literature, and allows for the application of abstract harmonic analysis in the strong duality proofs. A bonus of this choice is the discovery that many of the well-known properties of the ϑ-function have new and short (almost trivial) analytic proofs.

The theory extends effortlessly to graphs on homogeneous spaces, and in this way our framework captures essentially all of the examples mentioned in the first paragraph – the choice to work over Cayley graphs rather than some more general class of graphs on homogeneous spaces was basically taken to make the proofs cleaner. Furthermore, the sorts of Cayley graphs on $\mathbb{R}^n$ of interest to us can be understood by approximating them with Cayley graphs on large n-dimenensional tori; for this reason the compactness assumption is no big restriction either. However, even if some infinite graphs eventually come along whose ϑ-functions do not fit into our theory, it is our belief that the proof techniques given here could extend to accommodate them.

11.2 Preliminaries

In Part V we assume the reader possesses a mastery of the basics of abstract harmonic analysis on compact nonabelian groups up to and including the Peter-Weyl theorem. The books [Fol95] and [DE09] give clear and readable expositions of all the prerequisite material.

Let Γ be a compact group with Borel σ-algebra $\mathcal{B}$, Haar measure λ, and identity element e. Since Γ is compact, its Haar measure is finite, and we will always assume the normalization $\lambda(\Gamma) = 1$. A subset

$X \subset \Gamma$ with $e \notin X$ will be called a *connection set* if $X = X^{-1}$; that is $x^{-1} \in X$ whenever $x \in X$. For any connection set X, the *Cayley graph* $\mathrm{Cay}(\Gamma, X)$ is defined as the graph $\mathrm{Cay}(\Gamma, X) = (V, E)$ with $V = \Gamma$, and $E = \{\{x, y\} \subset V : y^{-1}x \in X\}$. The conditions defining the connection set guarantee that $\mathrm{Cay}(\Gamma, X)$ is undirected and without self-loops. Notice that by the definition of E, for every $a \in \Gamma$, left multiplication by a is a graph automorphism of $\mathrm{Cay}(\Gamma, X)$.

In addition to the independence number, we will also be interested in the *independence ratio* of $G = \mathrm{Cay}(\Gamma, X)$, defined as

$$\tilde{\alpha}(G) := \sup\{\lambda(I) : I \in \mathcal{B} \text{ is independent in } G\}.$$

Theorem 4.1 shows that the supremum in the definition of $\tilde{\alpha}$ should not be replaced with a maximum. We further define $\chi_m(G)$, the *measurable chromatic number* of G; this is the smallest number k such that Γ can be partitioned into sets $C_1, \ldots, C_k \in \mathcal{B}$ such that each C_i is independent, or ∞ if no such finite partition exists.

In what follows, when $X \subset \Gamma$ is a connection set and $G = \mathrm{Cay}(\Gamma, X)$, we will use the notation $X^c = \Gamma \setminus (X \cup \{e\})$, and $G^c = \mathrm{Cay}(\Gamma, X^c)$; the graph G^c is the complementary graph of G. We will use $\overline{X}$ to mean the closure of X in the topology of Γ, and we define $\overline{G} = \mathrm{Cay}(\Gamma, \overline{X})$, which we call the *closure* of the graph G.

We will be mainly interested in two sorts of Cayley graphs on Γ, which we call respectively dense and sparse. Roughly speaking, dense Cayley graphs have sufficiently dense edge sets to force the independence number to be finite. When the graph is infinite, the chromatic number is then also infinite. An example of an infinite dense Cayley graph is the kissing number graph on the circle, obtained by setting $\Gamma = \mathbb{R}/\mathbb{Z}$ and $X = (-1/6, 0) \cup (0, 1/6)$. Sparse Cayley graphs are roughly speaking those graphs whose edge set is sparse enough to guarantee independent sets of positive Haar measure. All the circle graphs from Theorem 4.1 are sparse. It is important to note that our definitions of dense and sparse have nothing to do with those often found in the graph theory literature. For instance, our sparse graphs $\mathrm{Cay}(\Gamma, X)$ can have $\lambda(X) > 0$; in other words a positive proportion of the possible edges can be present.

The formal definitions are as follows. We say that a connection set $X \in \mathcal{B}$ is *sparse* if $e \notin \overline{X}$, and that $\mathrm{Cay}(\Gamma, X)$ is *sparse* if X is sparse. We call a connection set $X \in \mathcal{B}$ *dense* if e is an interior point of $\{e\} \cup X$, and we say $\mathrm{Cay}(\Gamma, X)$ is *dense* if X is dense.[1] Notice that $G = \mathrm{Cay}(\Gamma, X)$ is sparse if and only if G^c is dense. There exist Cayley graphs which are neither dense nor sparse, and a Cayley graph is both dense and sparse if and only if it is finite. We next give some justification to these definitions.

Proposition 11.1. *Let $G = \mathrm{Cay}(\Gamma, X)$. If G is sparse, then $\tilde{\alpha}(G) > 0$ and $\chi(G) \leq \chi_m(G) < \infty$. If G is dense, then $\alpha(G) < \infty$.*

Proof. First suppose G is sparse. Then there exists an open symmetric unit neighbourhood U with $U^2 \subset \Gamma \setminus \overline{X}$; then U is an independent set with $\lambda(U) > 0$, proving the first assertion. For the second assertion,

[1]Dense Cayley graphs are examples of the topological packing graphs of de Laat and Vallentin [dLV13].

the set $\{xU : x \in \Gamma\}$ is an open cover of Γ, so by compactness there exists a finite subcover, say $x_1U, \ldots, x_nU$. Since left multiplication by a group element is a graph automorphism, each x_iU is an independent set. With Γ now covered by n independent sets, one easily constructs a proper colouring with at most n colours. The inequality $\chi(G) \leq \chi_m(G)$ is obvious.

If G is dense then we take U^2 to be contained in the interior of $\{e\} \cup X$. Since U is a clique, each translate x_iU of U can contain at most one point from an independent set. This proves the last assertion. $\square$

Suppose that Y is a Borel subset of Γ. We regard Y as a topological space with the topology induced from Γ, and we define $\mathcal{B}(Y) := \{E \cap Y : E \in \mathcal{B}\}$. We use $C_{\mathbb{R}}(Y)$ to denote the set of continuous $\mathbb{R}$-valued functions on Y, and $\mathcal{M}_{\mathbb{R}}(Y)$ to denote the set of regular signed Borel measures on Y having finite total variation. The set of continuous $\mathbb{R}$-valued functions on Y vanishing at infinity will be denoted $C_{\mathbb{R},0}(Y)$; this is the set of functions $f \in C_{\mathbb{R}}(Y)$ such that for each $\varepsilon > 0$, the set $\{x \in Y : |f(x)| \geq \varepsilon\}$ is compact in Y.

As in Part IV, we use $\mathcal{P}$ and $\mathcal{Q}$ to denote, respectively, the cone of $\mathbb{R}$-valued continuous functions of positive type, and the cone of finite signed regular Borel measures of positive type.

Now suppose further that Y is symmetric; that is $y^{-1} \in Y$ whenever $y \in Y$. Then $S_{\mathbb{R}}(Y)$ will be the set of continuous $\mathbb{R}$-valued symmetric functions on Y:

$$S_{\mathbb{R}}(Y) := \{f \in C_{\mathbb{R}}(Y) : f(y^{-1}) = f(y) \text{ for all } y \in Y\}.$$

We will use $S_+(Y)$ to denote the set of all functions in $S_{\mathbb{R}}(Y)$ taking only nonnegative values. We define $\Sigma_{\mathbb{R}}(Y)$ to be the set of all $\mathbb{R}$-valued finite symmetric signed regular Borel measures on Y:

$$\Sigma_{\mathbb{R}}(Y) := \{\nu \in \mathcal{M}_{\mathbb{R}}(Y) : \nu(E) = \nu(E^{-1}) \text{ for all } E \in \mathcal{B}(Y)\}.$$

We will use $\Sigma_+(Y)$ to denote the set of positive measures in $\Sigma_{\mathbb{R}}(Y)$. By Corollary 8.7 (b) and Proposition 9.6, we have $\mathcal{P} \subset S_{\mathbb{R}}(\Gamma)$ and $\mathcal{Q} \subset \Sigma_{\mathbb{R}}(\Gamma)$.

For $x \in \Gamma$, the Dirac point mass at x will be denoted δ_x.

Chapter 12

Sparse Cayley graphs over compact groups

In this chapter we define extensions of the Lovász ϑ-function and Schrijver's ϑ'-function for sparse Cayley graphs over compact groups, which give upper bounds for $\tilde{\alpha}$. The extension is given in terms of two mutually dual conic optimization programs. We prove both weak and strong duality, and we relate these ϑ and ϑ' to their finite counterparts. We furthermore mention how the theory extends to include graphs on Γ-homogeneous spaces when Γ is a subgroup of the graph automorphisms. Lastly we demonstrate the machinery by working out a few examples (in particular forbidden distance graphs on the unit sphere and on $\mathbb{R}^n$) to show how they fit into our theory.

12.1 ϑ_s and ϑ'_s: primal formulation

In this section we give definitions of the ϑ- and ϑ'-functions of sparse Cayley graphs over compact groups and we prove that they upper bound the independence ratio. When Γ is a finite group we prove that the new definitions give $\vartheta(G)/|\Gamma|$ and $\vartheta'(G)/|\Gamma|$, respectively; the difference in normalization comes from the fact that we use the normalized Haar measure in our definitions. We then briefly explain how the theory extends to include graphs on Γ-homogeneous spaces when Γ is a compact subgroup of the automorphism group.

We demonstrate our formulation with an example by recovering the linear programming upper bounds for the independence ratio of forbidden distance graphs on the unit sphere appearing in [BNdOFV09] and [dOF09]. (The dual will be investigated in the next section.) Just before giving the example, we prove a useful fact clarifying the relationship between positive type functions on the special orthogonal group and positive definite functions on the sphere.

12.1.1 Definition and relation to ϑ-function for finite graphs

Suppose Γ is a compact group, X is a sparse connection set, and $G = \mathrm{Cay}(\Gamma, X)$. We define $\vartheta_s(G)$ as the value of the following conic optimization program.

$$\vartheta_s(G) := \sup\left\{\int \phi \; d\lambda : \phi(e) = 1, \phi|_X \equiv 0, \phi \in \mathcal{P}\right\}, \tag{12.1}$$

where $\phi|_X \equiv 0$ means that $\phi(x) = 0$ for all $x \in X$. We also define

$$\vartheta_s'(G) := \sup \left\{ \int \phi \, d\lambda : \phi(e) = 1, \phi|_X \equiv 0, \phi \in \mathcal{P} \cap S_+(\Gamma) \right\}. \quad (12.2)$$

Our first proposition gives the most important property of ϑ_s and ϑ_s'.

Theorem 12.1. *If G is a sparse Cayley graph, then programs* (12.1) *and* (12.2) *are both feasible and*

$$\vartheta_s(G) \geq \vartheta_s'(G) \geq \tilde{\alpha}(G).$$

Proof. Since G is sparse, there exists an independent set $I \in \mathcal{B}$ with $\lambda(I) > 0$ by Proposition 11.1. Let I be any such set, and let $\phi = \lambda(I)^{-1} \mathbb{1}_I^* * \mathbb{1}_I$. Then clearly $\phi \in \mathcal{P} \cap S_+(\Gamma)$. Also

$$\phi(e) = \lambda(I)^{-1} \int \mathbb{1}_I(y) \mathbb{1}_I(y) \, dy = 1,$$

and if $x \in X$, then

$$\phi(x) = \lambda(I)^{-1} \int \mathbb{1}_I(y)^* \mathbb{1}_I(y^{-1}x) \, dy = \lambda(I)^{-1} \int \mathbb{1}_I(y) \mathbb{1}_I(yx) \, dy = 0$$

for since I is independent, at most one of y and xy can belong to I. Therefore ϕ is feasible for (12.1) and (12.2). Finally, letting 1 denote the trivial representation of Γ, we have

$$\int \phi \, d\lambda = \widehat{\phi}(1) = \lambda(I)^{-1} (\widehat{\mathbb{1}_I}(1))^* \widehat{\mathbb{1}_I}(1) = \lambda(I).$$

Clearly $\vartheta_s(G) \geq \vartheta_s'(G)$, so the proposition follows. $\square$

We note here that $\vartheta_s(G) = \vartheta_s(\overline{G})$ because of the requirement in (12.1) and (12.2) that ϕ be continuous. This observation combined with Theorem 12.1 also shows that the sparseness of G is actually *equivalent* to the feasibility of (12.1) and (12.2), because of the constraint $\phi(e) = 1$.

The next proposition explains how our definitions of ϑ_s and ϑ'_s relate to the usual definitions for finite graphs.

Proposition 12.2. *Suppose* Γ *is a finite group with* $n = |\Gamma|$. *Then* $\vartheta_s(G) = \vartheta(G)/n$ *and* $\vartheta'_s(G) = \vartheta'(G)/n$.

Proof. Essentially the same proof is given in [DdLV14]. We prove only that $\vartheta_s(G) = \vartheta(G)/n$, since the proof of $\vartheta'_s(G) = \vartheta'(G)/n$ is an easy modification of this one. We first construct a feasible solution of program (12.1) from an optimal solution of program (2.6). Let $A = (a_{x,y})_{x,y\in\Gamma}$ be a solution for program (2.6) with $\vartheta(G) = \sum_{x,y\in\Gamma} a_{x,y}$, and for each $z \in \Gamma$, let $\phi(z) = \sum_{\{x,y\in\Gamma : y^{-1}x=z\}} a_{x,y}$. Then clearly $\phi(e) = 1$, and $\phi(x) = 0$ for all $x \in X$, and $\frac{1}{|\Gamma|}\sum_{g\in\Gamma}\phi(g) = \frac{1}{n}\sum_{x,y\in\Gamma} a_{x,y} = \vartheta(G)/n$. To see that ϕ is of positive type, let points $x_1, \ldots, x_n \in \Gamma$ be distinct points and let $v \in \mathbb{C}^n$ be given. Then

$$v^t(\phi(x_j^{-1}x_i))_{i,j=1}^n v = \sum_{i,j=1}^n \phi(x_j^{-1}x_i)\overline{v_i}v_j$$

$$= \sum_{i,j=1}^n \sum_{g\in\Gamma} a_{gx_j^{-1}x_i,g}\overline{v_i}v_j = \sum_{g\in\Gamma}\sum_{i,j=1}^n a_{gx_i,gx_j}\overline{v_i}v_j,$$

and the latter is nonnegative since $(a_{gx_i,gx_j})_{i,j=1}^n$ is a principal submatrix of the positive semidefinite matrix A.

For the other direction, let ϕ be a feasible solution for (12.1), and define $A = (a_{x,y})_{x,y\in\Gamma}$ by $a_{x,y} = \phi(y^{-1}x)/n$, $(x, y \in \Gamma)$. We have $\mathrm{Tr}(A) = 1$, and $a_{x,y} = 0$ whenever $y^{-1}x \in X$. Also A is positive semidefinite since for $v = (v_x)_{x\in\Gamma}$, we have

$$v^t A v = \sum_{x,y\in\Gamma} a_{x,y}\overline{v}_x v_y = \frac{1}{n}\sum_{x,y\in\Gamma} \phi(y^{-1}x)\overline{v}_x v_y \geq 0$$

since ϕ is of positive type. Finally $\sum_{x,y} a_{x,y} = \frac{1}{n}\sum_{x,y}\phi(y^{-1}x) = \sum_{g\in\Gamma}\phi(g)$. □

12.1.2 Frequency domain formulation

By taking the Fourier transform, one can find an expression for ϑ_s in the frequency domain. This can help compute it. This formulation was explored in [DdLV14] when Γ is a finite group.

Again suppose Γ is a compact group and that X is a sparse connection set, and let $G = \mathrm{Cay}(\Gamma, X)$. The following equalities are easily verified from (12.1), using Theorem 9.1 (Bochner's theorem). We write $A \succeq 0$ to mean that A is Hermitian positive semidefinite. We also use the notation $\widehat{\Gamma}$, d_π, $\mathcal{H}_\pi$ from Part IV; the trivial representation is denoted 1.

$$\vartheta_s(G) = \sup\Big\{A_1 : \sum_{\pi\in\widehat{\Gamma}} d_\pi \,\mathrm{Tr}(A_\pi) = 1,\ \sum_{\pi\in\widehat{\Gamma}} d_\pi \,\mathrm{Tr}(A_\pi \pi(x)) = 0 \text{ for } x \in X,$$
$$A_{\overline{\pi}} = \overline{A_\pi},\ A_\pi \succeq 0 \text{ for } \pi \in \widehat{\Gamma}\Big\}. \qquad (12.3)$$

The constraint $A_{\overline{\pi}} = \overline{A_\pi}$ (which does not appear in [DdLV14]) is required to ensure that $\sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}(A_\pi \pi(x))$ sums to an $\mathbb{R}$-valued function.

Therefore $\vartheta_s(G)$ can be expressed as the value of a block-diagonal semidefinite program, possibly with infinitely many finite blocks. If Γ is abelian then all the irreducible representations are one-dimensional, and one sees immediately that $\vartheta_s(G)$ can be expressed as the value of a *linear* program.

The same may hold true even when Γ is not abelian; a sufficient condition for (12.3) to reduce to a linear program is that X be closed under conjugation. One can see this using Schur's lemma (see e.g. [Fol95, Theorem 3.5]): Given a feasible solution $\{A_\pi\}_{\pi \in \widehat{\Gamma}}$ for (12.3), set $A'_\pi = \int_\Gamma \pi(g^{-1}) A_\pi \pi(g)\, dg$ for each $\pi \in \widehat{\Gamma}$. One can then show that $A'_\pi \pi(x) = \pi(x) A'_\pi$ for all $x \in \Gamma$, and so A'_π must be a multiple of the $d_\pi \times d_\pi$ identity matrix. It is easy to check that $\{A'_\pi\}$ is a new feasible solution for (12.3) having the same objective value. Therefore, when solving (12.3), we may as well assume from the beginning that all the matrices A_π are multiples of the identity; this leads to a linear program.

Note that by adding the constraints $\sum_{\pi \in \widehat{\Gamma}} d_\pi \operatorname{Tr}(A_\pi \pi(x)) \geq 0$, $(x \in \Gamma)$, to program (12.3), we can also obtain a frequency domain formulation for $\vartheta'_s(G)$.

12.1.3 Extension to graphs on homogeneous spaces

The forbidden distance graphs on spheres are important examples that we wish to be able to handle within our theory. Such graphs are not Cayley graphs, but this is only a minor problem. We now show how essentially all of the theory presented in Chapter 12 transfers to graphs G on Γ-homogeneous spaces equipped with the induced Haar measure when Γ is a compact subgroup of $\mathrm{Aut}(G)$. This idea is simply to "blow up" each vertex in G, replacing it with an independent set on the vertices in its preimage under the quotient map $\Gamma \to V$; edges are replaced with complete bipartite graphs. Then independent sets blow up to independent sets of the same Haar measure, and hence the independence ratios are the same. We now give the precise statement and details.

Let V be a locally compact Hausdorff topological space with Borel σ-algebra $\mathcal{B}$ and let $G = (V, E)$ be a graph. Consider the canonical left-action of $\mathrm{Aut}(G)$ on V, and let Γ be a compact group which is also a subgroup of $\mathrm{Aut}(G)$. Suppose that the action of Γ on V is transitive, and that the restricted action map $\Gamma \times V \to V$ is continuous. Suppose further that for each $x \in \Gamma$, each of the maps $v \mapsto xv$ is a homeomorphism of V. Then G will be called a *topological Schreier graph over the group* Γ.[1] Taking $\Gamma = SO(n)$, one sees that forbidden distance graphs on the unit n-sphere are topological Schreier graphs.

[1] This term was chosen because G is isomorphic to a Schreier left-coset graph over the group Γ.

Note that topological Schreier graphs necessarily have compact vertex sets.

Fix $v_0 \in V$ and let H be the stabilizer subgroup of v_0 in Γ. When Γ/H is equipped with the quotient topology, it is homeomorphic to V by [Fol95] Proposition 2.44. Let λ' be the Haar measure on $V \approx \Gamma/H$ normalized so that $\lambda'(V) = 1$, and define the *independence ratio* of G as

$$\tilde{\alpha}(G) = \sup\{\lambda'(I) : I \in \mathcal{B},\ I \text{ is independent}\}.$$

Let

$$X = \{x \in \Gamma : \{v_0, xv_0\} \in E\} \tag{12.4}$$

and $\widetilde{G} = \mathrm{Cay}(\Gamma, X)$. The graph $\widetilde{G}$ is a blow-up of G. The next proposition says that blowing up preserves independence ratios.

Proposition 12.3. $\tilde{\alpha}(G) = \tilde{\alpha}(\widetilde{G})$.

Proof. Let $\Psi : \Gamma/H \to V$ be the homeomorphism given by [Fol95, Proposition 2.44] and let $q : \Gamma \to \Gamma/H$ be the canonical quotient map. Call a subset of Γ *full* if it is a union of left-cosets of H. Then $\Psi \circ q$ induces a bijection between full independent sets in $\mathrm{Cay}(\Gamma, X)$ and independent sets in G. To see this, it suffices to show that $\Psi \circ q$ respects adjacencies and nonadjacencies. But for $x, y \in \Gamma$, we have by (12.4) that

$$\begin{aligned}\{\Psi(q(y)),\Psi(q(x))\} \in E &\iff \{yv_0, xv_0\} \in E \\ &\iff \{v_0, y^{-1}xv_0\} \in E \iff y^{-1}x \in X.\end{aligned}$$

This argument also shows that if I is an independent set in $\mathrm{Cay}(\Gamma, X)$, then IH is a full independent set. Since IH is at least as large as I, the proof will be done if we can show that $\Psi \circ q$ preserves Haar measures of full sets. But this follows immediately from the quotient integral formula (see e.g. [Fol95, Theorem 2.49]):

$$\int_\Gamma \mathbb{1}_{IH}(x)\ d\lambda(x) = \int_{\Gamma/H} \int_H \mathbb{1}_{IH}(xh)\ dh\ d\lambda'(xH)$$
$$= \int_{\Gamma/H} \mathbb{1}_{IH}(x)\ d\lambda'(xH) = \int_V \mathbb{1}_{\Psi(q(IH))}(v)\ d\lambda'(v).$$

□

With Γ, X, and H as in the statement of Proposition 12.3, if ϕ is any feasible solution for program (12.1), one can check that $\phi'(x) = \int_H \phi(xh)\ dh$ defines another feasible solution with the same objective value. Therefore, when computing ϑ_s of the blow-up graph as in program (12.1), we can suppose that ϕ is constant on the preimages of the projection map $\Gamma \to \Gamma/H$.

12.1.4 Forbidden distance graphs on S^{n-1}

In [BNdOFV09] and [dOF09], an extension of the Lovász ϑ-function is defined for forbidden distance graphs on the unit sphere in $\mathbb{R}^n$. We now show how this extension can be obtained as a special case of ϑ_s.

We first recall the setup from [BNdOFV09]. Let $n \geq 2$, and use $\langle \cdot, \cdot \rangle$ to denote the standard inner product on $\mathbb{R}^n$. Let S^{n-1} be the unit sphere in $\mathbb{R}^n$.

$$S^{n-1} := \{\xi \in \mathbb{R}^n : \langle \xi, \xi \rangle = 1\}$$

Let $D \subset [-1, 1]$ and consider the graph $G = (V, E)$ where $V = S^{n-1}$, and $E = \{\{\xi, \eta\} \subset S^{n-1} : \langle \xi, \eta \rangle \in D\}$. We want to upper bound the largest possible measure of an independent set in G, where we regard S^{n-1} as being equipped with the surface measure normalized so that S^{n-1} gets measure 1. It is shown in [BNdOFV09] (after a renormalization) that the following supremum provides such an upper bound:

$$\sup \Big\{ y_0 \ : \ \sum_{i=0}^{\infty} y_i = 1, \ \sum_{i=0}^{\infty} y_i C_i^{(n-2)/2}(d) = 0 \text{ for all } d \in D, \quad (12.5)$$
$$y_i \geq 0 \text{ for } i = 0, 1, 2, \ldots \Big\}.$$

The functions $C_i^{(n-2)/2}(x)$ are the Gegenbauer polynomials, defined in Section 5.1. Recall that we always use the normalization $C_i^{\nu}(1) = 1$.

Our aim is to recover this bound from our framework. Notice that G is not evidently a Cayley graph, so the definition in (12.1) does not immediately apply. We resolve this difficulty by regarding S^{n-1} as a homogeneous space and associating to G an auxiliarly Cayley graph $\widetilde{G}$ as explained in subsection 12.1.3.

Let $SO(n)$ be the group of $n \times n$ orthogonal matrices having determinant equal to 1. The next proposition exhibits a useful relationship between positive definite functions on S^{n-1} and functions of positive type on the group $SO(n)$; it will be handy when working out the examples involving spherical graphs.

Proposition 12.4. *Fix any point $\xi_0 \in S^{n-1}$, and let H be the stabilizer subgroup of ξ_0 in $SO(n)$ under the canonical action of $SO(n)$ on S^{n-1}. For each $t \in [-1,1]$ choose $x_t \in SO(n)$ so that $\langle \xi_0, x_t\xi_0 \rangle = t$. Let $\phi \in C_{\mathbb{R}}(SO(n))$ and define $f : [-1,1] \to \mathbb{R}$ by*

$$f(t) = \int_H \int_H \phi(hx_th')\ dh\ dh', \quad t \in [-1,1],$$

where the integrals are with respect to the Haar measure on the closed subgroup H of Γ. Define $\phi' : SO(n) \to \mathbb{R}$ by

$$\phi'(x) = f(\langle \xi_0, x\xi_0 \rangle), \quad x \in SO(n).$$

Then f is continuous, and f is positive definite when ϕ is of positive type. Moreover, f is positive definite if and only if ϕ' is of positive type.

Proof. To see that f is continuous, note that the definition of f does not depend on the particular choice of x_t, for if $\langle \xi_0, x\xi_0 \rangle = \langle \xi_0, y\xi_0 \rangle$, then $HxH = HyH$. The choice can be made so that $t \mapsto x_t$ is continuous from $[-1,1]$ to $SO(n)$. Letting $\varepsilon > 0$ be given, choose a symmetric conjugate-invariant neighbourhood U of e such that $|\phi(x) - \phi(y)| < \varepsilon$ whenever $y^{-1}x \in U$; this is possible by [Fol95, Proposition 2.6 and Lemma 5.24]. If $t_0, t_1 \in [-1,1]$ are close enough, then $x_{t_1}^{-1}x_{t_0} \in U$, and

$$\begin{aligned}
&\left| \int_H \int_H \phi(hx_{t_0}h')\ dh\ dh' - \int_H \int_H \phi(hx_{t_1}h')\ dh\ dh' \right| \\
&\le \int_H \int_H |\phi(hx_{t_0}h') - \phi(hx_{t_1}h')|\ dh\ dh' \\
&< \varepsilon,
\end{aligned}$$

where the last inequality holds since $(hx_{t_1}h')^{-1}(hx_{t_0}h') = (h')^{-1}x_{t_1}^{-1}x_{t_0}h' \in U$. This proves that f is continuous.

For the second assertion, assume that ϕ is of positive type. We show that f is positive definite. For this let $\xi_1, \ldots, \xi_s \in S^{n-1}$ be given. We want to show that the matrix $M = (f(\langle \xi_i, \xi_j \rangle))_{i,j=1}^s$ is positive semidefinite. Choose $y_1, \ldots, y_s \in SO(n)$ so that $\xi_i = y_i \xi_0$ for each $i = 1, 2, \ldots, s$, and let $v = (v_1, \ldots, v_s) \in \mathbb{R}^s$ be arbitrary. Then

$$f(\langle \xi_i, \xi_j \rangle) = f(\langle \xi_0, y_i^{-1} y_j \xi_0 \rangle) = \int_H \int_H \phi(h y_i^{-1} y_j h')\, dh\, dh', \quad (i, j = 1, \ldots, s).$$

Therefore, applying Proposition 9.4 to obtain a function $\psi \in L^2(SO(n))$ such that $\phi = \psi^* * \psi$, we get

$$\begin{aligned} v^t M v &= \sum_{i=1}^s \sum_{j=1}^s v_i v_j \int_H \int_H \phi(h y_i^{-1} y_j h')\ dh\ dh' \\ &= \sum_{i=1}^s \sum_{j=1}^s v_i v_j \int_H \int_H \int_{SO(n)} \overline{\psi(z)} \psi(z h y_i^{-1} y_j h')\ dz\ dh\ dh' \\ &= \sum_{i=1}^s \sum_{j=1}^s v_i v_j \int_H \int_H \int_{SO(n)} \overline{\psi(z y_i h^{-1})} \psi(z y_j h')\ dz\ dh\ dh' \\ &= \int_{SO(n)} \left| \sum_{i=1}^s v_i \int_H \psi(z y_i h)\ dh \right|^2\ dz \geq 0. \end{aligned}$$

We next show that ϕ' is of positive type when f is positive definite. Let distinct points $y_1, \ldots, y_s \in SO(n)$ be given. We want to show that the matrix $(\phi'(y_i^{-1} y_j))_{i,j=1}^s$ is positive semidefinite. But

$$\phi'(y_i^{-1} y_j) = f(\langle \xi_0, y_i^{-1} y_j \xi_0 \rangle) = f(\langle y_i \xi_0, y_j \xi_0 \rangle),$$

so the matrix is positive semidefinite by the definition of positive definiteness for functions $[-1,1] \to \mathbb{R}$ applied to the points $y_1\xi_0, \ldots, y_s\xi_0$.

To see that f is positive definite when ϕ' is of positive type, just use $\phi = \phi'$ in the first part of the proof. $\square$

Let ξ_0, H, and $\{x_t\}_{-1 \le t \le 1}$ be as in Proposition 12.4. Let X be the union of the double cosets Hx_dH, for $d \in D$, and set $\widetilde{G} = \mathrm{Cay}(SO(n), X)$. The graph $\widetilde{G}$ can be regarded as a blow-up of the graph G. We now prove that the value of (12.5) is equal to $\vartheta_s(\widetilde{G})$.

Suppose $y_0, y_1, \ldots$ is a feasible solution to the conic program (12.5). We construct a feasible solution to program (12.1) having the same objective value y_0. Define $f : [-1,1] \to \mathbb{R}$ by $f(t) = \sum_{i=0}^{\infty} y_i C_i^{(n-2)/2}(t)$. The series defining f converges uniformly to a continuous function of positive type by the Weierstrass M-test and Schoenberg's theorem. Moreover, we have $f(1) = 1$ and $f(d) = 0$ for every $d \in D$.

Now define $\phi' : SO(n) \to \mathbb{R}$ by $\phi'(x) = f(\langle \xi_0, x\xi_0 \rangle)$ as in Proposition 12.4. Then ϕ' is of positive type, and $\phi'(e) = f(1) = 1$. If $x \in X$, then $x = hx_dh'$ for some $d \in D$, and $h, h' \in H$; therefore

$$\phi'(x) = f(\langle \xi_0, hx_dh'\xi_0 \rangle) = f(\langle h^{-1}\xi_0, x_dh'\xi_0 \rangle) = f(\langle \xi_0, x_d\xi_0 \rangle) = f(d) = 0.$$

This shows that ϕ' is a feasible solution to the conic program (12.1), and we have by the quotient integral fomula ([Fol95] Theorem 2.49) that

$$\int_{SO(n)} \phi' \, d\lambda = \int_{SO(n)} f(\langle \xi_0, x\xi_0 \rangle) \, dx = \int_{S^{n-1}} f(\langle \xi_0, \xi \rangle) \, d\xi$$
$$= \sum_{i=0}^{\infty} y_i \int_{S^{n-1}} C_i^{(n-2)/2}(\langle \xi_0, \xi \rangle) \, d\xi = y_0.$$

This proves that the value of program (12.5) is no more than $\vartheta_s(\widetilde{G})$.

For the reverse inequality, let ϕ be any feasible solution to program (12.1), and define $f : [-1, 1] \to \mathbb{R}$ by $f(t) = \int_H \int_H \phi(hx_t h') \, dh \, dh'$. Then f is continuous and positive definite by Proposition 12.4, and clearly $f(d) = 0$ for all $d \in D$. So by Schoenberg's theorem we have

$$f(t) = \sum_{i=0}^{\infty} y_i C_i^{(n-2)/2}(t)$$

for some $y_i \geq 0$. Also

$$y_0 = \int_{SO(n)} f(\langle \xi_0, x\xi_0 \rangle) \, dx = \int_{SO(n)} \int_H \int_H \phi(hxh') \, dh \, dh' \, dx$$
$$= \int_{SO(n)} \phi(x) \, dx.$$

Finally $0 \leq \sum_{i=0}^{\infty} y_i = f(1) = \int_H \int_H \phi(hh') \, dh \, dh' \leq 1$ since $\phi(e) = 1$, and since scaling the coefficients y_i to satisfy $\sum_{i=0}^{\infty} y_i = 1$ can only make y_0 larger, we obtain a feasible solution to (12.5) with objective value at least that of ϕ. This shows that the value of program (12.5) is at least $\vartheta_s(\widetilde{G})$, as required.

12.2 Dual formulation

In this section we write down dual cone programs for (12.1) and (12.2). We prove weak duality and zero duality gap for our programs, and we show that from our duals one recovers the linear programming duals of the linear programs in [BNdOFV09], [dOFV10], and [dOF09] which give upper bounds for measures of independent sets in forbidden distance graphs on S^{n-1}, and for densities of 1-avoiding sets in $\mathbb{R}^n$.

As before, we let Γ be a fixed compact group, $X \subset \Gamma$ a sparse connection set, and we let $G = \mathrm{Cay}(\Gamma, X)$. For a Borel subset Y of Γ and $\nu \in \mathcal{M}(Y)$, we denote by $\widetilde{\nu}$ the extension of ν by zero to all of Γ. Formally, $\widetilde{\nu}$ is defined via the Riesz representation theorem as the unique regular Borel measure satisfying $\int f \, d\widetilde{\nu} = \int_Y f \, d\nu$ for all $f \in C(\Gamma)$.

We define

$$\vartheta_s^*(G) := \inf \left\{ \gamma : \gamma\delta_e + \widetilde{\nu} - \lambda \in \mathcal{Q}, \ \gamma \in \mathbb{R}, \ \nu \in \Sigma_{\mathbb{R}}(X) \right\}, \tag{12.6}$$

and

$$\vartheta_s'^*(G) := \inf \left\{ \gamma : \gamma\delta_e + \widetilde{\nu} - \lambda \in \mathcal{Q} + \Sigma_+(\Gamma), \ \gamma \in \mathbb{R}, \ \nu \in \Sigma_{\mathbb{R}}(X) \right\}. \tag{12.7}$$

Our first theorem is trivial, but important.

Theorem 12.5 (Weak duality). *For a sparse Cayley graph $G = \mathrm{Cay}(\Gamma, X)$, we have $\vartheta_s(G) \leq \vartheta_s^*(G)$ and $\vartheta_s'(G) \leq \vartheta_s'^*(G)$.*

Proof. Let ϕ be a feasible solution for (12.1); that is $\phi \in \mathcal{P}$, $\phi(e) = 1$, and $\phi|_X \equiv 0$. Also let $\gamma \in \mathbb{R}$, $\nu \in \Sigma_{\mathbb{R}}(X)$ be feasible for (12.6). Then

$$0 \leq \int \phi \, d(\gamma\delta_e + \widetilde{\nu} - \lambda) = \gamma\phi(e) + \int_X \phi \, d\nu - \int \phi \, d\lambda,$$

from which $\int \phi \, d\lambda \leq \gamma$ follows. This proves $\vartheta_s(G) \leq \vartheta_s^*(G)$.

For ϑ_s', suppose further that $\phi \geq 0$, so that ϕ is feasible for (12.2). Let $\mu \in \Sigma_+(\Gamma)$ be such that

$$\gamma\delta_e + \widetilde{\nu} - \lambda - \mu \in \mathcal{Q}.$$

Then

$$\begin{aligned} 0 &\leq \int \phi \, d(\gamma\delta_e + \widetilde{\nu} - \lambda - \mu) \\ &= \gamma\phi(e) + \int_X \phi \, d\nu - \int \phi \, d\lambda - \int \phi \, d\mu \\ &\leq \gamma - \int \phi \, d\lambda, \end{aligned}$$

as required. □

The next theorem says that the inequalities in Theorem 12.5 are in fact equalities. Many (but not all) of the main ingredients of the proof are contained in proofs of Slater's condition from semidefinite programming (see e.g. [BV04, Section 5.3.2] or [GM12, Theorem 4.7.1]). These results however do not apply directly to our programs, and we therefore give a full proof.

Theorem 12.6 (Zero duality gap)**.** *We have* $\vartheta_s(G) = \vartheta_s^*(G)$ *and* $\vartheta_s'(G) = \vartheta_s'^*(G)$

Proof. We prove only the first assertion, since the proof of the second is very similar. Define $\varphi : \mathbb{R} \oplus \Sigma_{\mathbb{R}}(X) \oplus \Sigma_{\mathbb{R}}(\Gamma) \to \mathbb{R} \oplus \Sigma_{\mathbb{R}}(\Gamma)$ by

$$\varphi(\gamma, \nu, \mu) = (\gamma, \mu - \gamma\delta_e - \widetilde{\nu}).$$

Let K be the closure of $\varphi(\mathbb{R} \oplus \Sigma_{\mathbb{R}}(X) \oplus \mathcal{Q})$ in the product topology on $\mathbb{R} \oplus \Sigma_{\mathbb{R}}(\Gamma)$, where we think of $\Sigma_{\mathbb{R}}(\Gamma)$ as being topologized with the total variation norm. Suppose $(\vartheta_s(G), -\lambda) \notin K$. By Theorem 10.1, there exist $\sigma \in \mathbb{R}$ and $f \in C_{\mathbb{R}}(\Gamma)$ not both zero, and $\rho \in \mathbb{R}$, such that

$$\sigma\gamma + \int f \; d(\mu - \gamma\delta_e - \widetilde{\nu}) \geq \rho \quad (\gamma \in \mathbb{R}, \nu \in \Sigma_{\mathbb{R}}(X), \mu \in \mathcal{Q}), \tag{12.8}$$

and $$\sigma\vartheta_s(G) - \int f \; d\lambda < \rho. \tag{12.9}$$

If $\sigma\gamma + \int f \; d(\mu - \gamma\delta_e - \widetilde{\nu}) < 0$, for some choice of γ, ν, and μ, then scaling by a large positive number would give a contradiction to (12.8). Hence

$$\sigma\gamma + \int f \; d(\mu - \gamma\delta_e - \widetilde{\nu}) \geq 0, \qquad (\gamma \in \mathbb{R}, \nu \in \Sigma_R(X), \mu \in \mathcal{Q}). \tag{12.10}$$

Setting $\gamma = 0$, $\nu = 0$, and letting μ range over $\mathcal{Q}$, equation (12.10) shows that $f \in \mathcal{P}$, by Theorem 10.2. Setting $\mu = 0$, $\gamma = 0$, and letting ν range over $\Sigma_{\mathbb{R}}(X)$, equation (12.10) also shows that $f|_X \equiv 0$. Setting $\mu = 0$ and $\nu = 0$ in (12.10), we obtain $\sigma\gamma - \gamma f(e) \geq 0$ for all $\gamma \in \mathbb{R}$. Therefore $f(e) = \sigma$. Since f and σ cannot both be zero, and since $f \in \mathcal{P}$, it follows that $\sigma > 0$. We may therefore replace f by $\sigma^{-1}f$ in (12.9), obtaining $\vartheta_s(G) - \int f \; d\lambda < \rho$. Setting all variables to zero in (12.8) gives $0 \geq \rho$, whence $\vartheta_s(G) < \int f \, d\lambda$; this is a contradiction since f is a feasible solution to program (12.1).

We have just shown that there exist sequences $\{\gamma_n\} \subset \mathbb{R}$, $\{\nu_n\} \subset \Sigma_{\mathbb{R}}(X)$, and $\{\mu_n\} \subset \mathcal{Q}$ such that $\gamma_n \to \vartheta_s(G)$, and $\mu_n - \gamma_n\delta_e - \widetilde{\nu}_n \to -\lambda$ in the total variation norm as $n \to \infty$. Let

$$q_n := (1-\varepsilon)\gamma_n\delta_e + 2\varepsilon\delta_e + (1-\varepsilon)\widetilde{\nu}_n - \lambda.$$

We next claim that for every $\varepsilon > 0$, there exists $n_0 = n_0(\varepsilon)$ such that $q_n \in \mathcal{Q}$ for all $n \geq n_0$. Let

$$B = \{\omega \in \Sigma_{\mathbb{R}}(\Gamma) : \|2\delta_e - \lambda - \omega\| < 1\}$$

be the total variation norm ball of radius 1 about the measure $2\delta_e - \lambda$. Then $B \subset \mathcal{Q}$, for if $\omega \in B$, and $\phi \in \mathcal{P}$ with $\phi(e) = 1$, then

$$1 > \left| \int \phi \, d(2\delta_e - \lambda - \omega) \right| = \left| 2 - \int \phi \, d\lambda - \int \phi \, d\omega \right|,$$

and from this it follows that $\int \phi \, d\omega > 0$ since $0 \leq \int \phi \, d\lambda \leq 1$; so $\omega \in \mathcal{Q}$ by Theorem 10.2. Now let

$$p_n = \gamma_n\delta_e + \widetilde{\nu}_n - \mu_n - \lambda.$$

Then

$$q_n = (1-\varepsilon)p_n + (1-\varepsilon)\mu_n + \varepsilon(2\delta_e - \lambda),$$

and $p_n \to 0$ in total variation norm as $n \to \infty$. So $q_n \in (1-\varepsilon)\mu_n + \varepsilon B \subset \mathcal{Q}$ for sufficiently large n. This proves the claim, and we have therefore shown that putting $\gamma = (1-\varepsilon)\gamma_n + 2\varepsilon$ and $\nu = (1-\varepsilon)\nu_n$ gives a feasible solution for (12.6) for large enough n. Taking $\varepsilon \to 0$ we obtain $\vartheta_s^*(G) \leq \vartheta_s(G)$. The reverse inequality is Theorem 12.5. $\square$

12.2.1 Forbidden distance graphs on S^{n-1}, dual formulation

Suppose $n \geq 2$, and let G, D, $\widetilde{G}$, X, H, ξ_0, and $\{x_t\}_{-1\leq t\leq 1}$ be as in subsection 12.1.4. When D is finite, a dual of the infinite-dimensional linear program (12.5) is given in [dOF09]:

$$\inf\Big\{z : z + \sum_{d\in D} z_d \geq 1,\ z + \sum_{d\in D} z_d C_i^{(n-2)/2}(d) \geq 0 \ \ (i \geq 1),\\ z \in \mathbb{R},\ z_d \in \mathbb{R} \text{ for } d \in D\Big\}. \qquad (12.11)$$

We now show that (12.11) equals $\vartheta_s^*(\widetilde{G})$. By Theorem 12.6, this shows that there is no duality gap in the linear programs (12.5) and (12.11) defined in [dOF09]. The appropriate generalization of program (12.11) in the case when D is infinite will now also become clear.

Let $\gamma \in \mathbb{R}$, $\nu \in \Sigma_{\mathbb{R}}(X)$ be a feasible solution for (12.6). We show how to construct a feasible solution for (12.11) having objective value γ. The linear functional on $C_{\mathbb{R},0}(D)$ defined by $\varphi(f) = \int_X f(\langle \xi_0, x\xi_0\rangle)\, d\nu(x)$ is bounded, so by the Riesz representation theorem there exists a $\mu \in \mathcal{M}_{\mathbb{R}}(D)$ such that $\int_D f\ d\mu = \varphi(f)$ for all $f \in C_{\mathbb{R},0}(D)$.

For every $i \geq 0$, the polynomial $C_i^{(n-2)/2}(t)$ in t is a positive definite function by Schoenberg's theorem. Therefore if $\phi(x) = C_i^{(n-2)/2}(\langle \xi_0, x\xi_0\rangle)$, then by Proposition 12.4 and Theorem 10.2, we have

$$\begin{aligned} 0 &\leq \int_{SO(n)} \phi(x)\ d(\gamma\delta_e + \widetilde{\nu} - \lambda) \\ &= \gamma + \int_D C_i^{(n-2)/2}(t)\ d\mu(t) - \int_{SO(n)} \phi(x)\ d\lambda(x). \end{aligned}$$

But $\int_{SO(n)} \phi(x)\ d\lambda(x)$ is equal to 0 if $i \geq 1$, and it is equal to 1 otherwise. Therefore γ, μ is a feasible solution for the conic optimization program

$$\inf \Big\{ \gamma : \gamma + \int_D \ d\mu(t) \geq 1, \gamma + \int_D C_i^{(n-2)/2}(t)\ d\mu(t) \geq 0 \quad (i \geq 1), \\ \gamma \in \mathbb{R},\ \mu \in \mathcal{M}(D) \Big\}, \tag{12.12}$$

having value γ. When D is a finite set, program (12.12) is identical to (12.11).

We have shown that the value of (12.12) is no more than $\vartheta_s^*(\widetilde{G})$. For the reverse inequality, suppose that $\gamma \in \mathbb{R}, \mu \in \mathcal{M}(D)$ is a feasible solution for (12.12). We construct a feasible solution for (12.6) having the same objective value. The linear functional $\phi \mapsto \int_{-1}^{1} \int_H \int_H \phi(hx_t h')\ dh\ dh'\ d\mu(t)$ on $C_{\mathbb{R},0}(X)$ is given by integration against some $\nu \in \mathcal{M}_{\mathbb{R}}(X)$. It is easy to check that $\nu \in \Sigma_{\mathbb{R}}(X)$. To see that $\gamma\delta_e + \widetilde{\nu} - \lambda \in \mathcal{Q}$, we use Theorem 10.2. Let $\phi \in \mathcal{P}$ be arbitrary with $\phi(e) = 1$, and define $f : [-1, 1] \to \mathbb{R}$ by $f(t) = \int_H \int_H \phi(hx_t h')\ dh\ dh'$. Then f is positive definite by Proposition 12.4, and so by Schoenberg's theorem there exist numbers $y_i \geq 0$, $(i = 0, 1, 2, \dots)$ such that

$$f(t) = \sum_{i=0}^{\infty} y_i C_i^{(n-2)/2}(t),$$

with uniform convergence on $[-1, 1]$. Notice that $0 < \sum_{i=0}^{\infty} y_i = f(1) \leq 1$ since ϕ must be strictly positive in a neighbourhood of e.

We therefore have

$$
\begin{aligned}
\int \phi \, d(\gamma\delta_e + \widetilde{\nu} - \lambda) = {} & \gamma + \int_X \phi \, d\nu - \int \phi \, d\lambda \\
= {} & \gamma + \int_{-1}^{1} \int_H \int_H \phi(hx_t h') \, dh \, dh' \, d\mu(t) - \int_{S^{n-1}} f(\xi) \, d\xi \\
= {} & \gamma + \int_{-1}^{1} f(t) \, d\mu(t) - y_0 \\
= {} & \gamma + \sum_{i=0}^{\infty} y_i \int_{-1}^{1} C_i^{(n-2)/2}(t) \, d\mu(t) - y_0 \\
= {} & y_0 \Big(f(1)^{-1}\gamma + \int_{-1}^{1} \, d\mu(t) - 1 \Big) \\
& + \sum_{i=1}^{\infty} y_i \left(f(1)^{-1}\gamma + \int_{-1}^{1} C_i^{(n-2)/2}(t) \, d\mu(t) \right) \geq 0.
\end{aligned}
$$

This shows that γ, ν is a feasible solution for (12.6). Its value is γ, so $\vartheta_s^*(\widetilde{G})$ is at most the value of (12.12).

We finally remark that when D is nonempty, one can rewrite program (12.12) as

$$
\inf_{\mu \in \mathcal{M}(D), \int_D d\mu = 1} \frac{-\inf_{i \geq 1} \int_D C_i^{(n-2)/2}(t) \, d\mu(t)}{1 - \inf_{i \geq 1} \int_D C_i^{(n-2)/2}(t) \, d\mu(t)}. \tag{12.13}
$$

The upper bound (12.13) for the independence ratio of a forbidden distance graph on S^{n-1} appears in [BDdOFV14, Section 4.2]; it also recovers and extends Theorem 6.2 from [BNdOFV09].

12.3 Further examples

In this section we demonstrate the application of ϑ_s to a number of examples, in particular recovering the cone programs from [dOF09] and [dOFV10]. Strong duality for these programs is therefore established as a consequence of Theorem 12.6.

12.3.1 Finite cyclic groups

This is the simplest kind of Cayley graph. Let $n \geq 2$ be an integer, let $C_n = \mathbb{Z}/n\mathbb{Z}$ be the cyclic group of order n, and let $X \subset C_n$ be a connection set. Let $G = \mathrm{Cay}(C_n, X)$. The characters of C_n are the functions $x \mapsto e^{2\pi i x k/n}$ for $k \in C_n$. Using the frequency domain formulation (12.3) of $\vartheta_s(G)$, and performing some simplification, we obtain

$$\begin{aligned}\vartheta_s(G) &= \sup\Big\{a_0 : \sum_{\chi \in \widehat{C_n}} a_\chi = 1, \sum_{\chi \in \widehat{C_n}} a_\chi \chi(x) = 0 \ (x \in X), \\ &\qquad a_{\overline{\chi}} = a_\chi, a_\chi \geq 0 \ (\chi \in \widehat{C_n})\Big\} \\ &= \sup\Big\{a_0 : \sum_{k \in C_n} a_k = 1, \sum_{k \in C_n} a_k \cos(2\pi x k/n) = 0, \ (x \in X), \\ &\qquad a_{-k} = a_k, a_k \geq 0 \ (k \in C_n)\Big\}.\end{aligned}$$

This linear program can be solved analytically or with any linear programming solver, and Proposition 12.2 shows that it computes $\vartheta(G)/n$.

12.3.2 Circle group

Let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ be the circle group, let $X \subset \mathbb{T}$ be a sparse connection set, and put $G = \mathrm{Cay}(\mathbb{T}, X)$. The characters of $\mathbb{T}$ are given by $x \mapsto e^{2\pi i x k}$ for $k \in \mathbb{Z}$, and so (12.3) becomes

$$\begin{aligned}\vartheta_s(G) &= \sup\Big\{a_0 : \sum_{k\in\mathbb{Z}} a_k = 1,\ \sum_{k\in\mathbb{Z}} a_k e^{2\pi i k x} = 0\ (x \in X),\\ &\qquad\qquad a_k = a_{-k},\ a_k \geq 0\ (k \in \mathbb{Z})\Big\}\\ &= \sup\Big\{a_0 : \sum_{k\in\mathbb{Z}} a_k = 1,\ a_0 + 2\sum_{k\geq 1} a_k \cos(2\pi k x) = 0\ (x \in X),\\ &\qquad\qquad a_k = a_{-k},\ a_k \geq 0\ (k \geq 0)\Big\}.\end{aligned}$$

By Theorem 9.7, the dual formulation (12.6) becomes

$$\begin{aligned}&\inf\Big\{\gamma : \gamma + \int_X d\nu \geq 1, \gamma + \int_X e^{2\pi i k x}\ d\nu(x) \geq 0\\ &\qquad (x \in X, 0 \neq k \in \mathbb{Z}),\ \gamma \in \mathbb{R},\ \nu \in \Sigma_{\mathbb{R}}(X)\Big\}\\ &= \inf\Big\{\gamma : \gamma + \int_X d\nu \geq 1, \gamma + \int_X \cos(2\pi k x)\ d\nu(x) \geq 0 \qquad (12.14)\\ &\qquad (x \in X, 0 \neq k \in \mathbb{Z}),\ \gamma \in \mathbb{R},\ \nu \in \Sigma_{\mathbb{R}}(X)\Big\}.\end{aligned}$$

Assuming X is nonempty, we therefore have

$$\vartheta_s(G) = \inf_{\nu\in\Sigma_{\mathbb{R}}(X), \int d\nu = 1} \frac{-\inf_{k\geq 1}\int_X \cos(2\pi k x)\ d\nu(x)}{1 - \inf_{k\geq 1}\int_X \cos(2\pi k x)\ d\nu(x)}.$$

12.3.3 Circle group, one forbidden distance

We specialize Example 12.3.2 to the case $X = \{x, -x\}$. For this case, the computation of $\vartheta_s(G)$ also appears in [dOF09, Section 3.5a];

now, equipped with Theorem 4.1, we can also investigate when the inequality is strict. With the same setup as in Example 12.3.2, suppose $X = \{x, -x\}$ for some $x \in \mathbb{T}$. We then have $\vartheta_s(G) = \frac{\inf_{k\geq 1}\cos(2\pi kx)}{1-\inf_{k\geq 1}\cos(2\pi kx)}$. If x is irrational, then $\inf_{k\geq 1}\cos(2\pi kx) = -1$. If $x = p/q$ with p and $q > 0$ relatively prime integers, then $\inf_{k\geq 1}\cos(2\pi kx)$ is equal to -1 if q is even, and it is strictly greater than -1 otherwise. Therefore $\vartheta_s(G) = \tilde{\alpha}(G) = 1/2$ if either x is irrational or q is even. When $q = 3$, we have $\vartheta_s(G) = \tilde{\alpha}(G) = 1/3$. If $\cos(2\pi kp/q)$ is not equal to one of $\pm 1, \pm\frac{1}{2}$, then it is irrational; therefore $\tilde{\alpha}(G) = \frac{q-1}{2q} < \vartheta_s(G)$ when $q \neq 3$ is odd.

12.3.4 n-dimensional torus

Let $\mathbb{T}^n = (\mathbb{R}/\mathbb{Z})^n$ be the n-dimensional torus. The characters of $\mathbb{T}^n$ are the functions $x \mapsto e^{2\pi i(x\cdot k)}$, for $x \in \mathbb{T}^n$, $k \in \mathbb{Z}^n$. Fix some small $\delta > 0$, and let X_δ be the image of the set $\{(x_1, \ldots, x_n) \in \mathbb{R}^n : x_1^2 + \cdots + x_n^2 = \delta\}$ under the canonical projection $\mathbb{R}^n \to (\mathbb{R}/\mathbb{Z})^n$, and put $G_\delta = \mathrm{Cay}(\mathbb{T}^n, X_\delta)$. One can express $\vartheta_s^*(G_\delta)$ via the formulation (12.6):

$$\vartheta_s^*(G_\delta) = \inf \Big\{\gamma : \gamma + \int_X d\nu \geq 1, \gamma + \int_X e^{2\pi i(x\cdot k)}\, d\nu(x) \geq 0 \quad (k \in \mathbb{Z}^n), \gamma \in \mathbb{R}, \nu \in \Sigma_{\mathbb{R}}(X)\Big\} \tag{12.15}$$

Since X is just a sphere of radius δ, program (12.15) can be rewritten as

$$\vartheta_s^*(G_\delta) = \inf \Big\{ \gamma : \gamma + \int_{S^{n-1}} d\nu \geq 1, \gamma + \int_{S^{n-1}} e^{2\pi i \delta(\xi \cdot k)} \, d\nu(\xi) \geq 0 \tag{12.16}$$
$$(k \in \mathbb{Z}^n), \gamma \in \mathbb{R}, \nu \in \Sigma_{\mathbb{R}}(S^{n-1}) \Big\}.$$

12.3.5 One forbidden distance in $\mathbb{R}^n$

The next example was first treated in [dOFV10] and [dOF09], and later redone in [BDdOFV14]. Improved bounds have since been obtained in [BPT14] via a combinatorial strengthening of this method.

The *upper density* of a Lebesgue measurable set $E \subset \mathbb{R}^n$ is given by the formula

$$\limsup_{r \to \infty} \frac{1}{r^n} \int_{[-r/2, r/2]^n} \mathbb{1}_E(x) \, dx.$$

In this example we wish to upper bound the upper density of sets E which *avoid distance* 1; that is, for which $\|x - y\| \neq 1$ for all pairs of points $x, y \in E$.

Let $E \subset \mathbb{R}^n$ be a measurable set avoiding distance 1 and let $\varepsilon > 0$ be given. Let D be the upper density of E and choose $r > 0$ so that $\frac{1}{r^n} \int_{[-r/2,r/2]^n} \mathbb{1}_E(x) \, dx > D - \varepsilon$. Let $E' = r^{-1}(E \cap [-r/2, r/2]^n) \subset [-1/2, 1/2]^n$. Then $\|x - y\| \neq 1/r$ for all $x, y \in E'$. By removing a set of measure $O(1/r)$ from near the boundary of $[-1/2, 1/2]^n$, we can obtain from E' a set F whose image I under the canonical projection $\mathbb{R}^n \to \mathbb{T}^n$ is an independent set in $G_{1/r}$ from Example 12.3.4. The

Haar measure of I in $\mathbb{T}^n$ is equal to the measure of F. Therefore $D - \varepsilon \leq \tilde{\alpha}(G_{1/r}) \leq \vartheta_s(G_{1/r})$, and taking limits then gives

$$D \leq \limsup_{r\to\infty} \tilde{\alpha}(G_{1/r}) \leq \limsup_{r\to\infty} \vartheta_s(G_{1/r})$$

for all $\varepsilon > 0$. Since the rational points are dense in $\mathbb{R}^n$, we have from (12.16) that

$$\begin{aligned} &\limsup_{r\to\infty} \vartheta_s^*(G_{1/r}) \\ &= \inf \Big\{ \gamma : \gamma + \int_{S^{n-1}} d\nu \geq 1, \gamma + \int_{S^{n-1}} e^{2\pi i(\xi\cdot x)}\, d\nu(\xi) \geq 0 \\ &\qquad (x \in \mathbb{R}^n), \gamma \in \mathbb{R}, \nu \in \Sigma_{\mathbb{R}}(S^{n-1}) \Big\}. \end{aligned} \tag{12.17}$$

Denoting the normalized surface measure on S^{n-1} by $d\xi$, it is easy to see that if γ, ν is a feasible solution for program (12.17), then $\gamma, (\int d\nu)\ d\xi$ is another feasible solution having the same objective value. Therefore defining $\Omega : \mathbb{R} \to \mathbb{R}$ by $\Omega(t) = \int_{S^{n-1}} e^{2\pi i t(\xi\cdot\xi_0)}\, d\xi$, where $\xi_0 \in S^{n-1}$ is any point, we have

$$\begin{aligned} &\limsup_{r\to\infty} \vartheta_s^*(G_{1/r}) \\ &= \inf \Big\{ \gamma : \gamma + \sigma \geq 1, \gamma + \sigma\Omega(t) \geq 0\ (t \geq 0), \gamma \in \mathbb{R}, \sigma \in \mathbb{R} \Big\} \\ &= \frac{-\inf_{t\geq 0} \Omega(t)}{1 - \inf_{t\geq 0} \Omega(t)}, \end{aligned}$$

thereby recovering the upper bound for D given in [dOFV10]. An expression for Ω in terms of Bessel functions is also given there.

Chapter 13

Dense Cayley graphs over compact groups

13.1 ϑ_d and ϑ'_d: primal formulation

We will reuse the notation from Chapter 12. Let Γ be a compact group, and let $G = \mathrm{Cay}(\Gamma, X)$ be a dense Cayley graph. We define $\vartheta_d(G)$ as the value of the following conic optimization program.

$$\vartheta_d(G) := \sup\left\{1 + \int_{X^c} d\nu : \nu \in \Sigma_{\mathbb{R}}(X^c),\ \delta_e + \widetilde{\nu} \in \mathcal{Q}\right\} \tag{13.1}$$

We define $\vartheta'_d(G)$ as

$$\vartheta'_d(G) := \sup\left\{1 + \int_{X^c} d\nu : \nu \in \Sigma_{+}(X^c),\ \delta_e + \widetilde{\nu} \in \mathcal{Q}\right\}. \tag{13.2}$$

Notice that $\nu = 0$ gives a feasible solution to programs (13.1) and (13.2) of value 1; in particular both programs are always feasible and their values are at least 1.

Theorem 13.1. *If G is a dense Cayley graph, then $\vartheta_d(G) \geq \vartheta'_d(G) \geq \alpha(G)$.*

Proof. The first inequality is clear. For the second, let $I \subset \Gamma$ be a nonempty independent set in G. Then $|I| < \infty$ by Proposition 11.1. Set $\nu = \frac{1}{|I|}\mu^* * \mu - \delta_e$ where $\mu = \sum_{x \in I} \delta_x$. Then ν is feasible for program (13.2) by Corollary 9.8, and its objective value is $1 + \int_{X^c} d\nu = 1 + \widehat{\mu}(1)^* \widehat{\mu}(1)/|I| - 1 = \frac{1}{|I|}\left(\int_{X_c} d\mu\right)^2 = |I|$. □

Proposition 13.2. *When $G = \mathrm{Cay}(\Gamma, X)$ is a finite graph, we have $\vartheta(G) = \vartheta_d(G)$ and $\vartheta'(G) = \vartheta'_d(G)$.*

Proof. We give the proof only for ϑ, the proof for ϑ' being very similar. Let $\nu \in \Sigma_{\mathbb{R}}(X^c)$ be a feasible solution for (13.1) and let $\phi(x) = \delta_e(\{x\}) + \nu(\{x\})$. Then ϕ is a function of positive type. Using $n = |\Gamma|$, define the matrix $A = (a_{x,y})_{x,y \in \Gamma}$ by $a_{x,y} = \phi(y^{-1}x)/n$. Then A is symmetric, and we have $\mathrm{Tr}(A) = 1$ and $a_{x,y} = 0$ whenever $y^{-1}x \in X$. It is easily verified that A is positive semidefinite, so A is a feasible solution for program (2.6). Its objective value is

$$\sum_{x,y \in \Gamma} a_{x,y} = \frac{1}{n} \sum_{x,y \in \Gamma} \phi(y^{-1}x) = 1 + \int_{X^c} d\nu.$$

This shows $\vartheta(G) \geq \vartheta_d(G)$.

For the reverse inequality, let $A = (a_{x,y})_{x,y\in\Gamma}$ be a feasible solution for program (2.6), and for each $z \in \Gamma$, let $\phi(z) = \sum_{g\in\Gamma} a_{gz,g}$. Then ϕ is positive definite; indeed if $z_1, \ldots, z_s \in \Gamma$ and $v = (v_1, \ldots, v_s) \in \mathbb{R}^s$, we have

$$\begin{aligned} v^t(\phi(z_j^{-1}z_i))v &= \sum_{i,j=1}^{s} \phi(z_j^{-1}z_i)v_i v_j = \sum_{i,j=1}^{s} \sum_{g\in\Gamma} a_{gz_j^{-1}z_i,g} v_i v_j \\ &= \sum_{g\in\Gamma} \sum_{i,j=1}^{s} a_{gz_i,gz_j} v_i v_j, \end{aligned}$$

and the last quantity is nonnegative since $(a_{gz_i,gz_j})_{i,j=1}^s$ is a principal submatrix of the positive semidefinite matrix A for each $g \in \Gamma$ (possibly with multiplicities, but this makes no difference). It is now easy to check that $\phi\, d\lambda - \delta_e$ is a feasible solution for (13.1) with objective value $\sum_{x,y\in\Gamma} a_{x,y}$. Therefore $\vartheta_d(G) \geq \vartheta(G)$. □

13.2 Dual formulation

Once again let $G = \mathrm{Cay}(\Gamma, X)$ be a dense Cayley graph. We define

$$\vartheta_d^*(G) := \inf\left\{1 + \phi(e) : \phi \in \mathcal{P}, \phi|_{X_c} \equiv -1\right\} \tag{13.3}$$

and

$$\vartheta_d'^*(G) := \inf\left\{1 + \phi(e) : \phi \in \mathcal{P}, \phi|_{X_c} \leq -1\right\}. \tag{13.4}$$

Just as in the sparse case, the weak duality proofs are straightforward.

Theorem 13.3 (Weak duality). *When $G = \mathrm{Cay}(\Gamma, X)$ is a dense Cayley graph, we have $\vartheta_d(G) \leq \vartheta_d^*(G)$ and $\vartheta_d'(G) \leq \vartheta_d'^*(G)$.*

Proof. Let $\nu \in \Sigma_{\mathbb{R}}(X^c)$ be a feasible solution for (13.1), and let $\phi \in \mathcal{P}$ be a feasible solution for (13.3). By Theorem 10.2 we have $0 \leq \int \phi \, d(\delta_e + \widetilde{\nu}) = \phi(e) - \int_{X^c} d\nu$, so $1 + \int_{X^c} d\nu \leq 1 + \phi(e)$. The proof for ϑ'_d is similar. □

Theorem 13.4 (Zero duality gap)**.** *When G is a dense Cayley graph, we have $\vartheta_d(G) = \vartheta^*_d(G)$ and $\vartheta'_d(G) = \vartheta'^*_d(G)$.*

Proof. We prove only the first assertion, the proof of the second being very similar. Define $\varphi : \Sigma_{\mathbb{R}}(X^c) \oplus \Sigma_{\mathbb{R}}(\Gamma) \to \Sigma_{\mathbb{R}}(\Gamma) \oplus \mathbb{R}$ by $\varphi(\nu, \mu) = (-\widetilde{\nu} + \mu, \int_{X^c} d\nu)$. Let $I = \varphi(\Sigma_{\mathbb{R}}(X^c) \oplus \mathcal{Q})$, and let $\overline{I}$ be the closure of I in $\Sigma_{\mathbb{R}}(\Gamma) \oplus \mathbb{R}$, where we regard $\Sigma_{\mathbb{R}}(\Gamma)$ as topologized with the total variation norm.

We claim that $(\delta_e, \vartheta^*_d(G) - 1) \in \overline{I}$. Suppose not; then by Theorem 10.1 there exists $\phi \in C_{\mathbb{R}}(\Gamma)$ and $\sigma \in \mathbb{R}$, and $\rho \in \mathbb{R}$ such that

$$\int \phi \, d(-\widetilde{\nu} + \mu) + \sigma \int_{X^c} d\nu \geq \rho \quad (\nu \in \Sigma_{\mathbb{R}}(X^c),\ \mu \in \mathcal{Q}) \tag{13.5}$$

$$\text{and} \quad \int \phi \, d\delta_e + \sigma(\vartheta^*_d(G) - 1) < \rho. \tag{13.6}$$

We have

$$\int \phi \, d(-\widetilde{\nu} + \mu) + \sigma \int_{X^c} d\nu \geq 0 \tag{13.7}$$

for all $\nu \in \Sigma_{\mathbb{R}}(X^c)$ and $\mu \in \mathcal{Q}$, since otherwise scaling by a large positive number would give a contradiction to (13.5). Putting $\nu = 0$ and letting μ range over $\mathcal{Q}$ shows by Theorem 10.2 that $\phi \in \mathcal{P}$. Putting

$\mu = 0$ in (13.7), we find that $\int_{X^c}(\sigma - \phi)\, d\nu \geq 0$ for all $\nu \in \Sigma_{\mathbb{R}}(X^c)$, which is only possible if $\phi|_{X^c} \equiv \sigma$.

Now setting $\nu = 0$ and $\mu = 0$ in (13.5) shows that $0 \geq \rho$, and combining with (13.6) gives

$$\phi(e) + \sigma(\vartheta_d^*(G) - 1) < 0. \tag{13.8}$$

Since $\phi \in \mathcal{P}$, it follows that $\sigma < 0$. Therefore $-\sigma^{-1}\phi$ is a feasible solution to program (13.3), and (13.8) implies that $1 + \phi(e) < \vartheta_d^*(G)$; this contradiction proves the claim.

We have just shown that there exist sequences $\{\nu_n\} \subset \Sigma_{\mathbb{R}}(X^c)$ and $\{\mu_n\} \subset \mathcal{Q}$ such that $-\widetilde{\nu_n} + \mu_n \to \delta_e$ in total variation norm and $\int_{X^c} d\nu_n \to \vartheta_d^*(G) - 1$. Let $\varepsilon > 0$ be given. We claim that there exists $n_0 = n_0(\varepsilon)$ such that $\delta_e + (1-\varepsilon)\widetilde{\nu_n} \in \mathcal{Q}$ whenever $n \geq n_0$. To see this, let $B = \{\omega \in \mathcal{M}_{\mathbb{R}}(\Gamma) : \|\omega - \delta_e\| < 1/2\}$. Then $B \subset \mathcal{Q}$, for if $\omega \in B$ and $\psi \in \mathcal{P}$ with $\psi(e) = 1$, then $|\int \psi\, d\omega - 1| = |\int \psi\, d\omega - \psi(e)| < 1/2$, which can happen only if $\int \psi\, d\omega > 0$. Now,

$$\begin{aligned}\delta_e + (1-\varepsilon)\widetilde{\nu_n} &= (1-\varepsilon)(\delta_e + \widetilde{\nu_n} - \mu_n) \\ &\quad + (1-\varepsilon)\mu_n + \varepsilon\delta_e \in (1-\varepsilon)\mu_n + \varepsilon B \subset \mathcal{Q}\end{aligned}$$

for sufficiently large n, since $\delta_e + \widetilde{\nu_n} - \mu_n \to 0$. This proves the claim; therefore $(1-\varepsilon)\widetilde{\nu_n}$ is a feasible solution of program (13.1) when $n \geq n_0$, and its value tends to $1 + (1-\varepsilon)(\vartheta_d^*(G) - 1)$. It now follows that $\vartheta_d(G) \geq \vartheta_d^*(G)$. The reverse inequality is Theorem 13.3. $\square$

13.3 ϑ_d and the chromatic number of sparse Cayley graphs

The ϑ-number of a finite graph gives a lower bound for the chromatic number of the complementary graph. The analogous statement for infinite graphs relies on the fact that the complement of a sparse Cayley graph is dense.

Theorem 13.5. *Let* $G = \mathrm{Cay}(\Gamma, X)$ *be a sparse Cayley graph. Then* $\chi_m(G) \geq \vartheta_d(G^c)$.

Proof. Let $C_1, \ldots, C_k$ be a partition of Γ into Borel sets which are independent in G. Since we are claiming $k \geq \vartheta_d(G^c)$, it suffices to suppose that $\lambda(C_i) > 0$ for all $i = 1, \ldots, k$; we simply discard the nullsets in the partition. Let $\psi = \sum_{i=1}^{k} \mathbb{1}_{C_i}^{*} * \mathbb{1}_{C_i}/\lambda(C_i)$. Then

$$\widehat{\psi}(1) = \sum_{i=1}^{k} \frac{|\widehat{\mathbb{1}_{C_i}}(1)|^2}{\lambda(C_i)} = \sum_{i=1}^{k} \lambda(C_i) = 1.$$

Therefore by Theorem 9.1, it follows that $\phi = \psi - \mathbb{1}_\Gamma \in \mathcal{P}$. Since each C_i is an independent set, we have $\psi|_X \equiv 0$, and therefore $\phi|_X \equiv -1$. We conclude that ϕ is a feasible solution for program (13.3) for the graph G^c. Since $\mathbb{1}_{C_i}^{*} * \mathbb{1}_{C_i}(e) = \lambda(C_i)$ for each i, its value is $1 + \phi(e) = 1 + \psi(e) - 1 = \sum_{i=1}^{k} 1 = k$, and therefore $\vartheta_d(G^c) \leq k$ by Theorem 13.3. $\square$

Theorem 13.6 combined with Theorem 13.1 gives us a version of the famous “Sandwich Theorem” [Lov86, Theorem 5.4]:

$$\alpha(G^c) \leq \vartheta_d(G^c) \leq \chi_m(G).$$

The next theorem is the analogue of [Lov79, Theorem 12].

Theorem 13.6. *Let* $G = \mathrm{Cay}(\Gamma, X)$ *be a sparse Cayley graph. We have* $\vartheta_s(G)\vartheta_d(G^c) = 1$.

Proof. Let $\varepsilon > 0$ be small and let $\phi \in \mathcal{P}$ be a feasible solution to program (12.1) having objective value $\int \phi \geq \beta := \vartheta_s(G) - \varepsilon$, and let $\psi = \beta^{-1}\phi - \mathbb{1}_\Gamma$. Since $0 < \beta < 1$, we have $\widehat{\psi}(1) = \beta^{-1}\widehat{\phi}(1) - 1 \geq 0$, and it follows from Theorem 9.1 that $\psi \in \mathcal{P}$. One now easily checks that ψ is feasible for program (13.3) for G^c having objective value $1 + \psi(e) = \beta^{-1}$. Therefore $\vartheta_d(G^c) \leq \beta^{-1}$, and since ε was arbitrary we get $\vartheta_s(G)\vartheta_d(G^c) \leq 1$.

For the reverse inequality, again let $\varepsilon > 0$ be given, and this time let $\psi \in \mathcal{P}$ be a feasible solution to program (13.3) for G^c having objective value $1 + \psi(e) \leq \gamma := \vartheta_d(G^c) + \varepsilon$. Let $\phi = (\psi + \mathbb{1}_\Gamma)/(1 + \psi(e))$. Then ϕ is a feasible solution to (12.1) and since $\int(\psi + \mathbb{1}_\Gamma) \geq 1$, the objective value of ϕ is $\int \phi \geq \gamma^{-1}$. Therefore $\vartheta_s(G) \geq \gamma^{-1}$, and it follows that $\vartheta_s(G)\vartheta_d(G^c) \geq 1$ since ε was arbitrary. □

13.4 Extension to homogeneous spaces

In subsection 12.1.3, we saw that when G is a topological Schreier graph over the compact group Γ, its independence ratio $\tilde{\alpha}(G)$ is preserved after blowing up G to a Cayley graph $\widetilde{G}$ over Γ. The blow-up procedure involved replacing each vertex with an independent set. Except in trivial cases, this procedure *will* change the independence *number*. In order to upper bound $\alpha(G)$ within our framework, we therefore use a different kind of blow-up, which involves replacing each vertex with a clique rather than an independent set.

Let $G = (V, E)$ be a topological Schreier graph over the compact group Γ. Let $v_0 \in V$ be any point, and let H be the stabilizer subgroup of v_0 in Γ. We set

$$X = \{x \in \Gamma : \{v_0, xv_0\} \in E\} \cup (H - \{e\}),$$

and $\widetilde{G} = \mathrm{Cay}(\Gamma, X)$. (Compare with equation (12.4).) The next proposition shows that this kind of blow-up does not affect the independence number.

Proposition 13.7. $\alpha(\widetilde{G}) = \alpha(G)$*; this equation is also valid when one of the numbers is infinite.*

Proof. Let $q : \Gamma \to \Gamma/H$ be the canonical projection map and let $\Psi : \Gamma/H \to V$ be the homeomorphism given by [Fol95, Proposition 2.44]; so $\Psi \circ q(x) = xv_0$.

Let $I \subset V$ be an independent set in G, and for each $x \in I$, choose one $x' \in (\Psi \circ q)^{-1}(x)$; denote the set of all these choices x' by I'. Then I' is an independent set in $\widetilde{G}$, for if $x, y \in I'$ are distinct and $y^{-1}x \in X$, then $\{v_0, y^{-1}xv_0\} \in E$, and so $\{yv_0, xv_0\} \in E$. Moreoever $|I'| = |I|$ and, so $\alpha(\widetilde{G}) \geq \alpha(G)$.

Now suppose $J \subset \Gamma$ is an independent set in $\widetilde{G}$, and let $J' = \{xv_0 : x \in J\}$. Then $|J'| = |J|$, for if $xv_0 = yv_0$, then $y^{-1}x \in H$, which would mean that x and y were joined with an edge in $\widetilde{G}$. Moreover J' is independent since $\{xv_0, yv_0\} \in E$ implies $y^{-1}x \in X$. This shows $\alpha(G) \geq \alpha(\widetilde{G})$. □

13.5 Example: the Delsarte bound for spherical codes

Fix $n \geq 2$ and $A \subset [-1, 1)$. A *spherical A-code* (see e.g. [DGS77]) is defined as a subset $I \subset S^{n-1}$ such that $\langle \xi, \xi' \rangle \in A$ for all distinct $\xi, \xi' \in I$.

Let P denote the set of positive definite functions $f : [-1, 1] \to \mathbb{R}$ for S^{n-1}. Together with the Stone-Weierstrass theorem, [DGS77, Theorem 4.3] gives the following upper bound for the size of any spherical A-code when the elements of A are bounded away from 1.

$$\inf \{1 + f(1) : f \in P,\ f(t) \leq -1 \text{ for all } t \in A\} \tag{13.9}$$

The bound in (13.9) is known as the Delsarte bound for spherical codes.

Spherical A-codes are precisely the independent sets in the graph $G = (S^{n-1}, E)$, where $E = \{\{\xi, \xi'\} : \xi \neq \xi', \langle \xi, \xi' \rangle \notin A\}$. Since G is a topological Schreier graph over the group $SO(n)$, we can upper bound $\alpha(G) = \alpha(\widetilde{G})$ using our methods, where $\widetilde{G}$ is the clique blow-up graph described in Section 13.4. Provided A is bounded away from 1, the

Cayley graph $\widetilde{G}$ is dense. In this case, the upper bound of equation (13.9) is equal to $\vartheta'_d(\widetilde{G})$, and we now prove this.

Proposition 13.8. *$\vartheta'_d(\widetilde{G})$ is equal to the infimum* (13.9).

Proof. As in the hypotheses of Proposition 12.4, fix any $\xi_0 \in S^{n-1}$, let H be the stabilizer subgroup of ξ_0 in $SO(n)$ for the canonical action of $SO(n)$ on S^{n-1}, and let $\{x_t\}_{-1\leq t\leq 1} \subset SO(n)$ be a collection of points satisfying $\langle \xi_0, x_t\xi_0\rangle = t$ for each t, $-1 \leq t \leq 1$.

Let γ denote the infimum in (13.9). If $f : [-1, 1] \to \mathbb{R}$ is a feasible solution to program (13.9), define $\phi : SO(n) \to \mathbb{R}$ by $\phi(x) = f(\langle \xi_0, x\xi_0\rangle)$. By Proposition 12.4, ϕ is a feasible solution to program (13.4), and $\phi(e) = f(1)$. This shows $\vartheta'_d(\widetilde{G}) \leq \gamma$.

If $\phi \in \mathcal{P}$ is a feasible solution to program (13.4), define $f : [-1, 1] \to \mathbb{R}$ by $f(t) = \int_H \int_H \phi(hx_t h')\ dh\ dh'$. Then f is feasible for program (13.9) by Proposition 12.4, and $f(1) \leq \phi(e)$. Therefore $\gamma \leq \vartheta'_d(\widetilde{G})$, establishing the proposition. $\square$

Bibliography

[Ali07] C. D. Aliprantis. *Cones and Duality.* American Mathematical Soc., 2007.

[BDdOFV14] C. Bachoc, E. DeCorte, F. M. de Oliveira Filho, and F. Vallentin. Spectral bounds for the independence ratio and the chromatic number of an operator. *Israel J. Math.*, 202:227–254, 2014.

[BGSV12] C. Bachoc, D. C. Gijswijt, A. Schrijver, and F. Vallentin. Invariant semidefinite programs. In M. F. Anjos and J. B. Lasserre, editors, *Handbook on Semidefinite, Conic and Polynomial Optimization*, pages 219–269. Springer, 2012.

[BNdOFV09] C. Bachoc, G. Nebe, F. M. de Oliveira Filho, and F. Vallentin. Lower bounds for measurable chromatic numbers. *Geom. Funct. Anal.*, 19:645–661, 2009.

[BPT14] C. Bachoc, A. Passuello, and A. Thiery. The density of sets avoiding distance 1 in Euclidean space. `http://arxiv.org/abs/1401.6140`, 2014+.

[BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[CS93] J. H. Conway and N. J. A. Sloane. *Sphere packings, lattices and groups*, volume 290 of *Grundlehren Math. Wiss.* Springer-Verlag, New York, 3rd edition, 1993.

[DdLV14] E. DeCorte, D. de Laat, and F. Vallentin. Fourier analysis on finite groups and the Lovász theta-number of Cayley graphs. *Experiment. Math.*, 23:146–152, 2014.

[DE09] A. Deitmar and S. Echterhoff. *Principles of Harmonic Analysis*. Universitext. Springer, 2009.

[Del73] P. Delsarte. *An algebraic approach to the association schemes of coding theory*. PhD thesis, Université Catholique de Louvain, 1973.

[DGS77] P. Delsarte, J. M. Goethals, and J. J. Seidel. Spherical codes and designs. *Geometriae Dedicata*, 6(3):363–388, 1977.

[DKPS07] E. De Klerk, D. V. Pasechnik, and A. Schrijver. Reduction of symmetric semidefinite programs using the regular $*$-representation. *Mathematical Programming*, 109.2(3):613–624, 2007.

[dLV13] D. de Laat and F. Vallentin. A semidefinite programming hierarchy for packing problems in discrete geometry. arXiv:1311.3789, 2013.

[dOF09] F. M. de Oliveira Filho. *New bounds for geometric packing and coloring via harmonic analysis and optimization.* PhD thesis, CWI, Amsterdam, 2009.

[dOFV10] F. M. de Oliveira Filho and F. Vallentin. Fourier analysis, linear programming, and densities of distance avoiding sets in $\mathbb{R}^n$. *J. Eur. Math. Soc. (JEMS)*, 12:1417–1428, 2010.

[dOFV13] F. M. de Oliveira Filho and F. Vallentin. Computing upper bounds for the packing density of congruent copies of a convex body I. arXiv:1308.4893, 2013.

[DP15] E. DeCorte and O. Pikhurko. Spherical sets avoiding a prescribed set of angles. arXiv:1502.05030, 2015.

[DX13] F. Dai and Y. Xu. *Approximation theory and harmonic analysis on spheres and balls.* Springer Monographs in Mathematics. Springer, Berlin, 2013.

[Ell11] D. Ellis. Eigenvalue methods in extremal combinatorics: an overview. https://www.dpmms.cam.ac.uk/~dce27/evalmethods1.pdf, 2011.

[Fol95] G. B. Folland. *A Course in Abstract Harmonic Analysis.* CRC Press Inc., 1995.

[Fol99] G. B. Folland. *Real Analysis: Modern Techniques and Their Applications.* John Wiley & Sons Inc., 2nd edition, 1999.

[FW81] P. Frankl and R. M. Wilson. Intersection theorems with geometric consequences. *Combinatorica*, 1(4):357–368, 1981.

[GM07] B. Gärtner and J. Matoušek. *Understanding and using linear programming*, volume 168 of *Universitext*. Springer, Berlin, 2007.

[GM12] B. Gärtner and J. Matoušek. *Approximation algorithms and semidefinite programming.* Springer, 2012.

[HR94] E. Hewitt and K. A. Ross. *Abstract Harmonic Analysis*, volume 2. Springer Science & Business Media, 1994.

[Kal09] G. Kalai. How large can a spherical set without two orthogonal vectors be? Combinatorics and more (weblog), 2009.

[Knu94] D. E. Knuth. The sandwich theorem. *Electron. J. Comb.*, 1:#A1, 1994.

[Lev79] V. I. Levenshtein. On bounds for packings in n-dimensional Euclidean space. *Soviet Math. Doklady*, 20:417–421, 1979.

[Lov79] L. Lovász. On the shannon capacity of a graph. *Information Theory, IEEE Transactions*, 25.1:1–7, 1979.

[Lov86] L. Lovász. *An Algorithmic Theory of Numbers, Graphs, and Convexity.* CBMS Regional Conference Series in Applied Mathematics. SIAM, 1986.

[Lov03] L. Lovász. Semidefinite programs and combinatorial optimization. In B. A. Reed and C. L. Sales, editors, *Recent Advances in Algorithms and Combinatorics*, pages 137–194. Canadian Mathematical Society, 2003.

[Mon11] A. Montina. Communication cost of classically simulating a quantum channel with subsequent rank-1 projective measurement. *Physical Review A*, 84:060303, 2011.

[Mus08] O. R. Musin. The kissing number in four dimensions. *Annals of Mathematics*, pages 1–32, 2008.

[PZ04] F. Pfender and G. M. Ziegler. Kissing numbers, sphere packings, and some unexpected proofs. *Notices of the AMS*, 2004.

[Rai99] A. M. Raigorodskii. On a bound in Borsuk's problem. *Russian Mathematical Surveys*, 54(2):453–454, 1999.

[Rai00] A. M. Raigorodskii. On the chromatic numbers of spheres in $\mathbb{R}^n$. *Combinatorica*, 32(1):111–123, 2000.

[Rud87] W. Rudin. *Real and Complex Analysis.* Tata McGraw-Hill, 1987.

[Sch11] J. Schur. Bemerkungen zur Theorie der beschränkten Bilinearformen mit unendlich vielen Veränderlichen.

Journal für die reine und Angewandte Mathematik, 140:1–28, 1911.

[Sch79] A. Schrijver. A comparison of the Delsarte and Lovász bounds. *Information Theory, IEEE Transactions*, 25.4:425–429, 1979.

[Soi09] A. Soifer. *The Mathematical Coloring Book*. Springer, New York, 2009.

[Sze92] G. Szegö. *Orthogonal Polynomials*. American Mathematical Soc., 1992.

[Wit74] H. S. Witsenhausen. Spherical sets without orthogonal point pairs. *American Mathematical Monthly*, pages 1101–1102, 1974.

Index

Acknowledgements

The four years I spent in Delft as a PhD student were influenced by a number of people. The unique challenges that arise from living and working in a foreign country presented the opportunity for many to make a difference, and I would like to acknowledge some of them here.

My first thanks are owed to my advisor Frank Vallentin, who made this thesis possible, who supported me when I needed it, and who helped me to grow into an independent researcher. His kindness and generosity are unmatched. Karen Aardal has also been there the whole way, always ready with a joke and a smile, making sure the whole process stayed on course. A very special thanks is deserved by Oleg Pikhurko, who in addition to being a pleasure to work with, also taught me a lot about research, and about life in research; for me he served and continues to serve as an example when it comes to attacking hard problems, mathematical and otherwise.

There are a few more of my seniors I would like to mention. Markus deserves, in addition to my thanks, also a round of applause for being able to hold still while carefully reading the manuscript of this thesis, even when he didn't have to read it, and even when it probably didn't deserve to be read. Thanks to him, 100% of the errors in the first draft were corrected. The errors you find in the thesis now were introduced by me personally after he was done checking it over. Lex, Dion, Monique, and Marcel have been there at several key points during my PhD, inspiring, teaching, and supporting, and Jan Draisma was a big part of the reason my few months based in Eindhoven were so productive. Lastly I want to mention Fernando and Christine, who have been superb co-authors in the past, and who (I hope) will continue to be. Thank you all!

Much thanks is owed to my family, especially my parents, who influenced me in many ways, and who were always there with a helping hand whenever I needed it. Two others that deserve special mention in this section are Mahtab and Albaloo, whose presence in my life during the past four years was my main source of company and emotional support. This thesis would have been impossible without you. Thank you for being there.

Mathematically I got a lot out of the other PhD students around me, and I'd like to mention in particular David, Jan, and Nikita. We had a pretty awesome seminar together, and it would be hard to imagine a bunch of guys with more enthusiasm. From outside Delft, there was also Antonis, Guus, Jorn, and Miek. All told we had innumerable mathematical discussions and shared too many ideas to remember.

I'd like to personally thank some of the other students, past and present, at Delft and elsewhere, who helped create the positive atmosphere I benefitted from during my time as a PhD: (in no particular order) Sjoerd, Matthijs, Sonja, Pieter, Aida, Guido, Emil, Rob, Eddie, Shanfei, Chiara, Konstantin, Debanshu, Jeroen, Jonas, Shah, Seyran, and Özlem. And I'd like to mention Fred especially, mainly for being himself, but also for showing me around when I first arrived, and for always being ready to help with any administrative thing I got stuck with, which was basically every administrative thing.

Although almost all of my PhD years were spent being rather anti-social, I did on occasion have the good fortune to have some friends in my life outside of math. There was Rogier who courageously put up with me as a housemate, as well as Bing, Alex, Herman, Daniel, Milan, Cata, and Elif. All of you made a positve impact on me during my time in the Netherlands; it would not have been the same here without you.

My final thanks go to Liana, who encouraged me through the final stages of thesis writing, and with whom I hope to create many new memories.

Curriculum Vitae

Philip Evan Bigras DeCorte (Evan) was born on 13 September 1983 in Ottawa, Ontario, Canada. He attended secondary school at the Glebe Collegiate Institute in Ottawa, where he graduated with honours in 2001. After serving as a radio operator in the Canadian Forces (army), he began a Bachelor's degree in computing science in 2003 at Queen's University in Kingston, Ontario, Canada. He officially switched his major to mathematics in 2005, and graduated with honours in 2007, earning a B.Sc.H. in pure mathematics, with a computing science minor.

In 2007 Evan enrolled at the University of Toronto as a Master's student. He spent the 2008-2009 academic year visiting at the University of Amsterdam. He was awarded an M.Sc. in mathematics from the University of Toronto and completed the course requirements and comprehensive examinations for a Ph.D. in 2010, after which time he made the decision to move back to the Netherlands to begin Ph.D. research under the supervision of Frank Vallentin at the Delft University of Technology.

Besides writing this thesis, Evan has had a number of opportunities to grow as a teacher and researcher during the past four years. In particular, he has served as a teaching assistant and has himself taught undergraduate mathematics courses, both in English and in Dutch. During his third year he organized a reading group about abstract harmonic analysis for graduate students in Delft which continued into the beginning of his fourth year. He spent the summer after third year on a two month research visit with Oleg Pikhurko and Anusch Taraz and the Technical University of Munich, and made another visit to Oleg in November 2013. He has participated in a number of regularly occurring seminars and working groups, including the CWI algebra and combinatorics working group, the Leiden functional analysis seminar, and the Eindhoven discrete mathematics seminar. During the entire four years, he has been given the chance to participate in and present his research at numerous conference and workshops in the Netherlands, in other parts of Europe, and in North America.

Publication List

C. Bachoc, E. DeCorte, F. M. de Oliveira Filho, and F. Vallentin. Spectral bounds for the independence ratio and the chromatic number of an operator. *Israel J. Math.*, 202:227-254, 2014.

E. DeCorte, D. de Laat, and F. Vallentin. Fourier analysis on finite groups and the Lovász theta-number of Cayley graphs. *Experiment. Math.*, 23:146-152, 2014.

E. DeCorte and O. Pikhurko. Spherical sets avoiding a prescribed set of angles. `arXiv:1502.05030`, 2015.